

The Sharp Contour of Consciousness

By

Jonathan Aaron Simon

A dissertation submitted in partial fulfillment

Of the requirements for the degree of

Doctor of Philosophy

Department of Philosophy

New York University

May 2012

Paul A. Boghossian

© Jonathan Aaron Simon
All Rights Reserved, 2012

DEDICATION

To my mother, Cynthia Weisfeld Simon, my grandmother, Rhoda Weitzenfeld Simon, my grandfather Israel Simon, and my uncle Richard Keller Simon.

To Burton Weisfeld and Sylvia Weisfeld, and to William Simon and Cyndy Esty and Benjamin Simon.

To Heather Alexander.

ACKNOWLEDGEMENTS

Special thanks to the chair of my committee, Paul Boghossian, and to my advisors Ned Block, David Chalmers, and Sharon Street. Many thanks for time, insight and encouragement to Eli Alshanetsky, Michael Antony, David Barnett, Dean Chapman, Justin Clarke-Doane, Johnny Cottrell, Will Davies, Paul Egre, Kit Fine, Akiko Frischhut, Yu Guo, Grace Helton, Robert Howell, Hasen Khudairi, Geoffrey Lee, Dan Lopez De Sa, Phillippe Lusson, Colin Marshall, Farid Masrour, John Morrison, Mark Okrent, Joelle Proust, Michael Raven, Stephen Schiffer, Michael Schweiger, Miguel Sebastian, Jeff Sebo, Ted Sider, Michael Strevens, Peter Unger, Achille Varzi, Christopher Viger, Robbie Williams, Jarred Warren, Sebastian Watzl, Charlotte Witt and participants of seminars at NYU, ENS and Geneva.

Thanks to my parents, grandparents, brother and extended family for supporting me.

And finally, thanks to Heather Alexander for the love she has given me.

Table of Contents

DEDICATION	IV
ACKNOWLEDGEMENTS	V
INTRODUCTION	1
INTRODUCTION TO EACH CHAPTER.	3
<i>CHAPTER ONE: MICHAEL ANTONY ON THE VAGUENESS OF 'CONSCIOUSNESS'</i>	<i>3</i>
<i>CHAPTER TWO: GENERALIZED CONCEPTUAL GAP PRINCIPLES AND VAGUENESS</i>	<i>3</i>
<i>CHAPTER THREE: A PRIORI PROTOPHENOMENALISM AND NON-EXCLUDABILITY</i>	<i>5</i>
<i>CHAPTER FOUR: THE SIGNIFICANCE OF THE MENTAL PROBLEM OF THE MANY..</i>	<i>7</i>
<i>CHAPTER FIVE: THE SIGNIFICANCE OF CONSCIOUSNESS</i>	<i>8</i>
CONCLUSION	9
NOTE ON NOTATION.....	10
CHAPTER ONE: ANTONY ON THE VAGUENESS OF 'CONSCIOUSNESS'	11
§1.1 ANTONY'S CONCEPTUAL ARGUMENT.....	12
§1.1.1 <i>THE NECESSARY CONDITION</i>	<i>12</i>
§1.1.2 <i>'PHENOMENALLY CONSCIOUS' DOES NOT SATISFY MIDDLE GROUND.....</i>	<i>21</i>
§1.2 PROBLEMS WITH ANTONY'S CONCEPTUAL ARGUMENT.....	27
§1.3 ANTONY'S METAPHYSICAL CONCLUSIONS	37
§1.4 PROBLEMS WITH ANTONY'S METAPHYSICAL CONCLUSIONS.....	38
CHAPTER TWO:.....	43
GENERALIZED CONCEPTUAL GAP PRINCIPLES AND VAGUENESS.....	43
§2.1: GENERALIZED GAP PRINCIPLES.....	46

§2.2: THE INTEREST OF NON-EXCLUDABILITY, CANONICAL PRIMITIVITY AND CANONICAL INSCRUTABILITY.....	54
§2.3: GENERALIZED CONCEPTUAL GAPS AND THE STATUS OF PHENOMENAL CONCEPTS.....	59
§2.4 NON-EXCLUDABLE CONCEPTS ARE NOT VAGUE	62
CHAPTER THREE:	73
<i>A PRIORI</i> PROTOPHENOMENALISM AND NON-EXCLUDABILITY	73
§3.1: PROTOPHENOMENALISM AND A PRIORI PROTOPHENOMENALISM.....	76
§3.2: AGAINST <i>A PRIORI</i> PROTOPHENOMENALISM	84
§3.3: IF <i>A PRIORI</i> PROTOPHENOMENALISM IS FALSE, THE GENERALIZED NOTIONS OF CONCEPTUAL GAP APPLY TO PHENOMENAL CONCEPTS	102
CHAPTER FOUR:	112
THE SIGNIFICANCE OF THE MENTAL PROBLEM OF THE MANY.....	112
§4.1 THE MENTAL PROBLEM OF THE MANY.....	113
§4.2: FROM THE MENTAL PROBLEM OF THE MANY TO PROPERTY DUALISM.....	119
§4.2.1: <i>ONTOLOGICALLY CONSERVATIVE RESPONSES THAT DO NOT APPEAL TO VAGUENESS</i>	121
§4.2.2: <i>ONTOLOGICALLY CONSERVATIVE RESPONSES THAT APPEAL TO VAGUENESS</i>	123
§4.2.3: <i>ONTOLOGICALLY INFLATIONARY RESPONSES</i>	125
§4.3: CHALLENGING THE PREMISES.....	130
CHAPTER FIVE:	134

THE SIGNIFICANCE OF CONSCIOUSNESS.....	134
§5.1: NATURAL DIMENSION DETERMINABLES.....	135
§5.2: SIGNIFICANCE AND THE SEA CHANGE ARGUMENT	141
5.3: NORMATIVE SIGNIFICANCE.....	147
§5.4: IF PHENOMENAL CONSCIOUSNESS IS A NATURAL DIMENSION (OR SUB- DIMENSION) DETERMINABLE THEN PROPERTY DUALISM OR PANPSYCHISM IS TRUE.....	151
§5.5: THE NATURAL DIMENSION ARGUMENT VERSUS THE CONCEIVABILITY ARGUMENT	158
BIBLIOGRAPHY	163

INTRODUCTION

We have the intuition that experience - the qualitative character of the experience of color, shape, sound, pain, emotion ; *what it is like* to experience these things - is an on or off phenomenon. Though experience can be more or less intense, more or less rich, you are either experiencing (at some moment) or you are not. Asking whether someone is experiencing is like asking whether the light is on, rather than like asking whether the light is bright. Something can be more or less bright, but it can also be in a state indeterminate between being bright and not being bright. This cannot happen with subjective experience - you are either having some or you are not.

We also take there to be a significant contrast between experiential reality and non-experiential reality. Whether the light emits *exactly* 1542 lumens is a yes or no affair: it either is emitting exactly that amount or it is not. But there is no significant contrast here: there is nothing ontologically distinguished about emitting 1542 lumens rather than 1543, the way there is something ontologically distinguished about the light being on rather than off. Asking whether someone is experiencing is like asking whether the light is on, rather than like asking whether the light is emitting exactly 1542 lumens. There is a significant metaphysical contrast between experience and non-experience.

Putting these two observations together, we have that consciousness is a sharp contour in the world. But is there room for such a sharp contour? The material world is a gradual place. With a powerful enough microscope, everything looks cloudy, and there are only sharp contours in nature very rarely – some at the quantum level, some dealing with thresholds of fundamental phenomena like gravity. Most things that appear to be sharp and distinguished changes (like the firing of a neuron, or the splitting of an embryonic cell) are actually quite gradual on a fine enough scale.

So consciousness seems to be a sharp contour in the world, but material reality is too gradual to afford us any truly sharp contours (in anything like the appropriate places). This is a problem if we hope to find a place for consciousness in the material world. One strategy for solving the problem is to hold that ‘Phenomenally Conscious’ actually is vague – that the concept is really more like ‘Brightly Illuminated’ than like ‘Illuminated’ – even though it takes some work, and perhaps empirical discovery, to see how. Then we can say that consciousness is a contour, but deny that it is a sharp one. After a first chapter of literature review, in chapters two and three of this dissertation I argue that this strategy is not open to materialists: materialists must say that the concept ‘Phenomenally Conscious’ is not vague. In chapters four and five, I argue that we cannot dispense with the thought that consciousness is a contour, and I argue that this is genuinely inconsistent with

taking consciousness to be a part of ordinary, material reality. If consciousness is a sharp contour, we must either take the concept 'Phenomenally Conscious' to describe something that goes beyond ordinary material reality, or take it to describe far more of ordinary material reality than we might have supposed.

--

INTRODUCTION TO EACH CHAPTER.

CHAPTER ONE: MICHAEL ANTONY ON THE VAGUENESS OF 'CONSCIOUSNESS'

My work owes an intellectual debt to Michael Antony. In a pair of papers published in 2006, Antony develops an argument that 'Phenomenally Conscious' is not vague, and that consequently none of the leading materialist theories of consciousness may be correct. In this first chapter I discuss some pressing problems with Antony's argument and outline how I intend to overcome these problems, building on the core of his argument, in later chapters.

CHAPTER TWO: GENERALIZED CONCEPTUAL GAP PRINCIPLES AND VAGUENESS

In this chapter I consider several ways of generalizing the idea of a conceptual gap. Many materialists agree with dualists that there is a conceptual gap between phenomenal concepts and material concepts: zombies are conceivable, inverted spectra are conceivable, souls in worlds without much complexity are conceivable. As we usually understand it, the idea of a conceptual gap is the idea of a relation between two families of concepts – for example the family of phenomenal concepts and the family of material concepts. Here I consider three ways of generalizing from the relational to the absolute. Two of these, *Canonical Inscrutability* and *Canonical Primitivity*, are ways for a family of concepts to be autonomous from all other families of concepts, and one of them, *Non-Excludability*, is a way for a single concept to mark a kind of absolute autonomy in the realm of concepts. After I introduce these new distinctions I discuss their relevance to existing debates: they articulate distinct notions of conceptual primitivity, and they allow us to reframe the issues in certain debates like the debate over the autonomy of ethics. They also allow us to articulate respects in which phenomenal concepts may be distinguished even if we hold that the ordinary relational sort of conceptual gap does not distinguish them. I conclude the paper by showing one further application of the notion of Non-Excludability: I argue that if a concept is Non-Excludable then it is not vague (assuming that vagueness is, broadly speaking, a semantic phenomenon).

CHAPTER THREE: *A PRIORI* PROTOPHENOMENALISM AND NON-EXCLUDABILITY

It is commonly supposed that Democritus was right, and the rest of the Presocratic and Romantic philosophers were wrong: reality is made of Mere Matter and Void, rather than Sacred Fire, or Aether, or Will, or some other exotic thing. We suppose that the grand debate over the nature of being is closed, and materialists have won. But there is room in logical space for a Dual Aspect view according to which 'Sacred Fire' and 'Matter' are different guises of the same being. Many materialists embrace the Dual Aspect view that 'Mind' and 'Matter' are different guises of the same being. And this sort of materialist may have reason to embrace the more exotic sort of Dual Aspect theory. The view that there is a 'Sacred Fire' aspect to material reality might help to explain the connection between the non-mental and the mental – if the correlation between the 'Sacred Fire' aspect and the 'Mind' aspect is strong enough.

In this chapter I will describe a position, *Protophenomenalism*, which articulates a precondition on this exotic sort of Dual Aspect theory, and I will tentatively defend it. But my primary aim in this chapter is to show that there is a limit on how much of an explanation of the mental we can expect this

strategy to deliver. I will call the position which says that some description in exotic terms might actually *entail* a phenomenal description *A Priori Protophenomenalism*. I will argue that this view is difficult to defend in conjunction with a non-deflationary, naturalistic theory of concept application. Since that is the sort of theory of Concept Application you should embrace if you think that the mental is ultimately material, this means that if you think the mental is ultimately material, you should abandon the hope for an exotic conception of reality that entails anything about mental reality. This imposes a significant limit on the ways that an exotic Dual Aspect theory might help the materialist bridge the explanatory gap.

My secondary aim in this chapter is to use this result to show that phenomenal concepts are examples of the sorts of Generalized Conceptual Gaps that I discuss in chapter two. I will argue that if *A Priori Protophenomenalism* is false, then the concept 'Phenomenally Conscious' is *non-excludable*. I will also present some reasons to think that the family of phenomenal concepts is canonically inscrutable and canonically primitive, though for reasons I will discuss these claims are harder to establish. But *non-excludability* is the more newsworthy notion: In the second chapter I argue that if a concept is non-excludable then it is not vague, so the implication of my conclusion here will be that 'Phenomenally Conscious' is not vague (assuming a non-deflationary, naturalistic theory of concept application). And in the

fourth and fifth chapters I argue that if 'Phenomenally Conscious' is not vague then either Panpsychism or Property Dualism is true. There is thus a master argument here against Orthodox Materialism broadly construed: either a non-deflationary, naturalistic theory of concept application is false, or one of Panpsychism or Property Dualism is true.¹ I do defend the claim that Protophenomenalism is open to orthodox materialists (though *A Priori* Protophenomenalism is not), but this defense is modulo my argument in chapter four and five that Orthodox Materialism is untenable.

CHAPTER FOUR: THE SIGNIFICANCE OF THE MENTAL PROBLEM OF THE MANY

In this chapter I offer an argument that the property **phenomenal consciousness** is not identical to or grounded in any material property – in other words, for Property Dualism (characterized in a non-modal way²). My argument is an argument from the Mental Problem of the Many to Property

¹If one rejects a non-deflationary naturalist theory of concept application, I take the most compelling alternative to be a primitivist theory of concept application, which is a sort of primitivism about intentional states and relations. Since it is plausible that all mental states, including phenomenal states, are intentional states, this might amount to property dualism about phenomenal properties as well, but in any case it amounts to a property dualism about some mental properties.

²As I define it Property Dualism is compatible with a necessary supervenience claim (although the explanatory role of such a claim would be questionable). I take Substance Dualism to be compatible with (and probably to entail) Property Dualism. I do not think my own arguments establish Substance Dualism but they certainly do not rule it out. Thanks to Mike Raven for discussion on this point.

Dualism. There already are arguments from the Mental Problem of the Many to Dualism (Unger 2006, Zimmerman 2011). But these arguments are arguments for Substance Dualism, and they presuppose Property Dualism (Unger implicitly, Zimmerman explicitly). There is an argument from the Mental Problem of the Many to Property Dualism, but it calls for different premises than the existing arguments employ. My argument hinges on three premises. The first premise, which I defend in chapters two and three, is that ‘Phenomenally Conscious’ is not vague. The second is that there are not swarms of overlapping conscious beings wherever there is one. The third premise specifies a sense in which **phenomenal consciousness** is a sharp contour in the world: it is ontologically *significant*.

CHAPTER FIVE: THE SIGNIFICANCE OF CONSCIOUSNESS

In this chapter I further develop the idea that **phenomenal consciousness** is ontologically significant – that there is a contour in the world between conscious things and non-conscious things. Drawing on the premise that ‘Phenomenally Conscious’ is not vague, I argue that the property **phenomenal consciousness** is a Natural Dimension Determinable – a property demarcating an objective dimension of similarity and difference in quality space. I then argue that if **phenomenal consciousness** is a Natural

Dimension Determinable then either Property Dualism or Panpsychism is true (or both).

CONCLUSION

Chapters two through five of this dissertation amount to a master argument for the disjunction of Property Dualism, Panpsychism, or Primitivism about Intentionality. In chapter two I argue that if a concept is non-excludable then it is non-vague. In chapter three I argue that if *A Priori* Protophenomenalism is false, then 'Phenomenally Conscious' is non-excludable, and I argue that *A Priori* Protophenomenalism is false. This entails that 'Phenomenally Conscious' is non-vague. The argument hinges on the assumption of non-deflationary materialism about concept application. The non-deflationism in question is very mild, meaning that the most plausible way of rejecting the assumption is to reject materialism rather than non-deflationism. This rejection amounts to something like Primitivism about some intentional relations. Thus we face a choice between Primitivism about some intentional relations and the non-vagueness of consciousness. But in chapter four I use the non-vagueness of consciousness as a premise, along with the premise that consciousness is significant, in an argument that Property Dualism is true, and in chapter five I use it as a premise – again, along

with the premise that consciousness is significant – in an argument that either Property Dualism or Panpsychism is true. Thus, the master claim of my dissertation is that, if we grant that consciousness is significant, we must choose between Primitivism about intentional relations, Property Dualism or Panpsychism.

NOTE ON NOTATION

In this dissertation I will co-opt ordinary language to name particular theories or views, and I will mention a number of concepts and properties. I will refer to theories, views and principles with capital letters. I will refer to concepts with capital letters in single quotes, and I will refer to properties in lowercase bold. For example in chapter two I will discuss views like Materialism and Dualism, and I will discuss concepts like 'Phenomenally Conscious' and 'Good', and I will discuss properties like **phenomenal consciousness** and **goodness**.

CHAPTER ONE: ANTONY ON THE VAGUENESS OF 'CONSCIOUSNESS'

In two groundbreaking papers, Antony 2006a and 2006b, Michael Antony argues that the concept 'Phenomenally Conscious' is not vague,³ and that if 'Phenomenally Conscious' is not vague then most ordinary versions of materialism are false. In this chapter, I review Antony's argument, and isolate what I take to be its most pressing difficulties. There are serious problems with his case that 'Phenomenally Conscious' is not vague, and with his case that the non-vagueness of 'Phenomenally Conscious' implies that most ordinary versions of materialism are false. I do believe that these problems are solvable, and that there is a sound argument for Antony's conclusions, but this argument looks substantially different from Antony's. The subsequent chapters of this dissertation present that argument in full, and I will present it in outline here, as the critique of Antony makes it salient.

In §1.1 I will present Antony's argument that 'Phenomenally Conscious' is not vague. In §1.2 I will articulate what seem to me to be the most pressing problems with the argument, and indicate how I think these problems are best

³Antony's argument is for the concept 'Conscious State'. He spends a few pages explaining that by 'Conscious' he means 'Phenomenally Conscious'. And though he speaks of the concept 'Conscious State', he makes clear that he takes the argument also to work for 'Having a Conscious State'. But there are reasons to think the argument will work for the latter and not the former (we might say that it is indeterminate whether an unattended pain is a conscious state without saying it is indeterminate whether the experiencer having the unattended pain has conscious states). So for reasons of both substance and presentation I will speak as though his subject were the concept 'Phenomenally Conscious'.

dealt with. In §1.3 I will present Antony's argument that if 'Phenomenally Conscious' is not vague then most ordinary versions of materialism are false, and in §1.4 I will discuss problems with this argument and explain what I take to be the way forward.

§1.1 ANTONY'S CONCEPTUAL ARGUMENT

Antony claims that the concept 'Phenomenally Conscious' is not vague. His argument has two components. He first presents a condition which he argues is a necessary condition for a concept to be vague, and he then argues that the concept 'Phenomenally Conscious' does not satisfy this condition. I will present these components in turn.

§1.1.1 THE NECESSARY CONDITION

Antony does not name the condition. It is a complex condition with four components, and he names these C1 – C4. I hereby name the complex condition *Middle Ground*, because the core of the idea is that if a concept is

vague you have a conception of a middle ground between clear cases and clear non-cases.

Antony distinguishes between concepts in the psychological sense, concepts in the philosophical sense, and conceptions. A concept in the philosophical sense is “roughly, an abstract meaning, sense, etc” (2006b p.4), while a concept in the psychological sense is “roughly, a mental representation that can be a constituent of thoughts”(ibid.). A conception is “a mental representation structure that is semantically and syntactically *complex* – structures by means of which we identify, categorize and often simply think about objects, events, properties, etc” (ibid.). Antony distinguishes between philosophical and psychological concepts to make clear that his interest is with psychological concepts,⁴ and he distinguishes between psychological concepts and conceptions to remain neutral on the Fodorian debate over which concepts are complex and which are atomic. A classical theorist may identify conceptions with psychological concepts, while a Fodorian may hold that

⁴However, in the first footnote of his 2006a (which he wrote later but published sooner), Antony indicates that he takes mental representations to express abstract concepts, so he presumably draws these distinctions for the sake of clarity rather than because he takes something critical to hinge on the distinction.

conceptions are distinct from concepts but associated with them as a matter of psychological law.⁵

Antony tells us that his condition is a necessary condition on a thinker, S, possessing the capacity to conceive of borderline cases for some vague concept F. Here is the condition:

- “(C1) It must be psychologically possible for S to represent a series of individuals with individual conceptions $[\psi_1] - [\psi_n]$, such that:
- (C2) The series of individuals contains all of and only Fs, borderline Fs (i.e., individuals with respect to which v-dispositions are manifested), and not-Fs, and in that order.
- (C3) Individuals’ parts, properties, etc. are represented as gradually changing from individual to individual, either along some dimension, or with respect to (something like) their weighted sum, and
- (C4) There are at least some parts, properties, etc. represented in S’s conception [F], that the Fs, borderline Fs, and not-Fs are all represented by the $[\psi]$ s as clearly having.” (2006b p.9)

Take the concept ‘Tall’. If S is competent with ‘Tall’, then S can form conceptions of a series of individuals, each slightly less tall than the last, with the first being a clear case of ‘Tall’, the last being a clear case of ‘Not Tall’ and some intermediate ones triggering S’s v-dispositions (I say what these are just below). This satisfies C1-C3. For C4, we observe that each of these individuals

⁵Antony attributes the concept/conception distinction to Woodfield 1991, but see also Katz 1972, Davis 2003 and Higginbotham 1998. The distinction is particularly relevant if one is sympathetic with the conceptual atomism defended by, e.g. Fodor 1998 or Millikan 2000.

– the tall, the borderline, and the not-tall – has a height, and this means that ‘Height’ is the common conception that C4 demands.

By an ‘individual conception’ Antony means a conception (a complex mental representation) of an individual. Antony does not tell us exactly what he takes a borderline case to be – he says he wishes to remain neutral between competing theories and so will leave the matter intuitive – but he takes there to be a strong connection between a case of some concept F being a borderline case, and the community of competent users of F displaying a certain sort of disposition (a vagueness related disposition or *v-disposition*) in evaluating whether the case is an F. These dispositions will involve:

“...increased hesitation about whether to judge individuals as F or not-F, denying both that an individual is F and that it is not-F, changing one’s mind as to whether it is F or not-F, judging it to be neither clearly F nor clearly not-F (i.e., a borderline case), and so on... In addition, for those competent with a vague concept there is a characteristic phenomenology associated with conscious reflection on cases in and around the borderline region—a phenomenology distinct from that associated with certain kinds of ignorance about sharp boundaries, for example.” (2006b p.3)

Antony’s argument for C1 and C2 is that

“if S possesses a concept F, and F is vague, then S must be able to conceive instances falling under F, instances not falling under F, and borderline Fs. Since S can select and order what S conceives, C1 and C2 follow.” (2006b p.9) ⁶

In C3, Antony speaks of properties being represented as “gradually changing ... either along some dimension ... or with respect to their weighted

⁶For more on the psychological approach to vagueness see Wright 2001, Schiffer 2003.

sum." What does this mean? Antony does not really explain what it is to represent a change as gradual along some dimension. He does not say what a dimension is, though he suggests that a dimension is just a conception (2006b. p 10). As I will explain below, it is implicit in some of his arguments that he expects the variation in the relevant dimension to explain or ground the change from clear F to borderline F to non-F, though he does not make this explicit.

The distinction Antony draws between cases of gradual change along some dimension and gradual change with respect to some weighted sum is Alston 1964's distinction between Degree and Combinatory vagueness. According to Antony (and Alston), degree vagueness is vagueness where by a series of incremental changes in some dimension (for example height, or color) you can get from clear cases to borderline cases to clear non-cases. Combinatorial vagueness is vagueness in a cluster concept when it is vague just which conditions are a part of the cluster, or just how many of them something has to have in order to satisfy the concept. One example is 'Religion': this is a cluster concept with elements like 'Belief in a Deity', 'Belief in an Afterlife'. Borderline cases like 'Scientology' do not differ along some dimension (of metaphysically objective similarity) from clearer cases of religion: they differ over how many items in the cluster they satisfy (2006b p.8). Antony does not offer much of an argument that the division into degree

and combinatory types of vagueness is an exhaustive classification of the varieties of vagueness, but provided that we are relatively permissive in which concepts count as cluster concepts, and which changes are incremental changes in some dimension, it is hard to see what might count as a counterexample. Still, Antony may have to make these things more precise. For example, he must say why the series 'Phenomenally Conscious', 'Borderline Phenomenally Conscious', 'Not Phenomenally Conscious' does not specify suitably incremental changes in some dimension. These terms may be too general to count as individual conceptions, but the question is whether they may be elements of the relevant individual conceptions.

Antony's argument for C3 is that in conceiving a borderline F, one manifests v-dispositions

"... partly because the individual's properties, etc. come close to satisfying the requirements of both F-hood and not-F-hood. But that suggests an ability to imagine gradual changes to the individual's parts, properties, etc. (along a dimension or in terms of a weighted sum) sufficient to transform the individual into an F or a not-F." (2006b p. 9)

Note that this defense appeals to the idea that the relevant gradual variation along a dimension or in terms of a weighted sum explains or grounds the transition from clear F to borderline F to clear non-F. This constraint does not occur explicitly in C3. While it is clear that some sort of change must explain or ground the transition from F to borderline F to non-F, what is at

issue is whether it has to be gradual variation *along a dimension or in terms of a weighted sum*.

Antony concedes that C4 is the most contestable of the conditions of Middle Ground, and he only offers a defense of C4 for cases of degree vagueness. In those cases, he says that the common conception C4 asks for is just the conception of the dimension along which the individual conceptions gradually vary (2006b p. 10). In the case of ‘Tall’ this dimension is ‘Height’. A first temptation would be to hold that C4 is automatically satisfied if C3 is – or for that matter, even if it isn’t. There will always be general features common to any list of potential Fs, borderline Fs and non-Fs, for example, the feature ‘Exists’. But it is implicit in C4 that Antony wants the common conception to specify the parameter whose variation explains or grounds the transition from clear Fs to borderline Fs to clear non-Fs. As I have already indicated, this is not explicit – he does not even say that the common conception must be a parameter that varies in any way between the Fs, borderline Fs, and non-Fs, let alone in an explanatory way. But it emerges in his claim that for the cases of degree vagueness (where C4 is defensible) the common conception just is the specification of the dimension of variation mentioned in C3, and it emerges later in his argument that ‘Phenomenally Conscious’ does not satisfy C4.

The constraint C4 amounts to a strengthening of C3 – imposing some structure on what such a Dimension may be. For example, we might imagine a

series of things, some of which have spatial location and others do not, varying over the dimension of spatial volume – some things have spatial volume, while others have none at all. But ‘Spatial Volume’, though it might specify a Dimension of variation for such things, would not specify a common conception for them in the sense of C4 if some of those things do not have any spatial volume. Also, C4 denies that there might be dimensions of variation for *some* sequences of Fs, borderline Fs, and non-Fs that do not specify elements common to the conception of *all possible* Fs. The vagueness of ‘Red’ may hinge on what happens along the border between cardinal red and orange. There may be some feature common to both cardinal red things, and reddish orange things, but not common to crimson red things, that explains the variation from clear red to borderline red to not red. But Antony seems to be committed to saying that in such cases there must be some other conception of a dimension of variation that is common to all red things, as well as some borderline red things and some orange things.

As I will return to later, this is asking for a lot. As Brogaard 2010 points out, even with ‘Bald’ it is not obvious that there is any such common conception.⁷ Notions like ‘Quantity of Hair’ or ‘Has a Scalp’ may handle most cases, but what about birds, who may be bald even though they have neither

⁷Cf Brogaard 2010. Brogaard argues that Middle Ground cannot be a necessary condition on conceptual vagueness for precisely this reason.

scalps nor hair? Matters get worse when we consider trickier cases like ‘Cool’ or ‘Beautiful’ or ‘Funny’. While it may be possible to isolate *some* condition that all Beautiful things have in common with some borderline beautiful things and some non-beautiful things (e.g. ‘Being Perceptible or Intellectually Apprehensible’) it is hard to see that any such condition might explain the transition from beautiful to borderline beautiful to non-beautiful.⁸ In any case, Antony owes us an argument that conceptions simultaneously satisfying all of these constraints are available for any vague concept. A weakening of the demands of C4 might make Middle Ground more defensible, but as I will explain below, Antony’s argument that ‘Phenomenally Conscious’ does not meet the conditions of Middle Ground hinges on this strong reading of C4.

The overall structure of Middle Ground needs some clarification. Antony tells us that Middle Ground is a necessary condition on a thinker possessing the capacity to conceive of borderline cases for some concept ‘F’. But as Antony has stated it, C2 implies that someone competent with the concept F can form a series of individual conceptions, some of which are conceptions of borderline cases of F – conceptions that trigger v-dispositions in competent users. C2 demands the very thing for which C1-C4 are supposed to articulate a necessary condition. The circularity is not vicious, but it is confusing, since one usually thinks of independently stated conditions as being

⁸Thanks to Dan Lopez de Sa and Miguel Sebastian for discussion here.

independently satisfiable. But here, we must remember that part of Antony's claim is that C1-C2 cannot be satisfied for some concept F unless it is possible to also satisfy C3 for that concept, and also C4 if its vagueness is vagueness of degree.

Though I have serious doubts about the specifics, I take Middle Ground to articulate a compelling intuition. If a concept is vague then competent users of that concept must potentially be sensitive to that vagueness, and this sensitivity should manifest in some kind of conception of a middle ground – a border region – between clear cases of the concept and clear cases of its negation.⁹

§1.1.2 'PHENOMENALLY CONSCIOUS' DOES NOT SATISFY MIDDLE GROUND

Having claimed that C4 articulates a necessary condition on (degree) vagueness, Antony goes on to argue that 'Phenomenally Conscious' does not meet the condition C4, and also that 'Phenomenally Conscious' does not

⁹Antony claims that this condition is a necessary condition on vagueness regardless of whether vagueness is semantic, ontological or epistemic. It may seem that the condition is far more plausible on a semantic (or psychological) theory than on an ontological or epistemic one. But Ontological and Epistemic theorists tend to agree that vagueness manifests in the realm of concepts, even if its nature is epistemic or ontological. So this claim of Antony's is not necessarily indefensible.

exhibit combinatory vagueness. He does not offer an argument for his assumption that degree and combinatory vagueness exhaust the possible types of vagueness, but if that assumption is correct then these claims suffice for his conclusion.

He first clarifies what he means by 'Phenomenally Conscious'. He means qualitative, subjective, *what it is like*, what-the-phenomenal-zombie-lacks consciousness, rather than access consciousness, or self-consciousness, or attentional consciousness. He takes our concept of this sort of consciousness to be neutral with respect to whether materialism, dualism or idealism is true (2006b p.6). He leaves this notion of neutrality intuitive, but it later becomes clear that he takes it to imply that no representation of physical or functional properties may be part of the conception we have of 'Phenomenally Conscious' (2006b p.11).

Antony argues that there is no conception satisfying C4 for 'Phenomenally Conscious'.¹⁰ Antony proceeds by elimination. Such a conception must either entail consciousness, not entail consciousness, or be borderline whether or not it entails consciousness. Since by C4 the conception

¹⁰ In light of the question about the circularity in the formulation of Middle Ground that I discuss above, we might ask what Antony has in mind. If 'Phenomenally Conscious' cannot satisfy C4, but it can satisfy C1-C3 (or just C1-C2), then since that would mean that it is still possible to conceive of Borderline Cases, would that not be enough to ensure its vagueness (especially if, like Antony, you define vagueness in terms of borderline cases)? The answer again is that Antony takes it that there can be no borderline cases to conceive of unless C4 is satisfied.

must be a conception that is a component of at least one individual conception of a clearly non-conscious thing, it cannot entail consciousness. But it cannot be a borderline consciousness entailing element, Antony says, because this would mean that there are some clearly conscious states that we represent as involving borderline conscious states, or conscious states that are borderline realized. And Antony takes this to be clearly false. But why might there not be some fairly abstract conception that on its own neither clearly entailed consciousness nor non-consciousness, but in conjunction with some other conceptions entailed consciousness, and with yet other conceptions entailed non-consciousness? Again, because Antony is leaning on the (unarticulated) assumption that the conception not only specifies *some* element common to Fs, borderline Fs and not Fs, but it specifies the very parameter whose variation explains or grounds the variation from F to borderline F to not-F. There might be some parameter that on its own, no matter how we vary it, entails borderline consciousness, but then variation in that parameter cannot be what gets us from borderline consciousness to clear consciousness.

Most of the action is with the argument that the conception cannot be one that does not entail consciousness (or borderline consciousness or non-consciousness). This after all seems to be the most natural thought: 'Phenomenally Conscious' is vague if there is some parameter such that variation in it takes us from clear cases of non-consciousness, through

borderline cases of consciousness, to clear cases of consciousness. Antony does not claim to have a conclusive argument against this possibility. His more provisional argumentative strategy is to consider what seem to him to be the most natural candidates, and show that they do not, in fact, play the role that C4 requires. Antony first considers material conceptions – conceptions of physical or functional properties. These are ruled out, Antony claims, because ‘Phenomenally Conscious’ is a neutral concept (in the special sense of being a concept neutral on whether consciousness is material). Here, Antony is again leaning on the assumption that the common condition must specify that in virtue of which things vary from F to borderline F to not F. Otherwise some abstract functional level of description might suffice – the sort of description of the interrelationships of mental states that even a Cartesian would countenance. He is also leaning on the assumption that the common conception must be common to the general conception of F, rather than simply being common to at least one series of individual conceptions involving some Fs, some borderline Fs and some non-Fs. If the concept ‘Phenomenally Conscious’ is neutral in Antony’s sense, then no material concept is a component of that concept (or, as Antony might put it, no material conception is a component of the conception associated with that concept). But it does not follow that no material concept may be a component in each of a series of individual conceptions moving from conscious ones to non-conscious ones.

For example, consider the conception ‘State which is a conscious state iff its realizer oscillates at roughly 40 Hz’. This conception may not be a part of the neutral conception of ‘Phenomenally Conscious’ but it may figure in each of a series of individual conceptions, and indeed we seem to get an individual conception of a borderline case of consciousness by combining this conception with a conception of a borderline case of 40 Hz oscillation. Antony may be able to get what he needs here from the assumption that the parameter must explain and ground the transition from F to borderline F to non-F, but he must then tell us what this assumption amounts to and how it is defended.¹¹

He then considers two candidate neutral conceptions: ‘Intensity’ and ‘Temporal Duration’ and argues that neither of them suffice. Though we can apply ‘Intensity’ to both conscious and unconscious things, there is arguably an ambiguity here between ‘Phenomenal Intensity’ which implies ‘Phenomenally Conscious’ and other notions of intensity like ‘Stimulus Intensity’ which, arguably, are disqualified by Antony’s neutrality requirement. To handle ‘Temporal Duration’, Antony argues that a conception of a series of progressively shorter experiences does not give us a conception of borderline conscious states. A conception of an experience, however short,

¹¹Such an account will have to tell us how to filter through a lot of different potential trick responses, like ‘State that is an experience iff it is roughly X’, or the response setting the common conception for some series to be material concept M, and then taking the different determinations of that parameter to be different conjunctions of M: with ‘Conscious’, ‘Borderline Conscious’, and ‘Non-Conscious’, for example. I do not know if Antony has the machinery to adequately spell out what he needs here. I return to this point in section 2.

is a conception of an experience , and there is no conception of a sort of episode such that, if it is sufficiently long then it is an experience, if it is sufficiently short then it is not an experience, but if it is somewhere in the middle then it is a borderline case(2006b p.12). Again, Antony is assuming that the common conception must be the thing whose variation explains the transition from F to borderline F to non-F. He is also assuming that there can be no way of describing a state neutrally with respect to whether or not it is an experience, but so that the duration of that state determines whether or not it is an experience. He might show that no such description could allow 'Temporal Duration' to count as a common conception by assuming that some experiences may be instantaneous, and then appealing to his assumption that the common conception must characterize all Fs rather than just some of them. However, Antony believes that all experiences have some temporal duration (*ibid.*), so this strategy may not be open to him, and anyway even if it were it would not rule out that the relevant sort of neutral description itself gave us the terms we needed. However, Antony only claims to be considering potential common conceptions that we can actually already think of, so the onus may not be on him to cover all such exotic possibilities.

Finally, Antony argues that 'Phenomenally Conscious' does not exhibit combinatorial vagueness. Since the argument is quick, I reproduce it in entirety here. Antony says,

“Starting with Clear Fs, combinatory borderline cases are reached by imagining a series of individuals such that the number of parts, properties, etc. represented in the conception [F] is gradually reduced from individual to individual. Notice that as elements from [F] are removed, each successive [ψ] in the series not only represents its individual without the removed elements, but as *lacking the properties* the removed elements express. With that in mind, consider a [ψ] representing a clear conscious state. It will include consciousness-entailing elements like [phenomenology], [qualia], [subjectivity], etc. from the conception [conscious state]. Removing any element from that [ψ] will either leave consciousness-entailing elements in the next [ψ] in the series, or eliminate the last consciousness-entailing element. If any consciousness-entailing elements remain, the [ψ] will represent a clear conscious state. If none remain, then since as was argued above [consciousness] contains no borderline consciousness-entailing elements, the [ψ] will contain only clear non-consciousness-entailing elements. Since that [ψ] will represent its individual as lacking *all* properties expressed by all consciousness-entailing elements ... the individual will be represented as possessing no properties that suffice for consciousness.” (2006b p. 16)

The trouble with this argument is that the conception of a clear case of consciousness may include things other than unfactorizably consciousness-entailing elements. It may contain elements (component conceptions) that neither entail consciousness nor non-consciousness, but that when in combination with other such elements do entail consciousness. This is precisely what we should expect if ‘Phenomenally Conscious’ is in fact the sort of cluster concept that exhibits combinatory vagueness. So Antony’s argument seems to be begging the question, although again this may hinge on where we take the onus to lie, and how exotic we take the relevant possibility to be.

§1.2 PROBLEMS WITH ANTONY’S CONCEPTUAL ARGUMENT

My presentation of Antony's argument has already made clear some of the argument's internal problems. We need an account of what sorts of things may feed into the individual conceptions in the series, to rule out trivializing elements like 'Borderline F'. The argument that 'Phenomenally Conscious' exhibits no combinatory vagueness seems to beg the question. It is not obvious that there are no other sorts of vagueness besides combinatory and degree. There is so much loaded in to the notions of Dimension (from C3) and Common Conception (from C4), that it seems unlikely that these will really specify necessary conditions on the vagueness of concepts like 'Beautiful' or 'Funny' or even 'Bald'. And finally, Antony's case that 'Phenomenally Conscious' does not meet conditions C3 and C4 seems to rely on the loaded-in features that make it doubtful whether they really are necessary conditions on vagueness.

However, it is not my aim to dwell on these problems here. I suspect that they are symptomatic of a deeper problem, and it is of this deeper problem that I will speak. The deeper problem with the argument (with both the defense of Middle Ground, and the claim that 'Phenomenally Conscious' does not satisfy it) is that it is couched in the non-normative terms of descriptive psychology. I will argue that without appeal to some normative notion like license or justification, then either the Middle Ground condition is trivial, or the conclusion that 'Phenomenally Conscious' does not satisfy

Middle Ground is not nearly as interesting as it may appear. I will then argue that appealing to such a normative notion significantly weakens Antony's argument that 'Phenomenally Conscious' does not satisfy Middle Ground.

The risk of triviality arises first because of the possibility of false positives. Competent users make mistakes. Competent users may mistakenly display v-dispositions even when the case is not vague at all. In Antony's discussion of whether 'Phenomenally Conscious' satisfies C4, he considers 'Intensity' and 'Temporal Extent'. He argues that it would be confusion to think that 'experience of vanishingly small intensity' was a borderline case of experience. He is correct. But this does not mean that such confusion is impossible. And someone in the grip of this confusion would be able to produce a series that satisfied C1 – C4, using something like 'Phenomenal Intensity' as the common conception C4 demands.

A related risk of triviality arises because of a difficulty in saying what it takes to be a genuine Conception or dimension. Can there be gruesome or gerrymandered common conceptions, gruesome or gerrymandered representations of dimensions of gradual change? Without an effective filter of gruesome from non-gruesome conceptions and dimensions, C4 becomes a vacuous condition, and C3 approaches vacuity as well: Middle Ground collapses into the constraint that there be individual conceptions that happen to give rise to v-dispositions. And as I note in the previous paragraph, this can

happen by mistake. One reply would be to lean even more heavily on the idea implicit in C3 and C4 that the Conception or Dimension specifies the parameter whose variation *explains* or *grounds* the change from F to borderline F to non-F. But these are actually very tricky notions, and I do not know how we might get at them without either delving into the realm of the normative, or perhaps into speculative metaphysics.¹²

The most natural descriptive reply is to appeal to *our actual conceptions as we actually use them*. There may be gruesome conceptions out there, but we tend to not employ them in our mental representations. And though some people may have v-dispositions by mistake, most of us are on the right track most of the time. This is the line Antony seems to endorse. Though he does not present his position as a response to a worry about gruesomeness and mistaken v-dispositions, his is careful to restrict the central claim of his argument. It is not the claim that *the* concept 'Phenomenally Conscious' is non-vague but rather the claim that our *current* concept 'Phenomenally Conscious' is non-vague. Note that it is not obvious that the restriction of attention to our current *concepts* ensures that we need only consider our current *conceptions*.

¹²My own sense is that at very least, some notion of Conceptual Priority is called for. Antony might say that the relevant conception may not be conceptually posterior to any phenomenal concepts. The relevant generalization of this thought would make Middle Ground even more difficult to defend, but Middle Ground may wind up trivial otherwise. In my own argument, I evoke the Conceptual Priority relation but I put it to different (and I hope more defensible) use.

But we can perhaps think in terms of our current concepts as we currently conceptualize them, or just in terms of our current conceptions.

In any event, this restriction is critical to Antony's argument that 'Phenomenally Conscious' does not satisfy C4. Antony concedes that he does not have a knockdown argument for this conclusion. Instead, he considers a few of what seem to be the most plausible candidates for neutral conceptions that may be common to individual conceptions of conscious things, borderline conscious things and non-conscious things. Plausible by what measure? The examples he chooses are examples of the sort of neutral conceptions that we actually, currently tend to associate with 'Phenomenally Conscious'. It may be that future innovations of conceptual architecture will uncover new ways of conceptualizing consciousness, giving rise to radically different neutral conceptions that show us how consciousness may in fact be vague (either by degree, or by combination, or in some other way). But Antony does not consider this sort of case, since he takes it that such a discovery would only show that some future concept of consciousness is vague, and this is compatible with our current concept of consciousness being non-vague. The restriction also emerges in the way Antony discusses those cases he does consider. For example, he offers no positive argument that there will be no neutral conception of a state as either an experience or not an experience

depending on its temporal duration.¹³ And his case that ‘Phenomenally Conscious’ exhibits no combinatorial vagueness similarly contains no positive argument that the relevant sort of neutral conceptions will not emerge on future reflection.

To clarify how he understands the dialectic on this issue, Antony considers the example of the concept ‘Life’. Antony claims that the pre-modern concept ‘Life’ was non-vague. Our current concept ‘Life’ is vague, but according to Antony this is only thanks to biologists’ “...developing theories of such scope and power that earlier views of life could no longer be sustained.” (2006b p. 18). According to Antony, the upshot of the argument that our current concept ‘Phenomenally Conscious’ is non-vague is that it puts the onus on theorists of the nature of consciousness who think consciousness can be vague, to develop a theory of such scope and power that our earlier views of consciousness (as non-vague) may no longer be sustained.

But Antony makes an unwarranted assumption. He assumes that at some point our concept ‘Life’ went from being neutral with respect to the physical or spiritual nature of life to being non-neutral (and committed to the physical nature of life). This is the sort of change that he thinks of as warranted only by theories of suitable scope and power. He also makes the

¹³Not to mention that he still owes us a reason why ‘State that is a conscious state iff it last longer than, roughly, a second’ does not count. This is an example of the sort of thing I suspect Antony needs the Conceptual Priority relation to manage (see previous footnote).

unwarranted assumption that it is only in light of this change that our concept of 'Life' went from being non-vague to being vague (*ibid.*).

Antony owes us an argument for each of these claims. True, we may now all believe that Life is a physical phenomenon. But does this mean that it is now a part of our concept 'Life' that Life is a physical phenomenon? Perhaps on a psychological model of concepts where there is no distinction between what the concept entails and what is typically associated with the concept, this assumption is permissible, though the matter is unobvious. But more importantly, though it does seem that the pre-moderns took 'Life' to be non-vague and that we take it to be vague, it is not obvious that this transition turned on any scientific discovery. To be sure, we relied on discovery to learn that life has to do with carbon, and to learn that such things as viruses actually exist. But one might argue that it was open to the pre-moderns to observe, without the benefit of any scientific innovation, that 'Life' has many of the features of a cluster concept: that it always has been part of our conception of life that living things can grow, move, respond to stimuli, metabolize and reproduce. 'Vital Force' may have been an independently associated component of 'Life' but it was more commonly taken only to be a hypothesis about the underlying basis of these phenomena (Bechtel and Richardson 1998). The association may have been strong, but if the picture I sketch here is correct then even in the pre-modern age competent users of the concept 'Life'

were in a position to recognize that there might be borderline cases – things that could grow and move but that did not respond to stimuli or reproduce (or perhaps things that had the Vital Force but could not grow or move or respond or reproduce). This would mean that if people in those days thought that ‘Life’ was not a vague concept, then this was due more to failure of imagination than to ignorance of future discovery. The situation would be analogous to the one that Parfit 1984 claims we are in regarding Personal Identity: because we tend to think of Personal Identity in terms of the survival of the soul, we do not naturally confront those cases where some of the things that matter about survival disentangle from others. But all it takes to get us to do so are a few cleverly described thought experiments.

If this is the correct model for ‘Life’ then Antony may still say that ‘Life’ as it was used was non-vague, though now it is vague. But this would only be to say that people in pre-modern times failed to recognize the ways the concept was vague. In this case the onus never was on biologists to develop an elaborate theory according to which ‘Life’ can be vague, in order to convince us that the concept is vague. The ordinary folk were already in a position to see this, they just happened not to. If this is the correct model, then whether or not it took the biological revolution to change whether ordinary people happened to notice the vagueness in ‘Life’ was a matter of historical accident.

Antony may be able to establish that 'Phenomenally Conscious' as we currently conceptualize it is not vague, but if the model I have just sketched applies, then this only means that we have not been imaginative enough to see how it might be vague, though it might be. Some future scientific innovation might suffice to stimulate our imagination on this front, but we might also stimulate our imagination without such a revolution. Imaginative reflection on its own might serve equally well. Antony owes us an argument that this is not the right model on which to understand the apparent past non-vagueness of 'Life' and 'Personal Identity'. Otherwise, these examples show that this sort of failure of imaginativity really happens, and so it is not unreasonable to think that our current inability to see the vagueness of 'Phenomenally Conscious' is such a failure. And even if the model does not apply in those cases we still have reason to wonder whether it applies in the case of 'Phenomenally Conscious'. Antony owes us additional reasons to think that no such imaginative innovation is forthcoming. Without such additional reasons, the sense in which he has shown us that 'Phenomenally Conscious' is non-vague is much less newsworthy than it may appear.

The troubles we have seen here result because, if we adhere to the language of descriptive psychology, it is difficult to find a middle ground between a condition that almost any concept can satisfy (because people sometimes take things to be vague even when they *shouldn't*), and a condition

that too few concepts satisfy (because people sometimes fail to recognize things as vague even when they *should*). What we require, then, is a principle couched in more epistemic terms, a principle articulating a condition necessary for competent users to be *warranted* in taking that concept to be vague. In chapter two of this dissertation, I articulate such a principle (I call it *Excludability*¹⁴). But as we have just seen, the bar for showing that there can be no *warrant* of the right sort for thinking that ‘Phenomenally Conscious’ is not vague, is higher than the bar for showing that as we currently tend to conceptualize it, we do not see how ‘Phenomenally Conscious’ might be vague. We need to consider all possible warranted conceptualizations, not merely typical, current conceptualizations. In chapter three, I offer such an argument for this more powerful claim based on a consideration of all possible warranted conceptualizations.¹⁵

¹⁴A concept ‘F’ is *Excludable* if we can form some substantive, positive conception of a concept that excludes it (in the way that ‘Red’ is excluded by ‘Green’). I employ the conceptual priority relation to characterize the notion of a *positive* conception. My formulation of the argument circumvents the more particular problems of formulation that Antony faces. It allows me to abstract away from the difference between combinatorial and degree vagueness, and from having to say what exactly a gradual change is, and I do not rely on the claim that there is some common element, in any interesting sense, to all Fs, some borderline Fs and some non-Fs (nor on the claim that there is some common element to *some* Fs, some borderline Fs and some non-Fs). Also, I do not rely on the claim that vagueness necessarily involves the existence of an invariant class of borderline cases (a matter that some think Higher Order vagueness calls into doubt. Cf Fine 2009). My condition is also far more general, and it is far more difficult to generate cases that would have a chance of being counterexamples. For example, ‘Lame’ excludes ‘Cool’, ‘Ugly’ excludes ‘Beautiful’, and so on.

¹⁵One might wonder whether Antony’s focus on psychological rather than philosophical concepts was part of the problem here. Certainly the epistemic language comes more naturally when we speak of abstract immutable senses rather than occurrent mental representations,

§1.3 ANTONY'S METAPHYSICAL CONCLUSIONS

I turn now to a brief discussion of Antony's argument from the conceptual claim that 'Phenomenally Conscious' is not vague to a metaphysical conclusion. Antony observes that most of our ordinary formulations of materialism are couched in vague terms. Most of our functional notions, and most of our neurophysiological notions, are vague. But how, Antony asks, could an identity claim, 'C = N', be true, if 'C' has no borderline cases, while 'N' has borderline cases? In such an event we would have some case, x, such that it was borderline whether or not Nx, but not borderline whether or not Cx. This means that ' $\forall x.Cx \text{ iff } Nx$ ' is not determinately true. But if ' $\forall x.Cx \text{ iff } Nx$ ' is not determinately true then neither is 'C = N'(2006a pp.5-9)

Antony is not clear on whether his goal is the claim that such an identity cannot be determinately true, or rather the claim that such an identity must be false. Evans 1978 argues that if vagueness is ontological then there can be no vague identity, so if some identity statement is not determinately true then it is false. Antony's argument appears to be an attempt to generalize this argument – Antony invokes a second order quantificational schema, and

but one may hold that epistemic norms apply to psychological concepts, so there is no obvious connection here.

Leibniz's law, much as Evans does, to try to show that this result applies even if vagueness is semantic (2006a p.7). For those who take vagueness to be epistemic Antony offers a different argument: if vagueness is epistemic then both vague concepts and precise concepts refer to single precise properties, but it would be a great coincidence if some precise term happened to have the same exact denotation as some vague term (2006a p. 8-9).

Antony notes that the above reasoning is only of interest if the ultimate, correct concept of consciousness is non-vague. It is here that he argues that the non-vagueness of the present concept places the onus squarely on the shoulders of those who would argue that a future concept will be vague, to develop a theory of the nature of consciousness of such scope and power that our current (neutral) views of the matter can no longer be maintained (2006a pp. 12). Antony concludes that we have reason to disbelieve all versions of materialism that identify consciousness in vague terms (as well as versions of dualism that identify the material *correlates* of consciousness in vague terms). He takes this to imply that, for a true theory of the identity (or correlates) of consciousness, barring Eliminativism, we must look to the language of fundamental physics (2006a p. 13).

§1.4 PROBLEMS WITH ANTONY'S METAPHYSICAL CONCLUSIONS

A first problem with Antony's metaphysical argument is that, as I argue above, Antony has not given us compelling reasons to think that the ultimately correct future concept of consciousness is non-vague. Until we can show that the non-vagueness of our current concept does not hinge on a mere failure of imagination, it is difficult to say where the onus lies. But as I will argue in chapter three, I believe it is possible to show this. So it is worth considering Antony's argument on the assumption that the ultimately correct future concept of consciousness may well be non-vague.

The next major problem is Antony's claim to have established that vague identities like 'N=C' must be false (as opposed to merely not determinately true). If vagueness is ontological, the argument is simply Evans' (Evans 1978), and we may grant the point. But Antony attempts to generalize Evans' argument to the case where vagueness is semantic, and it is not clear that this generalization succeeds. The application of Leibniz' Law and schematic second order quantification is not obviously warranted if the term in question is semantically vague, because this may mean there is no specific thing that the term determinately refers to (Evans, *ibid.* and also Lewis 1988). There is also a problem with Antony's argument for the case that vagueness is epistemic. Antony says it would be a great coincidence if the semantic facts that determine which of the points in the borderline of 'N' is its true (hidden) exact border, determine it to be the very same as the border of 'C' (2006a p.9).

He relies here on the assumption that the mechanism which determines the true (hidden) border of 'N' is independent of the mechanism determining the border of 'C'. But this turns on the as yet forthcoming details of the Epistemicist account of how semantic mechanisms determine true (hidden) exact borders for vague terms.

So we may grant Antony that if vagueness is ontological, then identities like 'N=C' must be false, rather than merely indeterminate.¹⁶ I am less sure that he has a case if vagueness is epistemic. In that case, assuming the relevant mechanisms are properly coordinated, the identity might not only be indeterminate but in some sense perfectly true, in that the vague term 'N' might refer to exactly the same property as the precise term 'C'.

The semantic case raises interesting questions. I do not see that Antony can establish more than that such identities are not determinately true – they might all be false, or some of them might be indeterminate and not false. This result is still newsworthy. If we know that 'Phenomenally Conscious' is not vague, then assuming it denotes a property, we know there is a single precise property that it determinately denotes. If vagueness is semantic then some

¹⁶ But for a defense of the possibility of vague identity if vagueness is ontological, see Lowe 2001 and Barnes and Williams 2009. It is also worth noting that Antony claims his argument applies to the covariation claims a dualist makes as well as the identity claims a materialist makes. But the Evans argument does not extend beyond identity claims, so the dualist who embraces ontic vagueness may hold that the correlation claim is indeterminately correct rather than false.

identity like 'N = C' may be indeterminate rather than false, but it would still be the case that the language of the theory in which 'N' occurs is not capable of generating an exact description of exactly the property **Phenomenal Consciousness**. This, I take it, is the kernel of Antony's argument, and it is an eminently plausible point.

But there is still a question of whether this eminently plausible point means that the ultimate theory of consciousness must be couched in the language of fundamental physics. First of all, there are precise descriptions in the language of non-fundamental science, for example the description "oscillates at a frequency of exactly 40 Hz"¹⁷. Second of all, it might be that there is no way of describing the referent of 'Phenomenally Conscious' in the language of fundamental physics. It might call for some infinitary disjunction to generate an intensionally adequate characterization, which is bad enough, and there might be further hyperintensional distinctions that such an infinitary disjunction does not make (consider the way that 'Red' may covary with some disjunction of possible shades of red, even though **Red** is not identical to that disjunctive property). It may then be that there is no determinately correct theory of the material basis of consciousness, and the best available theory would be a (merely) roughly correct theory couched in

¹⁷ The usual version of the 40 Hz theory identifies 40 Hz as only a prototypical value rather than an exact value, but in any event such a theory is available.

the language of some higher level science. To make this work we could take an indeterminate identity between a vague and a precise general term to assert that the property denoted by the precise term is one of the candidate properties indeterminately denoted by the vague term, and we might take an identity between two vague general terms to assert that there is some overlap between the sets of candidate properties that each concept indeterminately denotes.¹⁸

However, Antony is still correct that if 'Phenomenally Conscious' is precise, then if any language is going to tell us *exactly* what material property **Phenomenal Consciousness** is (or perfectly correlates with), that language had better be precise. This is an interesting result, even if it does not imply that the ultimately correct theory of consciousness will be couched in the language of fundamental physics.

¹⁸ It is worth pointing out that we often do accept identities between terms that seem to differ over their borderline profiles, for example 'Water = H₂O'. Thanks to Dan Lopez De Sa for this point. A clear, drinkable mixture of H₂O and XYZ might be a clear case of 'Water' but an indeterminate case of 'H₂O'. However, if we accept the identity we accept that water really just is H₂O, which means that we take the borderline profile of 'Water' to inherit the borderline profile of 'H₂O'. If 'Water' is a rigid designator then presumably it denoted H₂O all along, and so our initial sense of disagreement in borderline profiles was really just confusing 'Water' with 'Watery'. Antony might put this in terms of our pretheoretic conception of water being replaced by a correct, scientific conception. Since Antony argues that this is only going to be an option in the case of 'Phenomenally Conscious' if scientists present us with a theory of irresistible power and scope, it is not fair to hold this revision model of vague identity statements as a counterexample to his claim. But for the counterclaim, see (Lopez De Sa, forthcoming).

My own most pressing grievance with Antony's argument however is not with anything he says but with something he does not say. I think the non-vagueness of 'Phenomenally Conscious' shows us something even more interesting than that a strictly correct theory must be couched in precise language. **Phenomenal Consciousness** is significant, in both descriptive and normative respects. But it turns out that the non-vagueness of 'Phenomenally Conscious' forces us to choose between the thesis that **Phenomenal Consciousness** is significant and the thesis that **Phenomenal Consciousness** is material. In chapters four and five of this dissertation, I argue that if 'Phenomenally Conscious' is non-vague then either it is neither descriptively or normatively significant, or one of Property Dualism or Panpsychism is true.

CHAPTER TWO:

GENERALIZED CONCEPTUAL GAP PRINCIPLES AND VAGUENESS

Many Materialists agree with Dualists that there is a gap between our way of thinking about the mind and our ways of thinking about the brain, the body and the rest of the physical world. As it usually occurs the idea of a gap is a relational idea: it relates one family of concepts to another family. And as it usually occurs one of these families is the family of phenomenal concepts, while the other is a family of material concepts. In this essay I will consider some ways of generalizing the idea of a conceptual gap. I hope to generalize in two respects: first, from a relational notion to an absolute one, and second, from a notion characterizing phenomenal concepts, to a notion characterizing arbitrary families of concepts, or a notion characterizing individual concepts, whether they belong to a prominent family or not.

In doing so I have two aims. My first aim is to show that sense can be made of notions of generalized conceptual gap, and that these notions are interesting. To this end, I argue that these notions frame debates elsewhere than just in the philosophy of mind – for example normative and metaphysical debates (such as the debate over the Autonomy of Ethics, or the question of whether we only understand terms like ‘Abstract’ or ‘Immaterial’ as negations of ‘Concrete’ or ‘Material’), and I will argue that these notions are important within the philosophy of mind, where they allow us to frame new hypotheses about what makes phenomenal concepts special.

My second aim is to show a link between one of these generalized notions of Gap, a notion I will call *Non-Excludability*, and vagueness. I will argue that if a concept is *non-excludable* then it is not vague, assuming that vagueness is a semantic phenomenon. In the next chapter I will argue that ‘Phenomenally Conscious’ is non-excludable, from which it follows that ‘Phenomenally Conscious’ is not vague.¹⁹

In §2.1 I introduce the Generalized Conceptual Gap principles. In §2.2 I suggest that these principles, though extremely selective, may apply to some concepts outside of the philosophy of mind – for example normative concepts. In §2.3 I argue that even if many different families of concepts exhibit gaps in the ordinary relational sense to other families of concepts, phenomenal concepts would be very distinguished if there were gaps between them and other concepts in any of the absolute senses I will consider in this chapter. In §2.4 I argue that if a concept is non-excludable then it is not vague (assuming that vagueness is a semantic phenomenon).

¹⁹ How, you ask, could this be, given the difficulty of defining ‘Phenomenally Conscious’ and in determining what it applies to? Concepts may be precise either *de facto* or *de jure*. A concept that is precise *de jure* is a concept that makes its own precise boundaries explicit. Concepts that are precise *de facto* are precise without making their own boundaries explicit. Non-excludable concepts are *de facto* precise. Reference magnetism (if there is such a thing) is another sort of *de facto* precision. But where reference magnets give rise to *de facto* precision because they are natural joints in metaphysical reality, non-excludable concepts are natural joints in conceptual reality. Both are ways in which some concept may be precise, even though masters of the concept may not know where the boundaries actually lie.

§2.1: GENERALIZED GAP PRINCIPLES

I will understand there to be an entailment relation between concepts – the relation that holds between them when there is no gap – that mirrors the entailment relation between propositions. In particular, I'll say that one concept 'C', entails another concept, 'D', just when for any object x, the proposition <for all x, if x is a C, then x is a D> is *a priori*. I take it that proposition <P> entails proposition <Q> just in case <if P, then Q> is *a priori*. I have in mind a very comprehensive conception of concepts: if something is the meaning of a monadic predicate in a possible language, finitary or infinitary, then it is a concept, and concepts may be distinct although *intensionally* equivalent, if they involve different logical forms, different components or different definitions.²⁰ Those who wish may replace 'Possible Predicate' for 'Concept'. I also have in mind a permissive conception of *a priority*: the difference between something that is knowable *a priori*, and something that is knowable on the basis of armchair knowledge, or something for which we rationally have extremely high credence, is not an important difference for my purposes.

²⁰ I suppress further generalizations for simplicity. We might relax the constraint that C and D be monadic, as well as the constraint that they be of matching arity, and we might call it entailment if the relation holds for *some* x rather than for all x, or for some course of values for C and some (potentially different) course of values for D, rather than demanding that they be the same. In any case, I mean to treat the claim that phenomenal facts are entailed by a specification of all of the physical facts, together with a specification that those are all the physical facts, as an entailment claim in my sense.

We express something in the vicinity of the intuitive idea of the Material-Phenomenal Conceptual Gap when we say: No phenomenal concept is entailed by any material concept.²¹

Whether this claim is true hinges on how we delineate the categories *phenomenal concept* and *material concept*. Are intentional concepts phenomenal concepts? Are functional concepts, or concepts of yet to be developed physics, material concepts? It depends on the distinctions we are interested in drawing. Usually this is clear enough from the context, but there may not be any single best partition of concepts into familial categories. If there were such a partition we could use it to characterize a Generalized Conceptual Gap: we might say that one concept (or family of concepts) exhibits such a gap when it is not entailed by any concept in any distinct family.

In special cases, such an approach yields fruit. The phenomenal concepts are such a special case. We have a good enough sense of which concepts are concepts of relatively specific types of experience – for example, the concept ‘Seeing Scarlet Red’ or the concept ‘Hearing a Trombone’ or the

²¹This is a very strong formulation, and to be made plausible must be qualified somewhat. For example, if Property Dualism is *a priori* then a physical description with a ‘That’s All’ clause (eg: ‘Has Physical Properties and Only Physical Properties’) will entail something about consciousness: its absence. Also, since many phenomenal descriptions will entail functional descriptions (which even many dualists grant), some negative functional descriptions will entail that certain phenomenal descriptions do not apply. Once we take these qualifications into account, we see how other analogous local gap principles may also be defensible, for example the thesis that macroscopic concepts (‘Water’) are not entailed by microscopic ones (‘H₂O’).

concept 'Feeling a Stinging Pain in one's Knee.' Let's call these the *canonical phenomenal concepts*. We might then characterize the claim that phenomenal concepts exhibit such a gap as the claim that no canonical phenomenal concept is entailed by any other sort of concept.

We need a notion of conceptual priority to give this characterization a fighting chance. 'Seeing Scarlet Red' is entailed by the conjunctive concept '((Seeing Scarlet Red or Being 5 feet Tall) and Not Being Five Feet Tall)' and by the negative concept 'Not Not Seeing Scarlet Red.' But these concepts are, in an obvious sense, *conceptually posterior* to the concept 'Seeing Scarlet Red' (I will say more about conceptual posteriority below). We may now characterize a Generalized Gap claim for phenomenal concepts:

(PHENOMENAL INSCRUTABILITY): No canonical phenomenal concept is entailed by any other concept of any other sort, except concepts conceptually posterior to canonical phenomenal concepts.

This claim has a fighting chance because we have a clear-enough category of canonical phenomenal concepts. This suggests that for any family of concepts for which there is some clear-enough canonical subset we can define

(CANONICAL INSCRUTABILITY): The family of concepts F is *canonically inscrutable* iff no canonical F concept is entailed by any other concept of any other sort, except concepts conceptually posterior to canonical F concepts.

A very closely related idea, leaning slightly more heavily on the notion of conceptual priority is:

(CANONICAL PRIMITIVITY): The family of concepts F is *canonically primitive* iff no canonical F concept is conceptually posterior to any concept of any other sort.

These definitions capture something but not everything that we might wish in a generalization of the idea of a Conceptual Gap. Such a generalization should capture, to the extent possible, the idea that our canonical phenomenal concepts are autonomous, and not to be defined or constructed or explained²² in terms of any more basic concepts. Suppose we had some conception of how to factor consciousness into components, but in such a way that each of these components depended on the others and so never occurred except as a part of some conscious whole. There is an intuitive sense in which the phenomenal concepts might be autonomous even though some concept like this exists, since this concept shows us no connection between things satisfying phenomenal concepts and things that do not. Similarly, suppose some very

²² Where one concept is explanatorily prior to another if the application of one helps explain the application of the other.

general precondition for consciousness, also a precondition for various forms of non-consciousness, turns out to be conceptually prior, or anyway not conceptually posterior, to some canonical phenomenal concepts. There is a sense in which the phenomenal concepts might be autonomous even if some concept like this exists, because such a concept might not give us any particularly substantive conception of the connection between things satisfying phenomenal concepts and things that do not (suppose the concept were 'Material Being' or 'Exhibits Some Functional Organization'). Finally, suppose that we have some conception, neither prior nor posterior to any phenomenal concepts, of a determinate way of being non-conscious. Such a conception might show us how a determinate sort of non-consciousness is related to a determinate sort of consciousness. Given such a conception, there is a sense in which the phenomenal is *not* autonomous, even if phenomenal concepts are both *canonically inscrutable* and *canonically primitive*. These are three respects in which the two defined notions do not capture all of the intuitions that we might hope for in a generalization of the notion of a conceptual gap. And then there is also our original point that there is no general way to partition concepts into families, and even when we do have a family of concepts there is no general way to tell canonical from non-canonical elements of that family. I will now develop a notion that is consonant with

these intuitions and does not rely on concepts' having well-defined familial relations.

This last notion I will develop is the notion of *Non-excludability*. Some concepts exclude other concepts. 'Red' excludes 'Green'. 'Dirt' excludes 'Water', 'Tall' excludes 'Short', 'Married' excludes 'Single' and so on. A very distinctive concept would be one that was not excluded by any other concept at all. No concept is that distinctive, since every concept is excluded by its own negation. But we may distinguish between those excluders that give us a positive idea of the alternative which does the excluding, and those excluders that merely negate the concept, or negate some necessary condition for its application:

(NON-EXCLUDABILITY): A concept 'C' is *non-excludable* iff the only concepts that exclude 'C' exclude it negatively

To exclude 'C' is just to entail ' $\sim C$ '. I say what negative exclusion is below. NON-EXCLUDABILITY is a generalized notion of a conceptual gap because a non-excludable concept is one that cannot be situated into any broader conceptual space: we have no positive conception of what such a concept contrasts with, no sense of how things that satisfy it relate to things that do not. It depicts a dimension all of its own. Also, non-excludability implies the ordinary relational sort of conceptual gaps: if 'C' is non-excludable then for any family of concepts F, no concept in F positively excludes 'C'. But if

F is a family of concepts that are not conceptually posterior to 'C', then there will be a conceptual gap (in the ordinary relational sense) between 'C' and the F concepts.

To properly draw the contrast between negative and positive exclusion we must again invoke the Conceptual Priority relation. One concept 'N' *negatively excludes* another concept 'P' just when 'N' excludes 'P' ('N' entails ' \sim P') and also 'N' is conceptually posterior to some concept that articulates an *a priori* necessary condition for 'P' to obtain (some concept entailed by 'P').

Many concepts entail the negations of other concepts. I am drawing here on the more substantive idea that some concepts *really are* negative, while others *really are* positive. 'Red' is positive and 'Not Red' is negative. '(Not Red) and Tall' is a mixture of negative and positive, but it is conceptually posterior to 'Red', and so its exclusion of 'Red' is negative exclusion.

We may further qualify the idea of negative exclusion. For example, if I enumerate your qualities, and then declare that I have enumerated all of your qualities, but I omit your sardonic wit, then in declaring that I have enumerated all of your qualities I am saying you do not have a sardonic wit. But my description only negatively excludes sardonic wit: I do not describe

some other virtue of yours that rules it out. Rather, I list only other virtues and then say 'that's all'. 'That's All' exclusion is also a form of negative exclusion.²³

But the key idea is the contrast between negative and positive concepts. This is a substantial contrast. It cannot be reduced to any lexical contrast, like a contrast in the logical form of the concepts involved. Nor ought it be reduced to some sociological or psychological contrast. Suppose there were a tribe that for reasons of contingent psychology found 'Not Red' or 'Color other than Red' to be a more graspable notion than 'Red'. In this tribe people first learn 'Not Red' (for which they have some single term, 'Nred') and define 'Red' as the negation of 'Nred'. Still, I say, 'Nred' is conceptually posterior to 'Red'.²⁴ This is

²³ If dualism is true *a priori* then the description, 'Has Only Physical Properties' entails 'Not Phenomenally Conscious' by means of such a That's All clause. Since I am open to the *a priori* truth of dualism (though my own argument for dualism is not fully *a priori*), but I also take 'Phenomenally Conscious' to be non-excludable, it is important for me that exclusion by that's-all clause counts as negative exclusion. You might worry that the that's-all clause makes the concept of exclusion too exclusive. For example you might worry that 'Sweet' only excludes 'Sour' given some sort of restricted That's All clause ("That's All of the Taste Qualities Here"). But this is mistaken. 'Sweet' excludes 'Sour' the way that 'Red' excludes 'Green'. 'Sweet and Sour' is a coherent predicate but so is 'Red and Green'. Nothing is both sweet and sour on the same taste bud at the same time. Even if it were, we have positive conceptions of a more determinate sort of taste space, and concepts like 'Bittersweet' exclude concepts like 'Honey Sweet'.

²⁴ How to support this claim? We might say that if two concepts negate each other, then at most one of them can seem to present a positive property. Intuitively, a property is positive iff having it is an achievement – something God had to do – rather than simply a failure to have some other sort of property. **Being red** is positive, and **not being red** is negative. **Being green**, though it necessitates **not being red**, is positive. We can also understand a positive property as one that you cannot lose simply by acquiring new properties. Chalmers (1996) defines a positive property as "One that if instantiated at a world W, is also instantiated by the corresponding individual at all worlds that contain W as a proper part". Both my formulation and Chalmers' suppress difficulties. In mine there is a contrast between one way of losing a property (simply acquiring another contrary one) and a more substantial way of losing a property. In Chalmers' we have the notion of what is instantiated at all worlds that contain W as a proper part. You might think that Positivity goes along with Naturalness or

not to say that for every pair of predicative concepts that negates each other, one must be negative and the other must be positive, only that this is often the case. 'Married' and 'Single' both seem to be positive, for example.²⁵ In such cases neither concept is non-excludable.

§2.2: THE INTEREST OF NON-EXCLUDABILITY, CANONICAL PRIMITIVITY AND CANONICAL INSCRUTABILITY

Very broadly speaking, the interest of these notions is that they give us alternative characterizations of respects in which some concept or family of concepts is *primitive*, or autonomous. Most philosophers these days are skeptical of definitional approaches to primitivity, but other approaches are available; the sense that some concepts are more basic than others survives the death of the hopes that the majority of concepts are *defined* in terms of some small core collection.

Fundamentality – we might say that positive properties just are fundamental properties. But a conjunction of fundamental properties is still a positive property, though it is not a fundamental property, and also on some accounts of fundamentality (for example, those connecting it with Causal Role) intuitively negative properties may come out fundamental. The distinction between a positive and a non-positive property may ultimately be a conceptually primitive distinction, but it is an important one. We might also characterize the conceptual priority relation more directly in terms of which concepts are more fundamental in the sense of carving more closely to natural joints (cf. Sider).

²⁵ The definition is not vacuous. If 'A' entails '~C', then 'C' entails '~A'. Might '~A' be the concept 'D' from the definition? We have that 'C' entails '~A'. 'A' automatically entails '~~A'. But 'A' is not posterior to '~A'. Only if 'A' were posterior to '~A' would we have some proof that every concept is non-excludable.

Equipped with our notion of conceptual priority we might characterize a primitive concept as one that is not posterior to any other concept. But there are other related notions to be had that do not rely quite so explicitly on the details of the conceptual priority relation, and the notions of Non-excludability, Canonical Primitivity and Canonical Inscrutability are among these: the satisfiability of each of these notions is compatible with the possibility that every concept is posterior to some other concept. Anyway, it is likely that there is no univocal notion of conceptual primitivity, but instead there will be different such notions that capture different aspects of the general idea.²⁶ If some family of concepts is canonically inscrutable, then its canonical elements may entail one another, but they are not entailed by any other concepts, and so, in one sense, primitive as a group. Likewise if a family of concepts is canonically primitive, then its canonical elements may be prior only to one another, which means that even if each of them is posterior to some other one of them, the family as a whole is primitive relative to everything else. A non-excludable concept is one that is primitive in the sense that it tells us little about how its instances must agree or contrast with the instances of other concepts. A non-excludable concept must also be something

²⁶ Chalmers (Constructing the World) is a sustained discussion of these issues, and a characterization of a few more conceptions of primitivity. Chalmers asks which concepts we require to articulate a basis of truths from which all truths would be *a priori* deducible (scrutable). Figuring in such a basis is a very natural sense of primitivity (albeit a relativizable one).

like the *head* of a family of concepts, in roughly the sense of describing a maximally determinable property rather than some more determinate one, since determinates are excluded by other determinates of the same determinable.

G.E. Moore famously asserted that the property of goodness is simple and unanalyseable (or at least unanalyseable in non-moral terms)²⁷. Moore's argument from this premise to Non-naturalism is often dismissed as a relic of a time before *a posteriori* necessity, but we ought not assume that this means that the concept 'Good' is no more distinguished than the concept 'Water'. 'Good' may not be non-excludable – it is presumably excluded by 'Bad' (though not if St. Augustine is correct that 'Bad' is a pure negation), but 'Has Normative Status' may be non-excludable (taking neutrality to be a kind of status), and the family of moral concepts (or anyway the family of normative concepts) may be canonically inscrutable or canonically primitive, as may be some subfamilies, like the family of deontic concepts.

Non-excludability raises some interesting questions in metaphysical debates. It is hard to find examples of concepts that are non-excludable. 'Water' is excludable: 'Solid Metal', 'Thin Air', 'Empty Vacuum' are all positive conceptions of 'Not Water'. 'Married' is excluded by 'Single', 'Red' by 'Green',

²⁷ Moore, *Principia Ethica*

'Tall' by 'Short', and so on. Things get interesting when we ask about more general terms like 'Concrete' or 'Material'. 'Immaterial' is certainly conceptually posterior to 'Material' but there are other concepts like 'Being the Number Two' that presumably are not (unless such concepts prove to actually be derivative themselves: mere negations of materiality). What about 'Concrete'? It is not obvious that 'Abstract' is downstream from 'Concrete'. It may be: some think we define abstract as the negation of concrete. But there do seem to be more positive conceptions of the abstract: Angels, Ghosts, Propositions, Numbers, Universals and Forms may give us positive conceptions of the abstract. If you take all of these notions to be negations, then 'Concrete' may well be non-excludable. This might be the position of a certain sort of nominalist, who held that there really are no such things as abstracta, and our abstract concepts do not really characterize existences distinct from concrete particular existences, but rather merely take characterizations of concrete particulars and negate key features of them. These are not the only interesting cases. According to some negative theologians, the concept of 'Ungodly' may be non-excludable – we only understand God by negation. Strict Finitists may feel this way about the concept 'Finite'. Modal concepts (possibility, necessity) may be as a family canonically primitive or inscrutable, and perhaps one of 'Possible' or 'Necessary' are prior to the other, in which the prior one may be non-

excludable. In chapter two, I will argue below that ‘Phenomenally Conscious’ is non-excludable.

There is a problem of easy exclusion. Concepts like ‘Normatively Evaluable’ or ‘Phenomenally Conscious’ seem to be excluded by concepts like ‘Being the Number Two’ or ‘Being the Conjunction Operation’. I do not think it is obvious that these are genuine exclusion relations²⁸ but in discussion I have found that few share my doubt on this point. Happily nothing hinges on this: non-excludability can be relativized without losing its interest. Too much relativization would trivialize things: Concept ‘C’ is not positively excluded except by those concepts that positively exclude it. But if we hold that a concept requires its instances to be of some extremely broad ontological category (or anyway that it does not rule out that anything in that category is an instance) then we may ask whether it is non-excludable within that category. Thus if you think that some things are excluded from being conscious because they are too abstract to be conscious, ask whether ‘Phenomenally Conscious’ is non-excludable by concepts whose instances are concrete beings. If you think that only actions, or only intentions, are normatively evaluable,

²⁸ If it is conceivable that rocks are conscious, and that ghosts and angels are conscious, and that God is conscious, why then should it be inconceivable that the number two or the conjunction operator should be conscious? Truly universal panpsychism seems to be a coherent possibility. Possibly on a Neo-Fregean conception of number the matter is determined by the implicit definition of number – though this is not obvious [cf Wright and Hale], and even if so it seems to amount to a case of That’s-All exclusion (which is a kind of negative exclusion).

ask whether 'Normatively Evaluable' is non-excludable by concepts whose instances are actions, or intentions.²⁹

§2.3: GENERALIZED CONCEPTUAL GAPS AND THE STATUS OF PHENOMENAL CONCEPTS

In the philosophy of mind our notions of Generalized Conceptual Gap are particularly interesting, as there are advanced debates about what makes phenomenal concepts special that hinge on the ubiquity of conceptual gaps. Many accept that there is a gap between phenomenal concepts and material concepts but there is a question about how special this makes the phenomenal concepts. Chalmers argues that it makes them extremely special, but many hold that it does not. Many materialists (among them Joseph Levine, Ned Block and Robert Stalnaker) hold that this sort of conceptual gap is commonplace (and that it therefore does not account for the *Explanatory Gap* which is a gap in a sort of explanation that is compatible with a lack of *a priori* entailment). Block and Stalnaker argue, for example, that there are analogous conceptual

²⁹ Other alternative refinements that I shall not pursue here, but that merit consideration: we might refine the notion of conceptual entailment. Thanks to David Chalmers for this suggestion. We might say: Concept 'C' *explanatorily entails* concept 'D' iff 'C' entails 'D' but also that something is C helps explain (display the grounds of) what it is to be D. Alternatively, we might ask not just for some positive conception of some way of being $\sim C$, but rather for a positive conception of a natural genus of which some ways of being C and some ways of being $\sim C$ are both species.

gaps between microphysical concepts and macrophysical concepts. If the conceptual gaps associated with 'Conscious' are no different from the gaps associated with 'Water' or 'Heat', then those conceptual gaps clearly do not account for the Explanatory Gap. This means for example that materialists can hold out hope that the Explanatory Gap will one day somehow be closed even though the Conceptual Gap will not be. But even if we grant Block, Stalnaker and Levine that the Explanatory Gap is not simply a matter of the conceptual gap between phenomenal concepts and material concepts, it does not follow that the Explanatory Gap is not a matter of some deep conceptual distinction. I propose here that it is at least in part a matter of the non-excludability of 'Phenomenally Conscious' and perhaps also of the canonical inscrutability or canonical primitivity of the phenomenal concepts in general.

Water facts may not be entailed by microphysical facts³⁰ but 'Water' is certainly excludable by other macroscopic concepts, for example 'Solid Ground' or 'Hard Steel'. This goes for just about any mundane concept we can imagine. Even the concept 'Macroscopic Object' is excluded by 'Microscopic Object, or failing that (because you think 'Microscopic' is just the negation of 'Macroscopic') it is excluded by any specific positive microscopic concept, for example the concept 'Electron'.

³⁰ I have stressed that as I understand it, 'entailment' is flexible between *a priori* and some slightly weaker notion like 'is armchair knowable from' or 'very strongly confirms'. Block and Stalnaker only claim that water facts are not *a priori* entailed by microphysical facts (plus totalizer and indexical facts), but my point holds no matter how we interpret 'entails'.

If 'Phenomenally Conscious' is non-excludable (as I argue in the next chapter), it stands apart from ordinary macroscopic concepts. It may even be that the Explanatory Gap is a function of this sort of conceptual divide. This would suggest that the Explanatory Gap is not bridgeable unless we change our concepts, either abandoning what led us to ask the questions in the first place, or replacing our current physical concepts with ones that are more accommodating to phenomenal consciousness.

Does the same hold if phenomenal concepts are canonically inscrutable or canonically primitive? Perhaps, although here it is harder to say. 'Electron' excludes 'Macroscopic Object' but if Block and Stalnaker are correct then canonical macroscopic concepts like 'Heat', 'Water' and 'Gold' may not be entailed by any other sort of concept (the trick here is that the canonical macroscopic concepts will be positive concepts. 'Not Water' may be entailed by 'Electron' but 'Not Water' is not a canonical macroscopic concept even though 'Water' is). Likewise for Canonical Primitivity: if you hold that concepts like 'Water' are not entailed by any microscopic concepts (however complex) then you may take them to not be posterior to any other sort of concepts. But I take these matters to be complicated, and there is at least some chance that the family of macroscopic concepts is neither canonically primitive nor canonically inscrutable, in which case the Phenomenal Concepts would stand apart if they were.

But of our three notions, it may be that Non-Excludability is the most exclusive. I will now develop one particularly exclusive consequence of a concept's being non-excludable – it implies that that concept is not vague.

I have yet to argue that 'Phenomenally Conscious' is non-excludable. This will be one of the tasks of the second chapter of this dissertation. Very broadly, the idea is that a completely universal panpsychism is conceivable. But there is more to it than that, as I shall explain there.

§2.4 NON-EXCLUDABLE CONCEPTS ARE NOT VAGUE

There are many deep questions surrounding the phenomenon of vagueness. Here, I will assume that vagueness is, very broadly speaking, a semantic phenomenon. Vagueness is a function of the workings of our concepts and predicates and terms – it is a phenomenon that is grounded in features of the inferential roles (broadly construed) of these semantic entities. This is not to say that context or psychology play no role – only that that role meshes with the inferential roles, broadly construed, of the terms involved. I take this to imply that characteristic symptoms of the vagueness of a concept should be available to anyone who knows how to get by with that concept, but not for example that vagueness is necessarily a matter of truth value gaps, or

semantic indecision. But I will also assume those additional things here:
vagueness is a matter of semantic indecision, and this means that vagueness manifests in truth value gaps. Vagueness is a matter of what happens when the inferential role associated with some concept is not sufficiently specific to determine a precise extension for it. In the memorable words of David Lewis,

“The only intelligible account of vagueness locates it in our thought and language. The reason it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word 'outback.' Vagueness is semantic indecision”. (Lewis 1986, 213)

I claim that if this plausible view of vagueness is correct, then a concept that is non-excludable cannot be vague. Recall:

(NON-EXCLUDABILITY): A concept 'C' is non-excludable iff the only concepts that exclude 'C' exclude it negatively.

Where one concept excludes another negatively if it excludes it by negating an a priori necessary condition for it to obtain, or by means of a That's All clause.

The two phenomena most essentially associated with vagueness are Borderline Cases and Sorites Susceptibility. On the semantic indecision view, these phenomena are different ways for concepts to give out. In both cases, this giving out is necessarily a matter of a contrast. In a borderline case, we see

a description that seems to have more or less equal claim to being a C and being a $\sim C$. We see some, but not all of the features that make for a clear case of a C, or we see some respects of similarity with a clear C, and we also see some, but not all of the features that make for a clear case of $\sim C$ or we see some respects of similarity with a clear $\sim C$. And we can tell – if we know how to get on with the concept 'C', and this really is a borderline case of 'C' – that the concept 'C' gives out for this case: the concept was not designed to say whether such a case falls inside or outside of its boundary. Similarly in a sorites series we are given a sequence of cases described in such a way that each of them seems to vary incrementally away from what a clear C is like toward what a clear $\sim C$ is like. If we know how to get on with the concept, we know that the concept will give out if we ask it to determine for us which of these transitions is the one that marks the final case of 'C' and the first case of ' $\sim C$ ': the concept was not designed to make distinctions that fine.

With Borderline cases we grasp a way of being somewhere between a clear C and a clear $\sim C$. With sorites series we grasp a series of gradual changes from a clear C to a clear $\sim C$ by steps too small for the concept to adjudicate on any change between them. But if we have either of these things, then 'C' is excludable. If our only conception of ' $\sim C$ ' is parasitic on the concept 'C' then we are not in a position to recognize something as having elements in common with $\sim Cs$ without judging those things to be $\sim Cs$, and we are not in a position

to recognize some series of changes as headed away from a clear C and towards a clear $\sim C$, without judging those things to be $\sim C$ s. If on the other hand we have ways of conceiving of things as somehow intermediate between C and $\sim C$, then we know how to factor ' $\sim C$ ' into components, or how to identify ' $\sim C$ ' as depicting a condition similar to other conditions that we can articulate in different terms. But this means we have an independent conception of ' $\sim C$.' If a concept is vague, for example, because its clear cases simultaneously meet two conditions (say, 'Being In The North' and 'Being In The East') and its indeterminate cases only meet one of them, then an excluder for that concept just is an excluder for each of the two conditions (here: 'Being in the South' and 'Being in the West'). In fact, an excluder might just be the negations of each of the conditions – this conjunction of negations would be conceptually posterior to each of the concepts it negates, but it would not be posterior to any single concept that states a necessary condition on 'Being in the Northeast' (all conjunctive concepts are excludable).

This is not to say that if a concept is vague, it must be *a priori* how to get by degrees from a clear instance to a clear instance of its negation. There may be ways of getting gradually from a clear instance to a clear instance of its negation that are only knowable a posteriori. But in order to be in a position to recognize such a thing as a genuine possibility you had better have some

positive conception of what it is this gradual series beginning with a C is changing into.³¹

For example, 'Water' is excluded by 'Thick Mud'. It is probably *a posteriori* that mud is a mixture of dirt and water. But it is only if you have some grasp of mud as a different thing than water that you are in a position to recognize a gradual transformation of water into mud as a transformation of a clear case of 'Water' into a clear case of 'Not Water'. Another way of putting the same point: it is only because we have some positive conception of an excluder of 'Water' that we are in a position to recognize something as a mixture of water and non-water. It is only by having some independent grasp on the sorts of things that water is not, that we can have any grasp on how water might potentially be mixable.

It certainly does not suffice that you have some sense of how a given concept can vary in degree, intensity, or some other magnitude. Even in the old days when we thought of life as Vital Force, we could make sense of the force

³¹ Even if supervenience is true (which Semantic Indecision theory does not imply), it certainly is not the case that it must be *a priori* what the admissible complete precisifications of some vague concept are. It is not implausible to ask that there be some *a priori* but partial conception of what makes for an admissible precisification, but this is a weak constraint, requiring only that some of the concept's penumbral connections are a part of its conceptual role. Anyway, I do not take the Semantic Indecision theory to imply Supervenience – especially if the latter is taken to be the view that precisifications are supposed to *extend* the meanings of vague terms while preserving what they already mean, and the idea that the vagueness of a vague term is bound up with what it means. For difficulties with this idea see Schiffer [2003]. For example, it is unclear how this picture could apply in intensional contexts, eg. 'She only likes bald men'

being stronger or weaker in some cases: presumably, the force diminishes before you die. But here we have no conception of how this force might be vague, and presumably you are alive, but weakly, until you have absolutely none of this force. We only get a grip on how 'Life' can be vague when we get a grip on some positive characteristics of not being alive – for example, being unable to process information, reproduce, adapt to circumstances as an organism, grow, etc. I have to have an independent conception of what something is heading towards, as it ceases to be alive, in order for there to be some space in the region between departure and arrival where it is unclear whether the border has yet been crossed.³²

This is not to say that we cannot in some thin sense understand a locution like 'Indeterminately C', where 'C' is a non-excludable concept. We can understand a locution like 'Indeterminately a Prime Number', after all. But if vagueness is in fact something that flows from the nature of concepts – from gaps in the way they delineate the differences between their instances and

³² We are now in a position to compare how my own proposal compares with the proposal of Antony considered in the previous chapter. My own claim is that if some concept 'C' is vague we must be able to conceive of a case as having something in common with a clear C, other things in common with a clear $\sim C$, but without this conception settling the thing as a C, or as a $\sim C$. I do not have to spell out exactly what this means, since it is clear that whatever it means, it will have to involve a conception of ' $\sim C$ ' that is not simply parasitic on the conception of 'C' – for if so then there would be no way of factoring out elements of our conception of ' $\sim C$ ' that did not entail ' $\sim C$ '. Antony has to give us an account of exactly what it is for gradual changes to be gradual in some dimension. Perhaps more importantly, he is committed to there being some element common to Cs, borderline Cs and $\sim Cs$. If there is then presumably there is some independent conception of $\sim C$, but the converse doesn't hold. I need not claim that there is anything common to a conception of C and a conception of $\sim C$.

non-instances – then there is nothing for such a locution to denote. If I have no appreciation of what it takes for a concept to fail to apply, I am in no position to appreciate where the application of the concept gives out. If vagueness flows from the nature of concepts then it is not something that we may only come to know about via empirical discovery, and a minimal constraint on vagueness not being something we may only come to know about via empirical discovery is that we have some independent and a priori (or at least armchair knowable) grasp on what an excluder of the vague concept must be.³³

This concludes the presentation of my argument. I will now consider an objection to it, and then I will conclude with an explanation of why my argument cannot be extended to show that families of concepts that are canonically inscrutable or canonically primitive are not vague.

The objection comes from David Papineau (and my response to it is heavily indebted to Michael Antony 2006c). In *Thinking About Consciousness*, Papineau argues directly for the conclusion that ‘Phenomenally Conscious’ is

³³ Non-excludability may call for relativization. As we noted, ‘Being the number two’ might exclude ‘Phenomenally Conscious’, but this does not exclude ‘Phenomenally Conscious’ from non-excludability, because we can think of non-excludability as relativized to some category restriction. There will be an ensuing constraint on the argument that a non-excludable concept cannot be vague. If the restriction is vague, then the relevant concept may be vague when it comes to the boundary of the restriction. I see no problem here, provided that our restriction is sufficiently broad. For example, if we only restrict the claim of the non-excludability of ‘Phenomenally Conscious’ to concrete particulars, this means that there will be no concrete particular that is an indeterminate case of phenomenal consciousness, and that is interesting enough.

vague or referentially indeterminate. Papineau takes the concept 'Phenomenally Conscious' (he calls it 'Consciousness-as-such') to be radically indeterminate. Papineau is responding to some worries raised by Ned Block in "The Harder Problem of Consciousness." In this paper, Block asks how Phenomenal Realists could ever hope to settle the question of whether superficial functional isomorphs of ourselves (like Commander Data from Star Trek: The Next Generation) are conscious. Papineau seeks to solve this problem on behalf of Phenomenal Realists by holding that there is no fact of the matter.

To justify his claim, Papineau appeals to Naturalistic Theories of Content – in particular to causal or teleosemantic accounts. It is plausible, he says, that the semantics of 'Phenomenally Conscious' (along with most phenomenal concepts) will be given by such an account, and,

"...any causal or teleosemantic account will leave it indeterminate exactly which of the correlated material candidates any given phenomenal concept refers to. For all the correlated material candidates will figure equivalently in the characteristic causes or biological functions of the relevant phenomenal judgments, and so causal or teleosemantic considerations will fail to pick out one mental candidate rather than another as the referent." (2002: 198)

The problem is that causal or teleosemantic accounts often imply indeterminacy where there is none, and this is a problem for those theories.³⁴ One salient corollary problem arises. Vagueness is very plausibly taken to be grounded in concepts, in such a way that symptoms of the vagueness of a vague concept should be detectable to those who are fully competent users of the concepts. But if we take causal and teleosemantic theories to constrain meanings so that concepts only determinately apply in cases where their causal or teleological sources are fully determinate, then they imply the prevalence of a sort of vagueness that is wholly undetectable to fully competent users of the relevant concepts. I note that this should bother even those who reject the theory that vagueness is semantic indecision – those who accept Epistemicism, and some of those who accept Onticism (those who take vagueness to be an ontological phenomenon rather than an epistemic or semantic one) also countenance the idea that indeterminacy has detectable linguistic or conceptual manifestations.

This is not to say that such theories are clearly wrong. But if ‘Phenomenally Conscious’ is vague because its semantics is given by a causal or teleosemantic theory of meaning, then we have larger things to worry about: most of the concepts and predicates we employ are probably

³⁴Cf. Boghossian, “The Rule Following Considerations”, “Naturalizing Content” and Kripke, “Wittgenstein on Rules and Private Language”,

indeterminate in ways we do not expect. It is also worth noting that if no physical description can have phenomenal implications, then the truth of the relevant naturalistic theory of meaning will have to be a posteriori, and it is not exactly obvious that a semantic theory can be a posteriori³⁵. All this is to concede that if there is this sort of referential indeterminacy, then concepts can be non-excludable and also be indeterminate. But this sort of indeterminacy is a far cry from the phenomenon that we ordinarily think of as vagueness.

That concludes my consideration of Papineau's objection. I conclude with an explanation of why the result here does not extend from non-excludability to canonical inscrutability or canonical primitivity. It is commonly assumed that primitive concepts, if there were any, would not be vague. If primitivity were a matter of being one of the basic building blocks in a grand definitional edifice then this might be so: arguably a concept is only vague if there is some more basic level of description along which it is possible to draw distinctions finer than those cut out by the less basic vague concept. But things are less obvious when we do away with the definitional model of primitivity. As I have indicated, someone who thinks that conceptual gaps are

³⁵ Cf Geoff Lee forthcoming. If the semantic theory is *a priori* then it is plausible that the relevant indeterminacies should be available to competent users of the relevant concepts, at least in conjunction with a specification of the causal or teleosemantic facts. But I take it the natural constraint should be that symptoms of the vagueness of a vague concept should be detectable to competent users whether or not they know the causal or teleosemantic facts. But in any case matters are even worse if the theory is *a posteriori*.

pervasive might think that the family of macroscopic concepts is canonically inscrutable, or canonically primitive. But 'Water' is vague, whether it is posterior to other concepts or not. Likewise, I will argue in the next chapter that the phenomenal concepts are canonically inscrutable and canonically primitive. But it is hard to say whether there are non-vague phenomenal concepts other than 'Phenomenally Conscious'. If there are concepts of maximally precise and determinate phenomenal states then these may well be precise. But if we have such concepts we do not employ them often in public discourse. For all we know, the space of such concepts is gunky, in the sense that for every determinate phenomenal concept, there is an even more determinate phenomenal concept.

CHAPTER THREE:

***A PRIORI* PROTOPHENOMENALISM AND NON-EXCLUDABILITY**

It is commonly supposed that Democritus was right, and the rest of the Presocratic and Romantic philosophers were wrong: reality is made of Matter and Void, rather than Sacred Fire, or Aether, or Will, or some other exotic thing. We suppose that the grand debate over the nature of being is closed, and materialists have won.

But there is logical space for materialists to embrace the Dual Aspect view that 'Sacred Fire' and 'Matter' are different guises of the same thing. Many materialists embrace the Dual Aspect view that 'Mind' and 'Matter' are different guises of the same thing. And this sort of materialist may have reason to embrace the more exotic sort of Dual Aspect theory. The view that there is a 'Sacred Fire' aspect to material reality might help to explain the connection between the non-mental and the mental – if the correlation between the 'Sacred Fire' aspect and the 'Mind' aspect were significant enough.

In this chapter I will describe a position, *Protophenomenalism*, which articulates a precondition on this exotic sort of Dual Aspect theory, and I will

tentatively defend it.³⁶ But my primary aim in this chapter is to show that there is a limit on how much of an explanation of the mental we can expect this strategy to deliver. I will call the position which says that some description in exotic terms might actually *entail* a phenomenal description *A Priori Protophenomenalism*. I will argue that this view is difficult to defend in conjunction with a non-deflationary, naturalistic theory of concept application. Since that is the sort of theory of Concept Application you should embrace if you think that the mental is ultimately material, this means that if you think the mental is ultimately material, you may have to abandon the hope for an exotic conception of reality that entails anything about mental reality. This imposes a significant limit on the ways that an exotic Dual Aspect theory might help the materialist bridge the explanatory gap.

My secondary aim in this chapter is to show that this has consequences for our understanding of phenomenal concepts. In the second chapter of this dissertation I characterize three different ways to generalize the idea of a conceptual gap: Canonical Inscrutability, Canonical Primitivity and Non-excludability. Here I will argue that if *A Priori Protophenomenalism* is false, then the concept 'Phenomenally Conscious' is *non-excludable*. I will also present some reasons to think that the family of phenomenal concepts is

³⁶It is only a precondition because, as I will discuss, one could be a Protophenomenalist without being a materialist.

canonically inscrutable and canonically primitive, though for reasons I will discuss these claims are harder to establish. However *non-excludability* is the more newsworthy notion: In the second chapter I argue that if a concept is non-excludable then it is not vague, so the implication of my conclusion here will be that ‘Phenomenally Conscious’ is not vague (assuming a non-deflationary, naturalistic theory of concept application).

In the fourth and fifth chapters I argue that if ‘Phenomenally Conscious’ is not vague then either Panpsychism or Property Dualism is true. There is thus a master argument here against Orthodox Materialism broadly construed: either a non-deflationary, naturalistic theory of concept application is false, or one of Panpsychism or Property Dualism is true.³⁷ I do defend the claim that Protophenomenalism is an option for orthodox materialists (though *A Priori* Protophenomenalism is not), but this defense is modulo my argument in chapter four that Orthodox Materialism is untenable.

In §1 I will introduce *Protophenomenalism* and explain the contrast between *A Priori* and *A Posteriori Protophenomenalism*. In §2 I will argue that *A Priori* Protophenomenalism is incompatible with a non-deflationary,

³⁷If one rejects a non-deflationary naturalist theory of concept application, I take the most compelling alternative to be a primitivist theory of concept application, which is a sort of primitivism about intentional states and relations. Since it is plausible that all mental states, including phenomenal states, are intentional states, this might amount to property dualism about phenomenal properties as well, but in any case it amounts to a property dualism about some mental properties.

naturalistic theory of concept application. In §3 I will argue that if *A Priori* Protophenomenalism is false then the concept 'Phenomenally Conscious' is *non-excludable*, and the family of phenomenal concepts may be canonically inscrutable and canonically primitive.

§3.1: PROTOPHENOMENALISM AND A PRIORI PROTOPHENOMENALISM

Say that a concept is *Exotic* just in case it is not a material concept nor posterior to any, not a phenomenal concept or posterior to any, and not a concept typically associated with any phenomenal concepts or posterior to any that are. I draw here on some terminology from the last chapter. The *material* concepts include concepts of matter and void, concepts of spatiotemporal and causal structure, functional concepts, and some other related ones.³⁸ The *phenomenal concepts* include positive concepts of

³⁸I leave this category vague. Is the Heraclitean concept of 'Fire' a material concept? This may depend. We have several different conceptions of the material. One very austere reading takes material concepts to be concepts only pertaining to the attribute of extension. A slightly less austere reading takes them to be concepts of the objects of an ultimate theory of physics, along with rigid designators that turn out to denote such objects. Finally there is the more permissive view that a concept is material if it is a concept that can be used to think about paradigmatically material objects (see Stoljar 2001). The Heraclitean concept would count as material in this third sense – he must have thought of 'Fire' as characterizing a special substantial form or attribute different from extension. This is not a material concept in either the first or second sense; this excludes 'Mere Matter' and it is not a rigid designator for whatever happens to cause Fiery experiences (it is not even obvious how to formulate the

particular phenomenal states, like ‘Seeing Red’ or ‘Hearing a Trombone’ or ‘Feeling Pain’, and also more general concepts like ‘Phenomenally Conscious’. I also take on board a relation of *Conceptual Priority*. In characterizing the generalizations of the conceptual gap I make heavy duty use of this relation, but here its function is mainly to screen out concepts *defined* in terms of other concepts. I will say more about which concepts are *typically associated* with the concept ‘Phenomenally Conscious’ below.

There might be exotic concepts that have nothing to do with consciousness one way or the other (for example, abstract mathematical concepts) and there might be unsatisfied exotic concepts. Our question is whether there are any exotic concepts with non-empty extensions that correlate in some systematic way with the extensions of phenomenal concepts. Call a theory asserting that there is at least one exotic concept with a non-empty extension that correlates systematically with the extension of some phenomenal concept a *Protophenomenalist* theory. Protophenomenalist theories come in two types: *A Priori Protophenomenalism* adds that some phenomenal concepts are entailed by some exotic concepts, while *A Posteriori Protophenomenalism* denies this (one final piece of terminology: I will speak of

claim that everything is fire if that claim is true only if everything is the sort of thing that typically causes fiery experiences). My own inclination is, if there is any doubt, and any chance that the concept might entail some phenomenal concepts, to think of such concepts as exotic rather than material. Thanks to Charlotte Witt for discussion on this point.

concepts entailing other concepts: one concept 'C' *entails* another concept 'D' just in case the proposition <for all x, if Cx then Dx> is *a priori*).

Protophenomenalist theories trace back as far as philosophy itself. If you think that everything is sacred fire, you probably think that consciousness is some special modification of that sacred fire. One possibility is that the exotic concepts prove to be completely alien – concepts we cannot even triangulate upon, pending some future innovation of thought. But the history of philosophy offers us many alternative conceptions of fundamental reality, aside from Matter: Aether, Air, Bubbles, Class Struggle, Dialectic, Earth, Fire, Force, Geist, God, The Good, the Indefinite, the Infinite, Information, Light, Love, Music, Number, Oppositionality, Power Relations, Representation, Self-Consciousness, Structure, Text, Unity, Vital Force, Water, Will, or any combination of these; for example, Will and Representation, or Text and Information, or Earth, Water, Air and Fire. Of course, any comprehensive theory should give us not only the fundamental principles of being, but also the fundamental principles of formation. If you think that everything is Air, you might take these principles to be Rarefaction and Condensation. If you think that everything is Light, you might take them to be Reflection and Refraction.

Protophenomenalism is a doctrine regarding concepts rather than properties. Not only Materialists, but also Dualists and Idealists may be

Protophenomenalists. Idealists may say that exotic concepts apply to phenomenal properties. Dualists may say that they apply to phenomenal properties, or to material properties, or to both. Also, Protophenomenalism does not automatically entail a Dual Aspect theory, at least not if we individuate aspects in terms of non-entailment. A Protophenomenalism according to which exotic concepts entail both material and phenomenal concepts is a Single Aspect Protophenomenalism. A Dual Aspect Protophenomenalism might take the exotic concepts to entail some phenomenal concepts, but not any material concepts (or to entail some material concepts and no phenomenal ones), then hold that the exotic concepts apply to the very same properties as some material concepts. A Triple Aspect Protophenomenalism would say that the exotic concepts entail neither material nor phenomenal concepts.³⁹ The materialist who hopes for a Protophenomenal way of bridging the gap between material and mental should hope for a Single or Dual Aspect version of the view – it is only against these that I will argue below.

Protophenomenalism goes naturally with the metaphysics of Russellian Monism or Neutral Monism. Russellian Monism is the view that the structural-dynamic and dispositional properties studied by physical science (the t-

³⁹ Nagel's presocratic philosopher, who presciently asserts that matter is energy, without having a conception of how this might be true, was such a compatibilist. In those days the material concepts would have been a narrower collection, and the concept 'Energy' probably would not have figured among them.

physical properties) are somehow tethered to categorical properties (the o-physical properties), which are either identical to or the grounds of phenomenal properties. Neutral Monism is the view that both mind and matter are constructed out of some third sort of thing that in at least some of its basic manifestations is neither mind nor matter.⁴⁰ Protophenomenalism goes very naturally with these views, as it is a precondition of there being a positive characterization of the o-physical properties that are not phenomenal properties, or of the neutral third sort of properties. However, Protophenomenalism is independent of both of these metaphysical views. Protophenomenalism is compatible with an ontology holding that structural-dynamic and dispositional properties are not tethered to categorical properties, and it is compatible with an ontology holding that everything is mental, with an ontology holding that everything is material, and with an ontology holding that everything is both mental and material. Thus, Protophenomenalism may be true even if Russellian Monism and Neutral

⁴⁰ Both Russellian and Neutral Monists may hold that the intrinsic nature of reality is neither mental nor material, but to the Russellian Monist this means that the o-physical properties are the categorical bases of the t-physical properties and ground (construct) the mental properties, while to the Neutral Monist this means that the neutral properties ground (construct) both material and mental properties. Proponents or sympathizers of Russellian Monism arguably include: Mach 1959, Russell 1921, 1924, 1927, Feigl 1958, Maxwell 1978, Lockwood 1989, Stoljar 2001, Chalmers 1996, 2002c, Strawson 1994, 2008. Proponents or sympathizers of Neutral Monism arguably include: Mach 1959, James 1904, Russell 1921, 1924, 1927, Sayre 1976 and Chalmers 1996, 2002c. The two positions converge if you hold that what it is for a categorical property to be the basis (or intrinsic nature) of a dispositional one just is for the categorical property to ground the dispositional one. The challenge for many of these views is to differentiate themselves from Panpsychism or Idealism. A case that no such differentiation is possible is made by V.I. Lenin 1909.

Monism are false. Also, Russellian Monism and Neutral Monism are compatible with it being impossible to adequately conceptualize the relevant o-physical or neutral properties, or anyway with it being impossible to conceptualize them except with concepts posterior to phenomenal concepts. So Russellian or Neutral Monism may be true even though Protophenomenalism is false. But insofar as russellian or neutral monists hope for a positive conception of the o-physical or neutral properties they posit, they hope that Protophenomenalism is true.⁴¹

One worry is that Protophenomenalism is dead in the water because there are not enough exotic concepts to go around. Our concept 'Fire' is material, not exotic. Our concept 'Will' is phenomenal, not exotic. But there may be some concept like 'Manifest Sacred Fire', or 'The Fiery Formal Attribute' loosely related to our ordinary concept 'Fire', but closer to what Heraclitus had in mind, that is not material,⁴² or a concept of 'Pure Will', only

⁴¹ I mean this claim both descriptively and normatively. They should, because if we could not conceptualize the relevant properties that would give us at least some *prima facie* reason to doubt their existence. And they do. See for example Chalmers 2002c.

⁴² As I indicate in a previous footnote, it is difficult to draw a line between material and exotic concepts. Stoljar's distinction between an object-based conception of materialism and a theory-based conception is salient here (Stoljar 2001). A concept like 'Manifest True Fire' would have counted as material on the object-based conception if ordinary fire really did amount to manifest true fire. The same goes *mutatis mutandis* for 'Life' and vital force. Also, I am not claiming that *all* exotic concepts must be compatible (in a single or dual aspect sort of way) with material concepts, only that there is theoretical room for some of them to be. Finally, nothing in my argument hinges on exactly how we demarcate the material concepts from the exotic concepts. So long as the concept clearly does not entail any phenomenal concepts, it does not matter for my dialectical purposes. If the concept may entail some

loosely related to our introspectively derived notion, that is not phenomenal. For example, our ordinary concept 'Life' is perhaps a material concept, but the concept 'Vital Force' apparently is not. In any case, some of the candidate exotic concepts I list above, like the concept of Power, or the concept of God, do not seem to be material, phenomenal or posterior to concepts that are. If there is some account of how such concepts can apply that reserves some neutrality on the metaphysical nature of what they apply to, then Protophenomenalism has a chance. I stress that I mean to take into account here the possibility of concepts that no actual sentient being will ever grasp. My question is whether there are exotic concepts suitably connected to phenomenal concepts that a sentient being *could* grasp (and truthfully apply).

One of the clauses of the definition of an exotic concept is: a concept not typically associated with phenomenal concepts or posterior to concepts that are. Which concepts are typically associated with 'Phenomenally Conscious'? I have in mind here concepts that are not just contingently psychologically associated with phenomenal concepts, but rather concepts that are both contingently psychologically associated with phenomenal concepts and also that entail, or are entailed by, phenomenal concepts. This category may include some concepts of intentionality and normativity : 'Being a moral

phenomenal concepts, classify it as exotic. I hold that phenomenal concepts are not entailed by material concepts.

agent', 'Being a moral patient', 'Being a rational agent', 'Being a Person', 'Having a Perspective', 'Believing that P', 'Desiring that P', 'Perceiving that P', 'Perceiving Redness'. This is not to say that any normative or intentional concept that entails some phenomenal concept is typically associated with it and therefore not exotic. For example, if there is some conception of an aesthetic value that gives us a positive conception of non-consciousness, it is not one that most of us have any grip on, nor is it obvious how we might go about getting a grip on it, and so it counts as exotic.

The category also includes concepts like 'Intensity of Experience' and 'Attentiveness of Experience', that describe potentially necessary conditions for experience. There are concepts like 'Semi-conscious', 'Marginally Conscious State' and 'Subconscious'. And then of course there are cognates of 'Phenomenally Conscious' itself, like 'Qualitatively Experiencing', 'There Being Something It Is Like', etc.

The category of concepts typically associated with phenomenal concepts, and therefore the category of exotic concepts, is parochial. Things that are exotic to us may be commonplace elsewhere. A concept that we do not now but one day will come to associate with phenomenal concepts (and that entails some of them) counts as exotic, since it is exotic for us now. My aim is to capture the idea that an exotic concept would be a concept different from anything that we ordinarily think of as a material (structural-dispositional)

concept and different from anything we ordinarily think of as a phenomenal concept: an exotic concept would reveal some fundamentally different vantage point on reality than the material or phenomenal vantage points.

In contrast, the notions of *non-excludability*, *canonical inscrutability* and *canonical primitivity* are not parochial notions. I will argue in §3 that if *A Priori* Protophenomenalism is false then ‘Phenomenally Conscious’ is non-excludable, and present some reasons in favor of phenomenal concepts’ being canonically inscrutable and canonically primitive. In doing so I will argue that none of the concepts typically associated with phenomenal concepts are counterexamples to any of these non-parochial claims. I turn now to an argument that *A Priori* Protophenomenalism is incompatible with a non-deflationary, naturalistic theory of concept application.

§3.2: AGAINST *A PRIORI* PROTOPHENOMENALISM

I take it to be at least coherent that there is a perceptual way for sentient beings to possess the sort of concepts that *A Priori* Protophenomenalism countenances. *We* may not have the perceptual faculties necessary to acquire such concepts, but other more philosophically insightful creatures might. And it may be that even we may acquire new perceptual

faculties by technological self-modification, or perhaps even by taking drugs (I mean 'perceptual' loosely here: a drug may assist me to acquire an exotic concept by giving me an experience as of something having an apparently new property that I had not experienced before, even though nothing with that property plays any causal role in bringing me to have the experience. For our purposes I consider this sort of experience to be perceptual).

I take the perceptual route to exotic concepts that entail phenomenal concept to be the most promising route. As I will argue below, introspection is unlikely to deliver us any exotic concepts that entail phenomenal concepts, and it is even more doubtful that some more theoretical will help us if the introspective and theoretical routes do not. I conclude that, if you want to defend *A Priori* Protophenomenalism, you should reject Non-Deflationary Materialism about Concept Application (which may mean that you should reject Materialism about at least some Intentional states or relations).

The trouble with introspective approaches is that introspection is a faculty for developing concepts of experiences (or of the objects of experiences or features closely connected to experience) – not of bits and components of experiences, unless those bits and components are experiences themselves. Basically, introspectively derived concepts are *too close* to phenomenal concepts to be exotic. If they entail phenomenal concepts it is because they are phenomenal concepts, or because they are concepts typically associated with

phenomenal concepts.⁴³ One might object that some of the concepts I list above as candidates for being Typically Associated with phenomenal concepts are very close to phenomenal concepts in just this sort of way (remember that as I use it, to be in this category a term must both be in fact typically associated with phenomenal concepts but also must entail some phenomenal concepts). The objection would continue that introspection might uncover future concepts that we will one day categorize as concepts Typically Associated with phenomenal concepts, but today do not. Such concepts might not be concepts of bits or components of consciousness but they might nevertheless be exotic concepts accessing material concepts. I take there to be at least some reason to think that introspection is not going to deliver any genuinely new concepts that entail phenomenal concepts without being conceptually posterior to phenomenal concepts (or the other concepts we currently Typically Associate with them). But I do not have a more decisive argument here, and accordingly (in light of the more principled reasons I am about to give against a Perceptual or Theoretical route) I suspect that a future introspective innovation may be the materialist *A Priori* Protophenomenalist's best hope. However, in the next

⁴³ We can introspect *aspects* of experience – the intensity of the pain, the saturation of the visual field. But this is far from what we are looking for. You do not forge new determinate concepts of experience by combining different determinable concepts of experience. And you certainly do not form a positive conception of non-experience by combining determinable concepts of experience. Schopenhauer takes himself to have introspective access to **will**. But it is precisely for this reason that it is obscure how any non-conscious thing might have **will**. Rather, it is hard to avoid interpreting Schopenhauer as a kind of panpsychist, holding that everything has **will** and therefore is minimally conscious.

section I will argue that such introspectively derived exotic concepts entailing phenomenal concepts would not threaten the thesis that 'Phenomenally Conscious' is non-excludable.⁴⁴

I turn now to the central argument of this section. If one thinks there is a perceptual route to exotic concepts that entail phenomenal concepts, then one must reject Non-Deflationary Materialism about Concept Application ('Concept Materialism' for short). I take Non-Deflationary Materialism about Concept Application to be the view, for some concept 'C', that there is a substantive (non-deflationary) account, in terms of material properties and relations (for example spatiotemporal relations, causal relations, functional relations, information theoretic relations, evolutionary-teleological relations)

⁴⁴ My remarks on the problems with introspective access to exotic concepts parallel his remarks on introspective access to a (concept of a) property that explains the psychophysical nexus. My remarks on the problems with perceptual access are distinct from his. Where he remarks simply that we have no such capacity (because all of our perceptual capacities are of essentially spatial phenomena), I am arguing in principle that no sentient creature could have such a capacity if Naturalism about Concept Application is correct. *A Priori* Protophenomenalism is compatible with McGinn's claim that the property **P** whose grasp will explain the psychophysical nexus, is ungraspable by us. McGinn speaks of *opaque logical supervenience*, and *A Priori* Protophenomenalism is a precondition for what he seems to have in mind. Otherwise if **P** is some material property we presumably do already have a grasp on it by means of some material concept from the relevant science. What we need is the *kind of* grasp on it that lays bare its connection to phenomenology and also to other things. That is, we need an exotic concept that entails both phenomenal and material concepts. In a way then, my own view is at odds with McGinn's, since McGinn in fact relies on the truth of an *A Priori* Protophenomenalism, since he thinks there is some property **P** out there adequate grasp of which would explain the psychophysical nexus, although we happen to be cognitively closed to it. Accordingly, his arguments that we are cognitively closed to *A Priori* Protophenomenalism do not all carry over to arguments that *A Priori* Protophenomenalism is false. E.g., his argument that if we had access to the relevant concepts then we would know what it is like to be a bat (but we cannot know what it is like to be a bat). This is a good argument, but only when 'we' is suitably restricted. Of course some sentient beings could know what it is like to be a bat (imagine a race of rational beings evolved from bats).

of what it is in virtue of which the concept 'C' applies to Cs. For example, a Concept Materialist account of the concept 'Red' might say that 'Red' applies to red things in virtue of those things' having been the normal cause in optimal conditions of the sort of experiences from which the concept 'Red' is acquired. In contrast, it is not Concept Materialism if you say that the full story is that, e.g., 'C' applies to Cs because Cs are in the extension of 'C', or that 'C' applies to object c in virtue of the fact that c is a C. If you hold that this is all there is to say about the grounds of application of a concept C, then you deny that there is a substantive explanation of what it is in virtue of which 'C' applies to Cs. A variant of this approach is to say that there is nothing in virtue of which 'C' applies to Cs. To say any of these things is to embrace Deflationism, and therefore to reject Concept Materialism.

The sort of Deflationism at issue here is not simply the view that **truth** is not a fundamental universal (as even many Metaphysicians will agree), nor simply the view that the grounds of application for different concepts have little or nothing metaphysically fundamental in common (as even many Use Theorists will agree), but rather the more radical view that concepts have no grounds of application in any non-disquotational sense at all. There is a non-factualist reading of this claim ("there are no facts of concepts application") and there is a primitivist reading ("the application relation is primitive and so it is a brute fact that 'C' applies to Cs") but the radical deflationist hopes for

some middle ground between these extremes. I am not sure that there is such a ground.⁴⁵

Another way to reject Concept Materialism is to embrace some sort of Primitivism or Non-Materialism – hold that the concept ‘C’ applies to Cs in virtue of some sort of primitive or non-material states of affairs. For example, one might say that ‘Red’ applies to red things because the concept ‘Red’ *reveals* the nature of manifest redness, and red things are things that really are manifest red (Johnston 1992).

Some care is required since Materialists may use the language of revelation, manifest color, and so on (Byrne and Hilbert 2006). The differentiating feature is whether one takes such language to articulate the final story of what it is in virtue of which a concept applies, or rather whether the facts about what concepts reveal are grounded in further facts about causal relations, information tracking relations, etc.

Having introduced Concept Materialism, I turn now to the central argument. If Concept Materialism is true, then if some perceptually derived concept applies to x, this is grounded in some substantive material relation (for example a causal or functional relation) between x and perceivers of x. But if one concept entails another then we may generally expect the grounds

⁴⁵ Horwich 1998a,1998b, Field 1994a,1994b, Boghossian 1990.

of one to also be among the grounds of the other – or at least that the grounds of the one would metaphysically entail the grounds of the other. This means that if some exotic concept that entails a phenomenal concept is perceptually derived, and Concept Materialism is true, then the phenomenal concept in question will apply to x in virtue of some substantive material relation between x and perceivers of x.

Versions of such a view may be attributable to Hegel, to the later Wittgenstein, and perhaps to Sellars and Brandom and Davidson.⁴⁶ But the claim is deeply implausible, and potentially circular. Say for concreteness that the phenomenal concept in question is ‘Feeling Pain’. Then the implication would be that the fact that the concept ‘Feeling Pain’ applies to x essentially involves some substantive material relation between x and perceivers of x. This is not quite to say that x’s feeling pain is actually grounded in some substantive material relation between x and other experiencers – x’s feeling pain may be identical to some intrinsic state of x’s brain – but the claim is almost as problematic. Concept Materialism does not imply that the fact that some object is red is grounded in that object’s material relations to perceivers (Red might be a surface reflectancy property). But it does imply that the fact that our concept ‘Red’ applies to an object is so grounded. And it is the redness

⁴⁶ Hegel 1977, Wittgenstein 1953, Sellars 1997, Brandom 1994, Davidson 1991. Thanks to Mark Okrent for discussion.

of red – the fact that **red** is the property designated by the mode of presentation given by concept ‘Red’; the fact that **red** has a manifest nature, or anyway a manifest aspect - that makes us want to study the property in the first place. **Red** would likely hold little interest for us if it were not the property designated by our perceptual concept ‘Red’. This is no objection to Concept Materialism. If one is sympathetic to Physicalism about Color, one may reasonably hold that **red** is a categorical property out there in the world⁴⁷ even though the reason it seems special to us is a function of the sort of conscious experiences we have in its presence. Dispositionalists about Color say that colors are relations to our minds. Concept Materialists need not say this, but they preserve the Dispositionalist idea that the applicability of color concepts is a function of the way those concepts, and colors, relate to our minds.

But we ought not say the same thing about conscious experiences themselves: it is surely something intrinsic to my conscious experience that makes it seem distinctive and worthy of study to me, not simply the fact that other people have conscious experiences in the presence of my conscious experiences. The fact that some states of me count as conscious states is not a function of some material relation between me and other people’s conscious

⁴⁷ Relative to our purposes, anyway. A Surface reflectancy property is a disposition to reflect or absorb photons. But it is not a disposition to cause certain mental states (though it may ground such a disposition).

states. This is implausible on its face. *Pace* Hegel, the recognition of the Other is not a precondition of the applicability of the concept 'Phenomenally Conscious'. Even Hegel himself may not have held this view – he held that a sophisticated conception of *self*-consciousness only applies in case of suitable recognition from the Other, but it is not obvious that he held this view also about basic phenomenal consciousness.

The view may in fact be circular; whether this is so depends on how many phenomenal concepts are entailed by perceptual concepts. If exotic concepts only entail a few phenomenal concepts (e.g. they entail 'Feeling Pain' but not 'Seeing Red') then there is probably no circularity, unless the relation grounding the application of the relevant exotic concept is one where it applies to me in virtue of my being in a state that causes you to feel pain. But if there are perceptually derived exotic concepts to entail any phenomenal concept whatsoever, and the applicability of each of these exotic concepts is grounded in some relation to perceivers' conscious states, then circularity looms.⁴⁸

This argument only presents a problem for Concept Materialism.

Consider again the view that some intentional relations are *revelatory*: that

⁴⁸ It also depends on whether the concept applies in virtue of the conscious states of perceivers or in virtue of some other fact about perceivers. Especially when the concepts in question are perceptual, it is plausible that the conscious states of the perceivers are what matter, but this is not inevitable. We may instead hold the critical factors to be the sub-personal processing states. This would allow us to avoid the circularity worry, though not the general implausibility worry. It is implausible that whether I count as conscious could be a function of how I interact with your subpersonal processing states.

perceptual experiences, and the concepts derived from those experiences, reveal the natures of the properties they describe. On such a view there is room to say that an exotic concept of the phenomenal reveals the nature of the phenomenal – the grounds of the concept’s application to some property **P** are simply that **P** is the property whose nature the concept reveals. Concept Materialists are welcome to this language, of course, but they must say in turn that such revelation is grounded in the usual sort of material relation, and then we are back in trouble. A non-naturalist may say that some exotic concept of the phenomenal applies to the things it does because *and only because* those are the things whose nature it reveals, and this is what allows the non-naturalist to countenance *A Priori* Protophenomenalism.

It is important to distinguish concept application from predicate application. Terms in a natural language only apply to things because they mean what they do. And they only mean what they do because of the intentions of their users, the relations between members of the linguistic community, etc. In some sense, then, a predicate only ever applies to something in virtue of some relation between the thing to which the predicate applies and the user of the predicate (or that user’s linguistic community). But

here I am talking about *concept* applicability – what it is in virtue of which the meaning expressed by the predicate characterizes the thing that it does.⁴⁹

Why is the argument restricted to Concept Materialism about Perceptually Derived concepts? It does not apply to introspectively derived concepts because it is perhaps plausible that some introspective phenomenal concepts only apply to the properties they do in virtue of some material relation between those properties and the experiencers who *instantiate* those properties. The implausibility only arises when the relation is between that property (or its bearers) and *distinct* concept users. Also, Concept Materialists may countenance abstract objects (or anyway take there to be concepts for those objects) but they may deny that our concepts of number, or our concepts of logical structure, are grounded in substantive material relations between bearers of the concept and users of the concept. I will argue now that the argument will extend about as far as Empiricism does – Concept Materialism about any non-theoretical Empirical concept implies that that concept applies where it does in virtue of some substantive material relation between the

⁴⁹ Someone might defend a more tenable and moderate deflationism about concept application by holding that what calls for metasemantic analysis is not the grounds of the application of some concept, but rather the grounds of the truth of a judgment or the veridicality of a perceptual experience. One then adds that the truth of judgments (at least of judgments expressed in natural language) always hinges on relational or communal elements, because language does. One concludes that the relationality of the relevant facts cannot have the implausible consequences I say. I respond that there is still a way to the implausible consequences, even if we take the action to be with judgments and perceptual experiences rather than with concepts. We distinguish between cases where the relational elements *exhaust* the metasemantic explanation and cases where they merely *enable* it.

instances of the concept (or the properties of those instances) and users of the concept. The exact extent of the argument is unclear since it is unclear where Empiricism is the appropriate response. Are our spatiotemporal and causal concepts ultimately empirical in origin or do they stem from some special non-relational faculty of rational intuition?

I will now defend the claim that what goes for perceptually derived concepts probably goes for all non-theoretical empirical concepts. The grounds of application of perceptually derived concepts tend to involve *type* relations between experiences and their causal or teleological objects, but in other cases the material relation may be one-off. Some have alleged that names, for example, acquire their meanings in dubbing ceremonies – which is one sort of causal relation between the instance of the concept and concept users – but then others are only in a position to correctly apply the concept to an individual in virtue of the right sort of transmission relations – a different sort of causal relation.⁵⁰ ‘Water’ need not cause a fixed sort of experience in most people who employ the concept, but ‘Watery’ does, and water had to be watery at least in some privileged class of cases. Even before we move to a discussion of more abstract concepts, there is a puzzle over concepts like ‘Square’. Unlike with ‘Red’ or ‘Water’, it is hard to digest the idea that what

⁵⁰ Of course many in the causal reference tradition argue that there are no *concepts* for names, but rather names directly denote, so names may not be an ideal example. In any case, terms derived in this way, for which there are no concepts, would *ipso facto* not stand for exotic concepts.

makes the property of **squareness** worthy of our attention is that the ‘square’ mode of presentation applies to it. Revelatory thinking comes much more naturally when it comes to shape than to color. Or anyway we seem to have two complementary conceptions of shape – both a more direct perceptual grasp and also a more abstract mathematical grasp. There is a puzzle over the relation of the two. This raises the question: do we embrace Rationalism, rather than Empiricism, regarding the origin of our concepts of shape? If Rationalism then there is perhaps room for the materialist to resist the idea that our shape concepts only apply in virtue of some substantive material relation between shaped things and users of shape concepts. But if we embrace Empiricism then our concepts of shape presumably are derived from perceptual experience.⁵¹

So I take it that the argument extends to concepts like ‘Water’ and ‘Square’ insofar as those concepts are Empirical. I do not think it is obvious that Rationalism gets anyone off the hook – this depends on what Rational Intuition is, and what grounds the application of the concepts it delivers. If Rational Intuition resembles a special sort of perceptual experience, then the materialist story about the application of rational concepts may well involve a

⁵¹ One obvious difference is that our experience of shape is cross modal, while our experience of color is not. The Empiricist can argue that the apparent difference between our perceptual and mathematical conceptions of space is actually a difference between uni-modal and multimodal conceptions of space.

substantive material relation between the concept's instances and its users. But it is unclear how exotic concepts entailing phenomenal concepts might be innate or delivered by rational intuition. So there is a problem here for any claim that there are rationally derived exotic concepts that entail phenomenal concepts, and there may also be a special problem for Concept Materialist Rationalists, hinging on what they can tell us about the grounds of application of rationally derived concepts.

I turn now to a consideration of theoretical concepts, concepts like 'Electron' or 'H₂O'. These deserve special treatment because it is not clear that they fall under the scope of my argument regarding perceptually derived concepts. Rather than trying to show that they do, my aim here will be to argue directly that such concepts cannot be exotic concepts that entail phenomenal concepts. Not everyone agrees. Chalmers suggests that we might come to possess new concepts entailing phenomenal concepts via "some sort of theoretical inference from the character of phenomenal properties to their underlying constituents", which he then glosses as "...some sort of inference to the best explanation of (introspected) phenomenology, subject to the additional constraints of (perceived) physical structure." (Chalmers 2002c)⁵²

⁵² Chalmers is not talking about exotic concepts in my sense here, but rather about concepts of what Stoljar calls o-physical properties – the sort of property that a russellian monist thinks serves as the intrinsic nature or categorical basis of the structural-dynamic properties that physical science studies. As I noted earlier, Russellian Monism may be true even if our only

Such a theoretical concept might be based only on introspective phenomenal concepts, in which case its application may not come under the scrutiny of my Concept Materialism argument above. But then such a concept is posterior to those phenomenal concepts (or the other related concepts I am calling "concepts typically associated with phenomenal concepts"), and therefore is not exotic in my sense. More generally I take it to be quite plausible that any theoretical concept that entails phenomenal concepts does so only because it is conceptually posterior to phenomenal concepts or concepts typically related to phenomenal concepts.

We need not understand theoretical concepts on an old-fashioned definitional model to countenance the idea that theoretical terms are dependent in meaning on the meanings of their various observation terms, even if only in indirect ways. Thus 'Electron' is dependent in meaning on 'Electricity', 'Photon' is dependent in meaning on 'Light', and 'H₂O' is dependent in meaning on 'Hydrogen' and 'Oxygen'. One way to try to avoid the implication would be to construe the theoretical information as only pre-conceptually, or even sub-personally, involved in users' grasping such concepts. Burge for example seems to hold that many of our concepts involve sub-personal processing of information in our mental files (Burge 2010).

concepts of o-properties are non-exotic (because derived from phenomenal concepts). However, Chalmers has indicated in conversation that he does hope that the theoretical approach might yield genuinely exotic conceptions of the o-physical properties.

Peacocke speaks of our grasping a concept as based in our grasp of that concept's fundamental reference rule, but where this rule may be couched in terms of pre-conceptual notions rather than full-fledged concepts (Peacocke 2008). But to the extent that the pre-conceptual notions, or the sub-personally processed information, have corresponding concepts that are not identical to the concept whose grasp is in question, the concept whose grasp is in question would be posterior to those other concepts.⁵³

So as concerns theoretical concepts, I take it to be straightforward enough that such concepts, if they were to entail phenomenal concepts, would do so only in virtue of being conceptually posterior to phenomenal concepts, or by being conceptually posterior to some other concepts that entailed phenomenal concepts, in which case we should be talking about them, rather than the theoretical concepts we derive from them. Theoretical concepts do not give *A Priori* Protophenomenalism any foothold.⁵⁴

⁵³This discussion again compares with some remarks in McGinn 1989 (pp. 358-359). McGinn is considering whether or not we might come up with an adequate conception as a theoretical inference from perceptual data about the brain (by which he means data that creatures like us are actually in a position to perceptually acquire). He says, "A certain principle of *homogeneity* operates in our introduction of theoretical concepts on the basis of observation" and he quotes Nagel (Panpsychism p. 183) as saying "it will never be legitimate to infer, as a theoretical explanation of physical phenomena alone, a property that includes or implies the consciousness of its subject." The thought expressed by McGinn and Nagel here is quite plausible. But it does not address the sort of suggestion Chalmers makes, that we might arrive at a conception of some such property as a theoretical inference from a *combination* of physical and psychical data.

⁵⁴ In the next section I will argue that if *A Priori* Protophenomenalism is false then 'Phenomenally Conscious' is *non-excludable*. In Chapter 1 I argue that if a concept is non-

This concludes the argument of this section. We will not derive exotic concepts that entail phenomenal concepts from introspection, since the concepts we derive from introspection are too close to phenomenal concepts. Nor will we derive exotic concepts that entail phenomenal concepts from theoretical inference, because any theoretical concepts that entailed phenomenal concepts would be conceptually posterior to phenomenal concepts or concepts typically associated with phenomenal concepts. Finally, if Concept Materialism is true then we cannot derive such concepts from perception. This is a point of principle rather than of contingent psychology – it is not merely the claim that *we* cannot have perceptual experiences of consciousness because *our* perception is limited to properties under spatial modes of presentation (cf. McGinn); it is the claim that no such perception is possible for any creature whatsoever. I argue that this point of principle extends to all non-theoretical Empirical concepts. This leaves open the possibility of some non-empirical, non-introspective basis for such a concept that even a Concept Materialist could accept. But the sorts of concepts that rationalists usually discuss in these contexts – concepts like space, time and

excludable then it is not vague. But If *A Priori* Protophenomenalism can be false even though there can be concepts like ‘Property that Grounds Both some forms of Consciousness and some forms of Non-consciousness’, then how can this result follow? Why cannot such a concept specify a borderline case of ‘Phenomenally Conscious’? The answer is that, precisely because such a concept is conceptually posterior to phenomenal concepts, it either leaves it epistemically open whether its instances are conscious or not (as in the concept I just expressed), or it explicitly adjudicates on the matter (as with “Property that Grounds Non-consciousness, but is Related to Properties that Ground Consciousness”).

causation – are conceptually divorced from phenomenal concepts. So in the absence of an actual example, or a principled reason to think there will be one, the onus seems to be on the Rationalist to show that any such concept might entail phenomenal concepts. I emphasize again that the argument regarding perceptual and empirical concepts is only an argument against Concept Materialist versions of *A Priori* Protophenomenalism. My argument does not apply to those who hold that perception is fundamentally a matter of Revelation. On that view, it is an open question whether creatures with more refined perceptual capacities than our own might be able to see the hidden commonality between mind and matter.⁵⁵ I now conclude this chapter with an argument that if *A Priori* Protophenomenalism is false then ‘Phenomenally Conscious’ is non-excludable and perhaps also phenomenal concepts are canonically inscrutable and canonically primitive.

⁵⁵I do not think it is obvious one way or the other where the onus then lies. Wherever there is a conceptual gap, there is a question of its full extent. Assuming there is a gap between phenomenal concepts and material concepts, what other concepts are a gap away from phenomenal concepts? Until we have an answer to this question we cannot rule out that it is a matter of principle that the only concepts that are not a gap away from phenomenal concepts are either phenomenal concepts, or concepts that stand no chance of displaying any real commonality between mental and non-mental reality. This might be the case, for example, if the criterion for non-gappiness relative to phenomenal concepts was that a concept display something about first personal point of view (cf Nagel 1974) or that a concept exhibit the tokening structure distinctive of phenomenal concepts, whereby tokening the concept in ordinary circumstances means tokening the property that the concept denotes (cf Block 2006, Chalmers 2006, Levine 2006, Loar 1990, Papineau 2002, Stoljar 2005).

§3.3: IF *A PRIORI* PROTOPHENOMENALISM IS FALSE, THE GENERALIZED NOTIONS OF CONCEPTUAL GAP APPLY TO PHENOMENAL CONCEPTS

Here again are our three generalized notions of conceptual gap, introduced in chapter two:

(PHENOMENAL CANONICAL INSCRUTABILITY): No canonical phenomenal concept is entailed by any other concept of any other sort, except concepts conceptually posterior to canonical phenomenal concepts.

(PHENOMENAL CANONICAL PRIMITIVITY): No canonical phenomenal concept is conceptually posterior to any concept of any other sort.

(NON-EXCLUDABILITY): A concept 'C' is *non-excludable* iff the only concepts that exclude 'C' exclude it negatively.

A canonical phenomenal concept is a positive concept of a more or less specific phenomenal state, like 'Experiencing Red' or 'Feeling Pain.' *A Priori* Protophenomenalism is the claim that there exist exotic concepts (concepts that are not material concepts or posterior to any of them, not phenomenal concepts or posterior to any of them, and not typically associated with phenomenal concepts or posterior to any that are) with non-empty extensions, by which some phenomenal concepts are entailed.

I am assuming that few or no phenomenal concepts are entailed by any material concepts. I assume this also means phenomenal concepts are not entailed by concepts posterior to material concepts. But non-excludability requires that some concepts that are neither phenomenal concepts nor posterior to phenomenal concepts nevertheless entail phenomenal concepts. If *A Priori* Protophenomenalism is false, the only concepts that could fit this description would be concepts typically associated with phenomenal concepts. The question is whether any of these concepts include any positive excluders of 'Phenomenally Conscious'. The Phenomenal Canonical Primitivity thesis claims that no concepts other than canonical phenomenal concepts are prior to canonical phenomenal concepts. Given the material-phenomenal conceptual gap (and assuming that systematic lack of entailment means lack of conceptual priority), no material concept is prior to any phenomenal concept. It follows that if *A Priori* Protophenomenalism is false then the only potential counterexamples are non-canonical phenomenal concepts, and concepts typically associated with phenomenal concepts. The Phenomenal Canonical Inscrutability thesis claims that the only concepts that entail canonical phenomenal concepts are canonical phenomenal concepts, or concepts posterior to them. Under our current assumptions, the only candidate counterexamples are non-canonical phenomenal concepts, and concepts typically associated with phenomenal concepts. I turn now to a discussion of

why that group of concepts does not include any positive excluders of 'Phenomenally Conscious' and why it may not include any concepts prior to canonical phenomenal concepts.⁵⁶

There is dispute over whether concepts of the normative and the intentional entail any phenomenal concepts. Must conscious beings be able to think? Must they be able to perceive? Must what happens to them matter morally? Must they be able to act and be responsible for their actions? If so then 'Not Phenomenally Conscious' may be entailed by concepts like 'Not Capable of Intentional States' or 'Not a Moral Patient' or 'Not a Moral Agent'. But these are negative characterizations, and so will not be counterexamples to our third principle, which states that 'Not Phenomenally Conscious' is only entailed by negative characterizations like these.

Regarding our first two principles the matter is less clear. For example, if you accept the doctrine of Intentionalism, and take it to be an *a priori* truth, you might think that a conception like 'Sensorily Aware of Red' is prior to and entails 'Experiencing Red'. But this is a controversial claim: it seems coherent to think of a very simple conscious being that experiences color without

⁵⁶ It is compatible with my argument in the previous section that there are exotic concepts that entail phenomenal concepts but that do not apply to anything. But such an eventuality would not be interesting, in the sense that such concepts would not help us understand much about consciousness as it is manifested in the actual world. In chapter one I argue that if a concept is non-excludable then it is non-vague. This thesis generalizes slightly: a concept that is only excludable by concepts that have empty extensions at some world, cannot have any vague instances at that world.

thinking. Adverbialists take all experiences to be structurally akin to tickles and itches – mere sensations that involve no fundamental intentional relation. You might also accept *A Priori* Intentionalism, but hold that the relevant intentional concepts are posterior to phenomenal concepts, or equivalent to phenomenal concepts. Regarding moral concepts it seems unlikely that there are any that are prior to canonical phenomenal concepts. Those that entail phenomenal concepts seem to do so at a very abstract level, for example entailing that someone is conscious somehow or other.⁵⁷ And if you think there can be simple beings who only experience color then you might think that things can be conscious without mattering morally.⁵⁸

‘Intensity’ and ‘Attention’ do not seem to be counterexamples to any of our three theses. To be sure, there is a notion of intensity – phenomenal intensity – that seems to be a prerequisite for phenomenal experience. Plausibly, every experience has some amount of phenomenal intensity or other – there is no such thing as an experience that has absolutely no intensity with respect to any parameter (like brightness, loudness, sharpness) at all. But

⁵⁷ Some argue that ‘Pain’ is itself a moral concept (Hewitt 2006). But this is no counterexample to the Primitivity thesis. Canonical Phenomenal Concepts may also admit of other classifications – the thesis claims that no other concepts that are not canonical phenomenal concepts are prior to concepts that are.

⁵⁸ Though ‘Not Being A Moral Patient’ may entail ‘Not Experiencing Pain’ even if it does not entail ‘Not Phenomenally Conscious’. But even this is not obvious: though some theories posit that pleasure is Good and pain is Bad, other theories deny this. Yet others accept it but deny that it is *a priori* – you should find this plausible if you take evaluative concepts to be autonomous.

no concept here will entail any specific canonical phenomenal concept, though some will entail the negations of such concepts, and certainly no canonical phenomenal concept will be posterior to such concepts. This takes care of our first two principles. As for Non-excludability, 'No Phenomenal Intensity' excludes 'Phenomenally Conscious', but it only excludes it negatively. What about 'Attention'? Perhaps every experience must be attended to to some degree or other, and no experience may be completely inattentive. This is far from obvious, but in any event the concept 'Not Attended to Any degree at all' is posterior to the concept 'Attended to Some Degree', and no concept of attention alone is likely to tell us whether some particular canonical phenomenal concept applies, or be prior to such a concept.

Terms like 'Semi-conscious', 'marginally conscious' and 'subconscious' are ambiguous between operational and phenomenal readings. On their operational readings, 'semi-conscious' and 'marginally conscious' indicate states displaying some but not all of the behavior usually associated with phenomenal conscious awareness, and 'subconscious' denotes a state that has some effect on the subject's behavior and function without the subject being occurrently aware of it. These terms may also be used to express more clearly phenomenal notions: 'semi-conscious' and 'marginally conscious' might then mean either a state of very low intensity experience, or a state of oscillation between experiencing and not experiencing, and 'subconscious' might denote

a state which, though it makes no direct contribution to the subject's experience, nevertheless indirectly (perhaps causally) influences the course of that experience. In any case, if notions such as these entail 'Not Phenomenally Conscious', it seems fairly clear that they will do so only through being posterior to 'Phenomenally Conscious'. And it is equally clear that such notions will not be conceptually prior to, or entail, any canonical phenomenal concepts.

Finally we have cognates of 'Phenomenally Conscious' itself: 'Being A Subject of Experience' or 'Being Such that there is Something it is Like to be You'. That the negations of such concepts only negatively exclude 'Phenomenally Conscious' is obvious, as it is obvious that such concepts do not entail canonical phenomenal concepts – concepts of specific positive phenomenal states. It is not completely obvious, however, that these concepts are not conceptually prior to some canonical phenomenal concepts. Does our general concept of consciousness derive from concepts of specific kinds of consciousness, or is the general concept a prerequisite for any specific conceptualization of phenomenology? This is an interesting issue, but nothing really turns on it: the spirit of our second thesis is intact if the only concepts

prior to canonical phenomenal concepts are cognates of 'Phenomenally Conscious.'⁵⁹

I take this to establish that none of the concepts we typically associate with phenomenal concepts are counterexamples to the thesis that 'Phenomenally Conscious' is non-excludable, and I take the case to be fairly strong that they also are not counterexamples to our other two generalized gap theses (with the possible exception of some Intentional concepts). But to conclusively establish the other two theses (Canonical Inscrutability and Canonical Primitivity) we would also have to establish that no non-canonical phenomenal concepts are counterexamples, and this is beyond my scope. Also, to establish Canonical Primitivity we would have to say more about the conceptual priority relation. For all I have officially argued, one concept could be prior to another even though they have no epistemic connection to speak of. But I take the remarks of this section to be progress toward a defense of those two theses.

⁵⁹ I must also consider here the sort of negative abstract concepts that exclude some phenomenal concepts. A phenomenal description will entail some functional descriptions, and perhaps also biological description, economical description, sociological description, etc. The negation of these functional, biological, economical or sociological descriptions will entail the negations of the relevant phenomenal descriptions. But this sort of case does not falsify any of our Generalized Gap theses. Such concepts are obviously not conceptually prior to canonical phenomenal concepts. The phenomenal concepts they do exclude, they exclude negatively, and moreover they only exclude *specific* phenomenal concepts negatively; they do not exclude 'Phenomenally Conscious'.

I conclude by addressing a loophole in my argument in the previous section. I argue there that introspection cannot give us access to exotic concepts that entail phenomenal concepts. But some of the concepts typically associated with phenomenal concepts seem to be introspectively derived (for example, our intentional concepts, our concepts of attention and intensity, and perhaps even our normative concepts). But this category is parochial. What about concepts that we will one day typically associate with phenomenal concepts (and that entail them) but that we have not noticed yet – or for that matter what about concepts that we currently typically associate with phenomenal concepts but that I have overlooked here?

If such concepts were perceptually derived (or anyway were Empirical) then this would be a problem for non-deflationary naturalism about concept application. But if they were introspectively derived, it would not be. Who says that there are no such exotic concepts out there waiting to be noticed (or perhaps already noticed, but overlooked by me)?

I take it to be plausible that there are limits to how many new concepts we may derive from introspection that we have not already encountered. We can make sense of beings with different perceptual machinery from us, and there is a seemingly infinite potential for new perceptual concepts, but this is not true of introspective concepts. But more importantly, even if there are such exotic concepts, we should expect them to not be counterexamples to any

of our generalized gap theses. The point I made in section two – that introspection does not give us access to components or building blocks of consciousness, but only to consciousness itself or to other states very much like it – may not strictly imply that introspection gives us no access to any exotic concepts. But it illustrates why we are not going to introspect a concept of a positive excluder of consciousness. We certainly will not directly introspect a positive excluder of consciousness, as you seem to have to consciously instantiate a state to introspect it, and you cannot consciously instantiate any state that excludes consciousness. The only hope would be to introspect different components of different states of consciousness that when assembled together exclude consciousness. But this is what it seems unlikely that introspection can deliver. Introspectively derived concepts, whether they make my list of concepts typically associated with phenomenal concepts or not, are not going to be the source of counter-examples to the claim that ‘Phenomenally Conscious’ is non-excludable. Likewise, introspection is unlikely to deliver us concepts that are prior to all canonical phenomenal concepts, and if it delivers a concept that entails some canonical phenomenal concept, it is likely that this would be because the entailing concept was an abstraction from (and hence posterior to) phenomenal concepts.

I concede, in any case, that my argument against introspectively derived exotic counterexamples to the generalized gap theses, is less powerful

than my argument against perceptually derived exotic counterexamples. The perceptual argument turns on a matter of principle – of why such concepts would be incompatible with a non-deflationary materialism about concept application. The introspection argument turns only on what seems to be a plausible observation about how introspection works, and what it can deliver, but it is not an argument of principle the way the perception argument is. Really, all I offer here is an argument of onus: the onus is on the friend of introspectively derived exotic counterexamples to generalized gap principles, to say what they are. Readers may wonder how this leaves my argument in comparison with McGinn's argument, which I have discussed in several footnotes, or with Antony's, which I discuss in chapter one. Antony is not concerned with undiscovered concepts at all, and so does not even discuss *prima facie* reasons to doubt that there will be any, be they derived introspectively or perceptually. McGinn does make a *prima facie* case. What I say about introspection is very similar to what he says (though my framework and my aims are different). My argument is more powerful than his primarily when it comes to consideration of perceptually derived concepts. He offers nothing analogous to my argument of principle (and in fact he is careful to say that the limitation in conceptual capacity he discusses is a limitation *for creatures like us* rather than a limitation on all possible conceptualizers).

CHAPTER FOUR:

THE SIGNIFICANCE OF THE MENTAL PROBLEM OF THE MANY

In this chapter I will offer an argument that the property **phenomenal consciousness** is not identical to or grounded in any material property – in other words, for Property Dualism (characterized in a non-modal way⁶⁰). My argument is an argument from the Mental Problem of the Many to Property Dualism. There already are arguments from the Mental Problem of the Many to Dualism (Unger 2006, Zimmerman 2011). But these arguments are arguments for Substance Dualism, and they presuppose Property Dualism (Unger implicitly, Zimmerman explicitly). There is an argument from the Mental Problem of the Many to Property Dualism, but it calls for different premises than the existing arguments employ. My argument hinges on the premise (which I have defended in chapters two and three) that ‘Phenomenally Conscious’ is not vague, and on a premise specifying a sense in which **phenomenal consciousness** is *significant*.

⁶⁰As I define it Property Dualism is compatible with a necessary supervenience claim (although the explanatory role of such a claim would be questionable). I take Substance Dualism to be compatible with (and probably to entail) Property Dualism. I do not think my own arguments establish Substance Dualism but they certainly do not rule it out. Thanks to Mike Raven for discussion on this point. Also, as I define it, Russellian Monism counts as a form of Property Dualism. Though it may count as a version of Physicalism on the object-conception of Physicalism (cf. Stoljar 2001), the view still is Property Dualist in the sense that it posits that the phenomenal properties (the o-physical properties) are neither identical to nor grounded in the structural-dynamic material properties (the t-physical properties). Arguably on this view the phenomenal properties ground the structural-dynamic properties, though not all Russellian Monists need think so.

In §1 I present the Mental Problem of the Many, and introduce the principles that my argument will employ. In §2 I present my argument from the Mental Problem of the Many to Property Dualism. In §3 I conclude by considering ways of rejecting my premises.

§4.1 THE MENTAL PROBLEM OF THE MANY

If Materialism is true then you and I are each composed of collections of material parts. We have organs as parts but also hairs, fingernails, skin cells, individual protein molecules, and so on. However nothing naturally singles out any exact collection of all of these parts over any other. Wherever there is one collection of material parts arranged person-wise, there are many other almost entirely overlapping such collections. Assume that 'Jon' names one of them. But then so does 'Jon#', which names the collection almost exactly like Jon except lacking one of the proteins from the piece of toast I had for breakfast, and containing one of the skin cells that is on the verge of falling off of my finger. And so does 'Jon*', which names the collection almost exactly like Jon except that it contains an extra protein from the piece of toast that I am on the verge of metabolizing, and lacks a different skin cell.

Given the vast swarm of interconnected parts that make up a person, there will be thousands or even millions of person-wise arranged collections of material parts, wherever there is one.⁶¹ And nothing distinguishes any one of these from all the others enough for it to be reasonable to hold that one of them is the true body of the person while the others are not, or that one of them has what it takes to be conscious while the others do not.

The problem I wish to discuss then arises because it is highly implausible that there are many phenomenally conscious beings wherever there is one. I do not almost entirely overlap a host of distinct phenomenally conscious beings, and neither do you. The problem also applies at the level of conscious states: though we may hold that there are disunified experiences realized within a given subject, we do not want to say that each of these is almost entirely overlapped by a host of others that also realizes an experience. But just as Jon is overlapped by Jon# and Jon*, so any collection of neurons or neuronal event N will be overlapped by similar ones N# and N*.

Again, the problem is that once we get specific enough to distinguish between all of the ever so slightly different collections of material parts that are arranged person-wise in the vicinity of every person (or collections of

⁶¹If we believe in temporal parts and special relativity then the number may in fact be infinite, since there will be infinitely many different foliations of spacetime into time-like hypersurfaces, and these will yield infinitely many different but overlapping temporal parts of each person.

material parts arranged brain-state-wise in the vicinity of every brain state), and we see that none of them is metaphysically singled out from the others in any way that physical science can reveal (that is, none is singled out from the others in terms of the structural-dynamic or causal properties that physical science studies), we appreciate the pressure to say that if one of these beings is conscious then they all are. But it is highly implausible that that is so. Let me call the rejection of this implausibility *No Swarms*:

(NO SWARMS): There are not swarms of overlapping experiencers wherever there is one (and in general experiencers do not overlap).

No Swarms is hard to resist. It is not obviously a conceptual constraint –it is perfectly coherent and conceivable that there are swarms of overlapping experiencers. But there is something metaphysically repugnant about the suggestion of Swarms. We know that there are cases of people whose brains are split, and that this may mean that those people embody two different conscious beings. These cases are difficult enough to fathom, but if No Swarms is false then we each share our brains not only with one other conscious being but with thousands, or millions.

Peter Unger suggests that the only way to countenance No Swarms is Substance Dualism. Unger says that it would amount to an unacceptable arbitrariness in nature if one of these beings got to be conscious while all of

the rest did not. What makes the one so special? But this, coupled with No Swarms, implies that none of them are conscious, and hence (since we know that some things, namely ourselves, are conscious) that Substance Dualism is true.

Even if Unger's reasoning is valid, we might object to his appeal to the unacceptability of arbitrariness. Nature is arbitrary and capricious. Why were the initial conditions of the universe just as they were, rather than slightly different? Why do bad things happen to good people? There seem to be violations of the principle of sufficient reason everywhere we look, and so it is difficult to assess the dialectical force of Unger's premise. In fact it looks as though the Substance Dualism Unger argues for will of necessity exhibit some of this arbitrariness in specifying the details of the interaction relation – why does the soul interact with Jon and not Jon*?⁶²

But while Unger's appeal to arbitrariness may falter, an appeal to the nature of **phenomenal consciousness** may succeed. Consider the material properties that differentiate Jon, Jon#, and Jon*. Whatever they may be, none of them will stand out. None will indicate any special sort of causal or

⁶²To address this problem Unger tells a complicated, speculative story about dispositional properties, holding that some things have dispositions that only manifest in the presence of single individual interaction partners (as opposed to manifesting in the presence of anything of some appropriate Type). Compare the fatalist suggestion that the initial conditions are not arbitrary because they were necessary. Perhaps Substance Dualism enables us to minimize the arbitrariness to some degree, but it is not going to rid us of it entirely.

structural threshold. But if **phenomenal consciousness** is a material property of Jon but not of Jon# or Jon* then it must be identical with one of these properties and not the others (or anyway it must be realized by one of these properties and not the others). But this is an implausibility of an entirely different order from the implausibility of the arbitrariness that Unger considers. It is implausible because the property **phenomenal consciousness** is at least somewhat distinctive – there is at least some substantive objective difference between any conscious thing and any non-conscious thing.

It is easy to confuse this principle with Unger's, but they are in fact very different. Note that if Unger's argument is sound it applies to Property Dualists as well as to Type Identity Materialists.⁶³ Unger's principle is that the *material* differences between material beings must be significant if one of them is conscious and the other is not. But the principle I appeal to holds simply that the *differences* between (material) beings must be significant if one of them is conscious and the other is not. If Property Dualism is true then my principle is automatically satisfied, because a conscious thing and non-conscious thing differ over the *sui generis* property **phenomenal consciousness**. We might articulate the principle I appeal to as something like

⁶³In fact Unger's argument may work better against Property Dualists than Type Identity Materialists, since Materialists may reply that identities, being necessary, are automatically non-arbitrary.

(SIGNIFICANCE): The property **phenomenal consciousness** metaphysically stands out. It constitutes a major difference between things that have it and things that do not.

Suppose that the 2784th straw is the one that breaks the camel's back. **Being the straw that breaks the camel's back** constitutes a major difference between the straw that has it and the previous straw. But **being the 2784th straw** does not. And to be sure, being the straw that breaks the camel's back involves more than simply being the 2784th straw – it relies also on the structural properties of camel's back, the gravitational constant, and so on. The analogue of Unger's principle would imply that no straw may be the one that breaks the camel's back. The analogue of my principle would imply only that the property **being the straw that breaks the camel's back** cannot be identified with the property **being the 2784th straw**.

I take differences like the stated physical differences between Jon, Jon# and Jon* (differences of a protein in the stomach or skin cell on the periphery of a finger) to be paradigmatic examples of non-major differences. The principle Significance asserts that the difference between being conscious and not being conscious cannot be grounded in some difference like that. In the next chapter I develop a more general framework for comparing Major from Non-Major differences. The principle advantage of the argument from the Problem of the Many Minds over the argument I give in the next chapter, is

that this argument is independent of the details of the framework I develop there. Here, we only require a grip of the paradigmatic examples of non-major differences; the ordinary physical differences between Jon, Jon# and Jon* in number of skin cells, number of proteins, etc.

§4.2: FROM THE MENTAL PROBLEM OF THE MANY TO PROPERTY DUALISM

So the situation for Type Identity Materialists is that we do not want to say that each of Jon, Jon# and Jon* are conscious – because it is implausible that there are many almost entirely overlapping conscious experiencers wherever there is one – but we also do not want to say that only one of them is conscious, at the expense of the others, because to do so would be to identify the property **phenomenal consciousness** with some material property that only one of Jon, Jon# or Jon* has, and this would violate Significance.

The question is whether there are any material properties with which to identify **phenomenal consciousness** that allow us to uphold both Significance and No Swarms. We have seen that none of the ordinary physical properties differentiating Jon from Jon# and Jon* (properties like **having 2784 skin cells in your left pinkie** will do). But it does not follow that no material properties will do.

A number of proposals to this effect are on offer. I will divide the space of possible responses into three separate groups. The first are those responses that hold that we can countenance Significance and No Swarms without positing the existence of any properties that are not grounded in the ordinary physical properties differentiating Jon from Jon# and Jon*, and without appealing to the vagueness of 'Phenomenally Conscious'. The second are those responses that hold that we can countenance Significance and No Swarms without positing the existence of any properties that are not grounded in the ordinary physical properties differentiating Jon from Jon# and Jon*, but only by appealing to the vagueness of 'Phenomenally Conscious'. The third are those who hold that we can identify **phenomenal consciousness** with some material property while countenancing both Significance and No Swarms, but only by positing properties that are not grounded in the ordinary physical properties differentiating Jon from Jon# and Jon* (or individuals not of a kind with Jon, Jon# and Jon*). I argue that responses from the first group either fail outright, or actually belong to the second or third groups. Responses of the second group cannot succeed if the premises of this chapter are correct - one premise of this chapter is that 'Phenomenally Conscious' is not vague. Regarding solutions from the third group, I argue that Property Dualism may be a more ontologically conservative hypothesis than any of these solutions.

§4.2.1: ONTOLOGICALLY CONSERVATIVE RESPONSES THAT DO NOT APPEAL TO VAGUENESS

One classic response from the first group is Geach 1980's Relative Identity theory. According to this theory, two things may count as distinct relative to predicate F but identical relative to predicate G. The salient application would be to say that Jon, Jon# and Jon* may count as distinct masses or mereological composites, but as the same conscious being. Such a solution allows us to say that there is only one conscious being. However, it also entails that there are distinct masses or mereological composites, each of which is a conscious being. If it is implausible to say that there are swarms of overlapping conscious beings, it is equally implausible to say that there are swarms of overlapping but distinct mereological composite beings, each of which is a conscious being. The Relative Identity solution is no solution to the Mental Problem of the Many.⁶⁴

Another approach from the first group is to say that only the *maximal* collection of particles arranged person-wise is conscious (Sider 2003). Strictly, this is to appeal to a property, **being maximal**, that is not grounded in the

⁶⁴A related response is in Hudson 2001. Hudson argues that each of the composites in the swarm is numerically identical to the others – a different (albeit overlapping) whole location of the same entity, capable of being multiply wholly located. But if it is implausible that there are swarms of conscious beings almost entirely overlapping me, then it is implausible that there are swarms of conscious beings overlapping me that happen to be alternative incarnations of me. Also, the metaphysics of this approach is sufficiently controversial that it may belong in group three rather than group one.

differences between Jon, Jon# and Jon*, but it is at least not a property whose instantiation is metaphysically controversial. The first problem is that this solution would not obviously countenance Significance. Why is being *maximal* any more of a distinction than being second-to-maximal, or third-to-maximal? A more pressing problem is that there may be no maximal collection – there is no guarantee that the fusion of all overlapping collections of particles arranged person-wise is itself a collection of particles arranged personwise: it may be too big.⁶⁵ But the most pressing problem of all is that it is vague which collection is the maximal one. Even if we are guaranteed that the fusion of overlapping collections of particles arranged person-wise is always a collection of particles arranged person-wise, it is vague which collections of particles are collections arranged person-wise. This solution will thus imply that it is vague which are conscious individuals, meaning that the present solution is really in the second group.

It is doubtful that there are any metaphysically uncontroversial extrinsic features that stand any chance of determinately distinguishing one of the person-wise arranged collections of particles from all of the others. But there certainly are no metaphysically uncontroversial intrinsic features that

⁶⁵This is a particularly pressing point if we focus on the spatiotemporal version of the problem (see footnote two).

do so. I conclude that no approach from the first group is likely to countenance both Significance and No Swarms.

§4.2.2: ONTOLOGICALLY CONSERVATIVE RESPONSES THAT APPEAL TO VAGUENESS

Approaches of the second group appeal to the vagueness of 'Phenomenally Conscious' to countenance Significance and No Swarms without appealing to any controversial metaphysics. For if 'Phenomenally Conscious' is vague then we may cut the Gordian Knot that confronts us: we may say that there is only one conscious being here, though it is indeterminate which of Jon, Jon#, or Jon* it is. We may then avoid having to say that one of the insignificant differences between Jon, Jon# and Jon* constitutes the difference between consciousness and non-consciousness (thereby upholding Significance) without surrendering to the absurd conclusion that each conscious being is surrounded by a swarm of other overlapping conscious beings.

It is unclear how exactly to implement the details of such a suggestion, and it may depend on what we take vagueness to be. For example, the suggestion may not succeed if we take vagueness to be epistemic rather than

semantic or ontological. But at least on some versions of the view that vagueness is semantic, there are proposals for how to handle the problem of the many in the event that the concept in question is vague. The details of how the Supervaluationist theory of vagueness accounts for the problem of the many have been worked out in some detail. See for example McGee and McLaughlin 2001, Weatherson 2003.⁶⁶ No such account is perfect, of course, and it might always turn out that it was in the end not possible to exploit the vagueness of ‘Phenomenally Conscious’ to evade the problem. But there is at least the prospect of an adequate solution here.⁶⁷

Unfortunately, the approach is not available to us, because ‘Phenomenally Conscious’ is not vague, as I have argued in chapters two and three. Because ‘Phenomenally Conscious’ is not vague, it is not open to us to say that there is only one conscious being in the vicinity of Jon, Jon# and Jon* though it is indeterminate which of Jon, Jon# and Jon* it is.

⁶⁶Thanks to Phillip Keller for discussion on this point.

⁶⁷There does seem to be a difference between the sort of indeterminate cases that arise in the context of the problem of the many and ordinary borderline cases. Ordinary borderline cases of F tend to be, in some sense or other, on the border between being F and being non-F, exhibiting features of both. In the problem of the many each of the many seems to be a perfectly qualified candidate for being F – the only thing counting against its F-hood is the presence of other overlapping qualified candidates. It seems more plausible to say that these things are *almost* Fs rather than *indeterminate* Fs (Lewis 1993). And we might argue that it is implausible that there are many Almost conscious things overlapping each other whenever there is one. But anyway, since we know that ‘Phenomenally Conscious’ is not vague, we do not have to worry about just how implementable this solution, which would hinge on such vagueness, would be.

This means we must consider responses from the third group: responses that appeal to more speculative metaphysics to countenance both Significance and No Swarms. The most popular such approach is the view that Composition is not Identity, but there are also a few options here that say composition is identity. I will consider these in turn.

§4.2.3: ONTOLOGICALLY INFLATIONARY RESPONSES

To say that composition is not identity is to say that there is some being composed by a collection of particles arranged person-wise but not identical to that collection (or to its mereological fusion). Such a view is typically (but not necessarily) supported by the thought that true Substances may survive the loss or gain of parts, while mere fusions may not. Thus, substances have different (and more interesting) persistence conditions from the things which compose them.

There are different ways of implementing the proposal. First, does only one of Jon, Jon# or Jon* compose a True Substance? If we say yes, it probably has to be a brute fact which of them it is. If we say no, how do we avoid the implication that it is indeterminate where exactly substances are located, or indeterminate what exactly are their parts? One option is to say that it is vague

which of the composites composes the substance. This is still a different view from any in the second group I consider above: here we say there is a substance that is determinately non-identical to any of the composites Jon, Jon# or Jon*, but indeterminately composed of each of them. Another option (suggested to me by Paul Boghossian) is to say that the substance is *multiply composed* – in rough analogy to the way that some properties are multiply realized.

These views all face problems. If the composition relation is not vague, then it probably will be metaphysically arbitrary which of the composites compose the substance. On the other hand there is a strong case (made by Sider 2001) that any vagueness here would have to be ontological vagueness, and there is also a strong case (made by Evans 1978) that there can be no such thing. But the most pressing problem here is that it is hard how a Substance gets to be conscious unless the composite(s) that compose it are conscious. Theorists who deny that Composition is Identity, but who nevertheless call themselves Materialists, owe us an account of what makes them materialists rather than Epiphenomenalist Substance Dualists (in the general sense that there are two fundamentally different sorts of concrete particulars rather than the more specific sense which adds that one sort is material and the other is mental). The usual answer is that the properties of the substance – with the important exception of its persistence properties – are grounded in the

properties of the composites that compose it at that time. But if a composite like Jon* has properties that ground the property of being conscious, then why does Jon* not have the property of being conscious? It does not do to say that the property of being conscious is multiply realizable. A physical property P may ground (realize) functional property F without being identical to that property. But it does not follow that some individual may instantiate property P without instantiating property F.

This objection could of course be advanced against attempts to use Composition Is Not Identity to solve any problem of the many – the problem of the many Cats, or the many Clouds. But one attractive response to the objection (due to Lowe 1995) is that the Sortal Properties in terms of which the problem of the many is usually posed are in fact properties pertaining to the persistence conditions of the substances involved. Having the property of **being a cloud** involves having the ability to survive changes of your parts; an ability that is not shared by any composite of particles arranged cloud-wise.

Perhaps there is some sortal property like **being a conscious being** which is sensitive to the persistence facts in this way. But the property **phenomenal consciousness** (or the property **having phenomenal consciousness**) is an occurrent property, one that apparently has nothing to do with persistence conditions. The usual reply does not work. To explain why the composites are not also conscious, the proponent of the Composition is not

Identity theory must drive a further wedge between the substance and the composites that compose it. This makes the theory look even more like a form of Substance Dualism (albeit, presumably, an epiphenomenalist one). And it undercuts the theory's claim to be already motivated by the need to solve other problems in the theory of persistence.⁶⁸

What about metaphysically speculative proposals that do not deny that Composition is identity? We might restrict composition, holding that the collection corresponding to Jon composes a sum, but the collections corresponding to Jon# and Jon* do not. Alternatively we might hold that Jon, but not Jon# or Jon*, is capable of surviving the change of his parts.⁶⁹ Or, if we are tired of proposals cast in terms of formal ontological properties, we might invoke other sorts of metaphysically contentious emergent properties. For example, we might hold that Jon, but not Jon# or Jon*, has *Elan Vital*. Another option would be to simply take it that one of the material properties differentiating Jon from Jon# and Jon* *turns out* to be a significant property – a hidden sort of higher level naturalness, or something in that spirit.

The suggestion that composition is restricted has been defended (Markosian 1998, Van Inwagen 1990) but it has serious problems – not least of which being that it compromises the popular claim that mereological

⁶⁸Cf Olson 1995.

⁶⁹Cf Casati and Varzi 1999

composition is ontologically innocent (Lewis 1991).⁷⁰ The suggestion that Jon has different persistence conditions from Jon# and Jon* does not obviously help us since it is not obvious what persistence conditions have to do with properties like **phenomenal consciousness**, and again, in a different way it calls into question the idea that mereological composition is ontologically innocent. If we already thought there were such a thing as Vital Force, it might make sense to postulate its involvement with consciousness, but as we do not, the suggestion would be extravagant. In section five, I consider the suggestion that some ordinary physical property that differentiates Jon from Jon# and Jon* (for example the property **having exactly 2784 skin cells on the tip of the left pinky** finger) just *turns out* to mark some major difference between Jon on the one hand and Jon# and Jon* on the other. I argue there that such a suggestion actually amounts to a version of Property Dualism. But even if it does not, it is clearly a highly speculative proposal.

I conclude that none of the proposals from group three obviously have an advantage over Property Dualism. Each of the other proposals involves taking some contentious metaphysical hypothesis and putting it to work even

⁷⁰Van Inwagen 1990 attempts to mitigate some of the harshness of the view by holding that mereological composition relations are vague: it is vague which of Jon, Jon# or Jon* is the one that composes the sum. Sider 2001 argues that this vagueness would have to be ontological, which is problematic enough. The view may also imply that Jon, Jon# and Jon* are indeterminately conscious, which would conflict with our premise. However, my case for that premise hinges on my taking vagueness to be broadly speaking a semantic phenomenon, so the implication that it might be ontologically vague (though highly problematic in its own right) does not actually conflict with my argument in chapters two and three.

more contentious than that for which it was intended. The present considerations seem to directly motivate Property Dualism. Since that is not obviously a less extravagant option than its rivals, we should take it seriously.⁷¹

§4.3: CHALLENGING THE PREMISES

This concludes my argument from the premises No Swarms, Significance and the non-vagueness of 'Phenomenally Conscious' to Property Dualism (or anyway to the view that Property Dualism is no more ontologically extravagant than any of the other theories that countenance all of these premises). It remains for us to consider whether rejecting any of the premises might turn out to be the most plausible response.

My argument that 'Phenomenally Conscious' is not vague hinges on the assumption of a non-deflationary, naturalistic theory of concept application.

One might resist the conclusion of the argument of this chapter by rejecting

⁷¹How would Property Dualism be preferable to the sort of Composition Is Not Identity view we considered just above, supposing we took some more moderate version of Composition Is Not Identity to be called for to handle problems of persistence and the general problem of the many? You could take one of the composites Jon, Jon# or Jon* to be fundamentally phenomenally conscious and then take the Substance to inherit the consciousness of the composite that composes it, just as the Substance inherits the spatial properties of the composite that composes it. This arguably still conflicts to some degree with the spirit of No Swarms, but arguably the proponent of this strategy needs an account of property inheritance to assuage discomfort about this sort of thing anyway.

such a theory. But the sort of deflationism one would have to embrace to escape the result would be a particularly unstable variety, hovering between Primitivism and Non-factualism about the grounds of concept application (see chapter three). So the most natural strategy for resisting the conclusion would be to embrace a non-naturalistic theory of concept application – which means accepting that the intentional relation, or some other relation in which it is grounded, is a primitive (or anyway non-materially-grounded) relation. Given the plausible theory of Intentionalism, this actually implies Dualism about phenomenal properties, but even if we doubt that, we cannot doubt that it implies Dualism about intentional properties, so this would not be much of a defense if the objective is to salvage a Type Identity theory of the mental.⁷²

The only other options are rejecting Significance or No Swarms. I do not see how to reject Significance without in effect changing the subject. If the difference between being phenomenally conscious and not being phenomenally conscious does not mark some sort of objective threshold, I do not know what does. It seems to be a part of the concept ‘Phenomenally

⁷²Also, my argument for the non-vagueness of ‘Phenomenally Conscious’ is an argument that ‘Phenomenally Conscious’ does not have borderline cases in the ordinary sense, or give rise to sorites series. But one might think that the sort of vagueness a concept needs to be amenable to a vagueness-based solution to the problem of the many is different: it has to do with the cogency of a special way of being *almost a C*, and this is different from there being borderline Cs, or there being sorites series for C. I suspect that my argument does extend to an argument that ‘Phenomenally Conscious’ is not vague in this way, and also, more generally, that nothing can be vague in this special way without being vague in the more ordinary ways.

Conscious' that it is significant in at least this respect.⁷³ It is also worth noting a difficult meta-semantic question that arises here. How does the concept 'Phenomenally Conscious' come to select one from the set of insignificantly different candidate referent properties? Some hold that surprising, hidden semantic facts like this abound, usually in the context of an epistemic theory of vagueness (see Williamson 2000) but many take this sort of thing to be highly implausible.

Perhaps we should consider rejecting No Swarms. Those who accept Panpsychism may be amenable to such a suggestion, though it is not obviously a consequence of Panpsychism. But there certainly are precedents for the suggestion. For example, on one interpretation of quantum mechanics, the Many Minds interpretation, there are a potential infinity of minds associated

⁷³But cf. Lee forthcoming for an argument to the contrary. Lee argues for what he calls a *Deflationary Stance* on the question of the significance of consciousness (Block introduces the terminology of *Inflationism* and *Deflationism* in this context in his 2002. For an example of Deflationism similar to what Lee may have in mind see Dennett 1991), and he considers both normative and descriptive forms of significance, as I do (the descriptive form I consider is Significance, I consider normative significance in §4). But Lee's argument is of the hypothetical form: If a certain sort of materialism is true then consciousness is not significant (in each of various senses). I am inclined to agree with something in the vicinity of this hypothetical claim. But Lee does not offer us any explanation of why Consciousness seems to be significant, even if it really is not. There is a difficult dialectical question here: if we must choose between materialism and the principle that consciousness is significant, which should it be? Materialism may be an attractive view, but so is the thought that the difference between being conscious and not being conscious is more substantive than any of the differences between Jon, Jon# and Jon*. We are at minimum entitled to an explanation of why it seems so plausible to say that consciousness is significant, even though it is not. I also note that Lee's argument does not draw on the non-vagueness of 'Phenomenally Conscious', and the form of Descriptive Significance he considers is stronger than the one I consider. **Phenomenal Consciousness** might fail to be significant in the descriptive sense he considers, yet still be Significant in my sense.

with a single body – each registering a different possible projection of the relevant Hilbert space.⁷⁴ And of course there are examples of different overlapping minds, as we see in Brain Bisection cases.⁷⁵ But this is strange territory, and again, it may be more reasonable on the whole to embrace Property Dualism. But for those who toy with rejecting No Swarms, in the next and final chapter of this dissertation I will consider an arguments for the disjunction of Property Dualism or Panpsychism that does not presuppose No Swarms. This argument also will more conclusively rule out competing (Inflationary but not Property Dualist) options. Here I tried to show that Property Dualism is at least as reasonable as other adequate solutions to the problem, but in the next chapter I will argue that Property Dualism and Panpsychism are the only reasonable solutions.⁷⁶

⁷⁴Albert 1993.

⁷⁵Nagel 1971.

⁷⁶Zimmerman 2011 advances an argument similar to Unger's, from the problem of the many collections of parts arranged person-wise, via consideration of phenomenal properties, to Substance Dualism. Zimmerman's argument does not appeal to No Swarms, but he does assume that Property Dualism is true (and he takes this to imply that 'Phenomenally Conscious' is not vague, and that phenomenal properties are fundamental). Zimmerman argues for Substance Dualism much as I do here for Property Dualism – by arguing that it is at least as reasonable as other adequate solutions. I do not think it is obvious that Property Dualism entails that 'Phenomenally Conscious' is not vague (my own argument hinges on the materialist assumption of a naturalistic theory of concept application, but see Goff forthcoming and Simon forthcoming for discussion) or that phenomenal properties are fundamental (see Chalmers 2002c for discussion). However, I suspect that something along the lines of Zimmerman's argument may be advanced even if we relax these assumptions. I agree with him that Substance Dualism is not obviously less reasonable than the other options open to us, though I disagree with Unger that it is the only reasonable solution available.

CHAPTER FIVE:

THE SIGNIFICANCE OF CONSCIOUSNESS

I argue that the property **phenomenal consciousness** is a Natural Dimension (or Sub-Dimension) Determinable – a property demarcating an objective dimension (or Sub-Dimension) of similarity and difference in quality space. I then argue that if **phenomenal consciousness** is a Natural Dimension (or Sub-Dimension) Determinable then either Property Dualism⁷⁷ or Panpsychism is true (or both). My argument hinges on the premise that ‘Phenomenally Conscious’ is not vague, and on a premise specifying a sense in which **phenomenal consciousness** is *significant*.

In §1 I explain what a natural dimension or sub-dimension determinable is. In §2 I introduce my premise that **phenomenal consciousness** is significant and give my argument that **phenomenal consciousness** is a natural dimension or sub-dimension determinable. In §3 I

⁷⁷As I define it non-modally in the previous chapter: Property Dualism is the view that phenomenal properties are not identical to or grounded in material properties. This is compatible with a necessary supervenience claim (although the explanatory role of such a claim would be questionable). I take Substance Dualism to be compatible with (and probably to entail) Property Dualism. I do not think my own arguments establish the disjunction of Substance Dualism or Panpsychism, but they certainly do not rule Substance Dualism out. Thanks to Mike Raven for discussion on this point. And again, as I have defined it Russellian Monism counts as a form of Property Dualism. Though it may count as a version of Physicalism on the object-conception of Physicalism (cf. Stoljar 2001), the view still is Property Dualist in the sense that it posits that the phenomenal properties (the o-physical properties) are neither identical to nor grounded in the structural-dynamic material properties (the t-physical properties). Arguably on this view the phenomenal properties ground the structural-dynamic properties, though this is a very weak sense of ‘ground’.

offer further support for my premise that **phenomenal** consciousness is significant. In §4 I argue that if **phenomenal consciousness** is a natural dimension or sub-dimension determinable then either Property Dualism or Panpsychism is true. In §5 I conclude by considering parallels between this argument and the sort of Conceivability argument advanced by Descartes, Kripke, Jackson and Chalmers.⁷⁸

§5.1: NATURAL DIMENSION DETERMINABLES

I assume that there is a structure of objective similarity and difference in quality space. Just what this assumption amounts to is the subject of much debate, but there is some consensus that we must distinguish between predicates that carve at the natural joints and those that do not, between ‘Green’ and ‘Grue’. Following David Lewis, I will say that joint-carving predicates express *natural* properties.⁷⁹ The only work to which I put these notions is to capture relations of objective similarity and difference. Two things have some objective similarity if they instantiate a common natural

⁷⁸Descartes 1641, Kripke 1980, Jackson 1982, Chalmers 1996.

⁷⁹Lewis 1983a, 1984, 1986. If you think that there is an abundance of properties, you do best to think of naturalness as a sort of selective property of properties (or classes). If you think that properties are sparse, as Armstrong 1988, 1997, 2004 does, you need only distinguish between the predicates that express properties and those that do not.

property, and they have some objective difference if there is a natural property that one of them has and another does not.

Many writers take the starting point of this sort of investigation to be a notion of Fundamental or Ungrounded properties rather than Natural ones.⁸⁰ Each of these notions is confusing, but there is an important reason that I use 'natural' here rather than 'fundamental' or 'ungrounded'. Talk of what is fundamental or ungrounded is entangled with modal talk. It is very hard to think of property P as fundamental, if property P is always necessitated by some other property, the way that **red** is necessitated by **scarlet** or **spin** is necessitated by **up-spin**. But properties like **red** and **spin** carve at the natural joints, and specify respects of objective similarity that their more fundamental determinates may not.⁸¹

Some have maintained that all objective similarity and difference supervenes on the identity or distinctness of maximally determinate natural

⁸⁰Many have contributed to the literature on the objective structure of reality, beginning with the Presocratics, Plato and Aristotle. But for some recent literature see the Lewis and Armstrong papers mentioned above, Bennett forthcoming, Dasgupta forthcoming, Fine 2001 and 2009, Heil 2003, McDaniel 2009, Mulligan, Smith and Simons 1984, Raven forthcoming, Rosen 2010, Schaffer 2004 and 2009, Sider 1995, 2001 and forthcoming, Smith and Simon 2007, Wilson forthcoming.

⁸¹The main pitfall of using the term 'Natural' is that one risks giving the impression of some association with Naturalism. It is my hope that the context makes it clear that this association would be illicit here: a property may be natural in my sense, but nevertheless be spooky, sui generis, and inaccessible to scientific study. Another more minor pitfall is that for Lewis, Naturalness is a matter of *intrinsic* similarity, not just of (objective) similarity. I do not assume this: there can be both natural and gerrymandered respects of extrinsic similarity, just as there can be both natural and gerrymandered relational properties.

properties (properties like **red**₂₇ rather than properties like **red**). On this view, the numerical specification of how many natural determinates I share with you (or perhaps the specification of the ratio of how many natural determinates I share with you, to how many natural determinates I differ with you over) fixes the facts about how objectively similar I am to you. But this picture ignores the sort of objective similarity that we have when we instantiate *nearby* determinates of the same determinable: if I am crimson red, and you are scarlet red. If I am crimson and you are scarlet then we are objectively similar, and we are more similar than we would be if I were green and you were scarlet. There are also facts of exclusion to account for: I cannot be both scarlet (all over) and crimson (all over).

These observations show that there is at least a supporting role for determinable properties to play in the story of natural structure.⁸² Taking the fundamental properties to be a special subset of the natural properties, we may grant that only the maximally determinate natural properties are fundamental. And it may be that these fundamental properties ground or

⁸²The *locus classicus* for the modern discussion of the metaphysical status of determinables is Johnson 1921. Wilson 2011 offers a sustained defense of the fundamentality of determinables, drawing on considerations of objective similarity but also on other arguments. Johansson 2000 offers defends the existence of maximal but not intermediate determinables. Gardenfors 2000 discusses the metaphysics of quality spaces (though he thinks of them as conceptual spaces). Armstrong 1978 denies the existence of determinable universals. Yablo 1992 and Shoemaker 2001, 2007 invoke determinables to do causal work, while Gillet and Rives 2005, and Kim 1992 express skepticism. Stazicker (forthcoming) appeals to determinables to solve a problem in the theory of perception.

determine all of the objective similarity facts. But it would be a mistake to think that the mere facts of identity and distinctness of such properties are all that this grounding or determining requires. Rather, to hold that the determinates do all the grounding, we must think of the determinates as *containing* their determinables. If in making the determinates God had to already have determinables on hand, then of course once he has made the determinates he has no more work to do to make the determinables. To allow that the determinates do all of the grounding is not to say that the determinable facts are conventional, or that talk of determinables is merely talk of disjunctions or higher order quantification over determinates.

Determinates may be the onstage properties that interface directly with particulars, but their determinables are the backstage coordinators who make the production possible. To the extent that Fundamentality has to do with being an ungrounded grounder, an unnecessitated necessitator, we may deny that determinables like **red** and **spin** are fundamental. But we must hold that some determinables are natural (or that some predicates pick out true determinables while others do not) if we think that **scarlet** is more like **crimson** than it is like **mauve**.⁸³

⁸³ There may be reasons for holding that determinables really may not be subordinated to their determinates, reasons for thinking they amount to more work that God had to do. One such reason is that determinables may play a causal role that their determinates do not (cf. Yablo 1992, Shoemaker 2001). Another reason has to do with the possibility of Gunk. You might think that it is possible for objects to fail to have maximally determinate properties. You

The upshot of this is that the natural determinates come organized into natural dimensions of quality space, dimensions of comparative similarity and of exclusion. Some determinables will mark regions in such spaces, while others, the maximal determinables, will mark the dimensions, or the spaces, themselves. The general picture will be of a tree with entangled branches: the dimension determinable (say, **color**) at the top, and then a series of levels of overlapping region determinables (**warm color, dark color, reddish color, greenish color, red, orange, purple, blue,...**) in the middle and then finally, if there are any, the maximal determinates (**red₂₇, orange₁₆, blue₃₈**) at the bottom. There may be reasons to privilege some of the determinables over the others, depending on whether they carve natural joints within their quality space. Though **color** may specify a natural space, it is not obvious that **warm color** carves at a natural joint in that space.⁸⁴

might think this on general theoretical grounds – you might think there can be no such things as maximally determinate properties – or you might think that some objects sometimes do have them and some objects sometimes do not. For example, you might think this because you think that time can be mereologically gunky (there are increasingly short moments of time, but no truly instantaneous points of time) and you think that continuous change is possible. (but cf. Arnztenius and Hawthorne 2005, who draw a different conclusion from this sort of possibility). This possibility deserves elaboration. It is compatible with, but does not imply, that an object may fail to be fully determinate (where this would mean that for some object there is a property that is the most determinate property (in some dimension) that it instantiates, even though there are more determinate properties in this dimension that it does not instantiate).

⁸⁴Cf Johansson 2000. For Johansson the question is whether intermediate determinables exist, rather than whether they are natural, but presumably in a sparse property framework existence means naturalness (although cf. Schaffer 2004). On the other hand, suppose you think that **green** is *really* a mixture of **yellow** and **blue**, but **yellow** is not really a mixture of **green** and **orange**. You might countenance this by holding that **yellowish** and **blueish** carve at the joints but **greenish** and **orangeish** do not.

We may ask questions about the structural properties of such quality spaces.⁸⁵ Does a given space have topological or metric structure? Is it continuous or dense? Does it divide into non-path connected subspaces? Is the space bounded or unbounded? I will not go into any of these questions in depth here, but it will be useful to distinguish between a Dimension determinable and a Sub-Dimension determinable. Suppose we have some quality space like **spatiality**. Now there are many different sorts of spatial properties and they all have spatiality in common, but one-dimensional spatial properties seem to be different in kind from two-dimensional spatial properties. We might think of properties like **one-dimensional spatial object** as sub-dimension determinables, to acknowledge that the regions of quality space that they designate are not only natural but also separated from other regions of the same quality space.

I have used color space as an example of an objective quality space. This is for heuristic purposes. It may be very difficult to determine which are in fact the objective quality spaces – just as it is difficult to determine which are the fundamental determinate properties. They may actually be thin on the ground. For example, it may be that the only natural dimension determinables instantiated at our world are things like **spatiotemporality** and **causality** and **mass** and **charge** (or perhaps only something even more general than this, for

⁸⁵ See for example Gardenfors 2000.

example a characterization of pathways in some multidimensional phase space). There is also the skeptical worry that 'Grue' in fact corresponds to a natural region in some quality space, while 'Green' does not (cf. Goodman 1955). But I take it that we have various paths to justification for beliefs about which things are objectively similar to which.⁸⁶ It is certainly a challenge to say what these are, but the enterprise is far from hopeless.

§5.2: SIGNIFICANCE AND THE SEA CHANGE ARGUMENT

In the previous chapter I invoke a principle, Significance, that speaks of a property constituting a *major difference* between things that have it and things that do not have it. In this chapter I develop another argument from that premise to metaphysical conclusions. There, because the context made things clear enough (the ordinary material differences between Jon, Jon# and Jon* are clearly *not* major differences), I left the matter intuitive. Now I will have need to say more precise things about the comparative difference between properties and their instances.

⁸⁶One area where the matter gets tricky: debates over the nature of Functionalism. Do our functional predicates ('Mousetrap', 'Adding Machine' etc) carve at the joints or not? Block 2002 and Kim 1992 hold that in general they do not, while Shoemaker 2001, 2007 holds that they sometimes do.

Having done so, I will present my first argument that **phenomenal consciousness** is a natural dimension determinable. Strictly, my argument is that **phenomenal consciousness** is a natural dimension or sub-dimension determinable, but for stylistic purposes I will suppress this detail sometimes in what follows. I take it to be quite an accomplishment to be a natural sub-dimension determinable, and my argument in the final section will be that if **phenomenal consciousness** is a dimension or sub-dimension determinable then one of Property Dualism or Panpsychism must be true.

Though I take it that we have some beliefs about which properties are natural dimension determinables, for the most part these concepts are too abstruse to figure in the beliefs of all but a few select metaphysicians. But we do have beliefs about objective similarity and difference, and the abstruse framework is only there to regiment our vocabulary for articulating this sort of belief.

I claim that we in many cases have the sort of belief that my principle Significance (from the previous chapter) articulates: that some property specifies a major objective difference between things that instantiate it and things that don't. What is a major difference? A major difference is not merely an objective difference – even the most inconsequential difference, such as the difference between Jon and Jon#, will presumably be grounded in some fundamental difference or other. A major difference is a difference over more

than just an adjacent (or very close) determinate or two. The difference between **red**₂₇ and **red**₂₈ is not a major difference, even though it is an objective difference.

There is a very special circumstance in which two things may seem to differ only over a single determinate, but still the difference is a major difference. This is the case when some quality space is bounded, and the loss of the determinate in question takes the thing from having a peripheral or boundary determinate from the space, to not having any determinate from that space. Imagine that something shrinks in size, until finally it no longer occupies any space at all. The difference between being extremely small, and not occupying any space at all, is in one sense only a difference in a single determinate (we may say it is only a difference of a Planck length or two) but in another sense it is a sea change – something that before was spatial, now is non-spatial. The same goes for a change between sub-dimensions. Something that is almost perfectly one dimensional may only differ from something that is perfectly one dimensional by a little bit of spatial extent, but this difference is nevertheless a sea change (if you doubt this, then you doubt that **one dimensional spatial object** really does specify a natural sub-dimension

determinable, not that ceasing to instantiate some sub-dimension determinable is a sea change).⁸⁷

With this in mind we may say that a major difference is a difference that is not a minimal difference within some natural quality space – a difference of at least several determinates within some natural quality space, or a difference between having and not having a property belonging to such a space. Formulated in this matter, the phenomenon of a property being Significant – that having it involves a major difference from anything that does not have it – is quite typical. There are a great many properties that seem to be significant: properties like **being spatiotemporal, having mass** but also properties like **being intelligent, being alive** and perhaps even properties like **being tall, being rich**.

In each case, there seems to be a major difference between anything that has the property and anything that does not – a difference of either at least a few natural determinates, or a difference that seems to be a sea change. Even if things may be more or less alive, the difference between being alive and not being alive seems to be a sea change.

⁸⁷Cf the distinction Johansson 2000 draws between Lack of Resemblance and Being Very Dissimilar. Two properties Lack Resemblance if they are not determinates of a common determinable, while two properties are dissimilar if they are determinates of a common determinable but not very close to one another in the relevant property space.

It is hard to deny that **phenomenal consciousness** is significant in this sense as well – the difference between being conscious and not being conscious is a sea change, or anyway it is not simply a matter of moving one determinate down the line. But as I have just indicated, this sort of Significance is rather common. How do we get from the Significance of **phenomenal consciousness** to any deep metaphysical conclusions?

Even though there are many cases where it is correct to assert that a given property is Significant, there actually are not that many properties out there that are Significant. Determinates, by definition, are not Significant, but neither are region determinables: if a region determinable is determined by some fixed set of determinates, something fails to have the region determinable by being one determinate down the line from the last determinate to determine that determinable.

How then can it be correct to say, 'Life is significant' or 'Intelligence is significant' or 'Being tall is significant'? The answer is that these terms are vague. This means that between any clear instance and any clear non-instance there is a no-man's land of borderline cases. If we take vagueness into account we may reformulate Significance:

Property **p** is *significant** iff there is a major difference between any clear instance of 'P' and any clear instance of '~P'.

I use 'clear' here as a substitute for 'determinate' to avoid confusion, since in this discussion the term 'determinate' is already in use. 'Clear' does not stand for any further restriction. If the predicate 'P' is vague then it may be slightly misleading to speak of the property **p** since at least on most semantic accounts of vagueness, a vague predicate does not correspond to any single property; instead there are a number of precise properties and the predicate indeterminately corresponds to each of them. But for ease of formulation I will sometimes speak in the material mode, unless the situation calls for clarification.

I claim, then, that a great number of properties are significant*. In almost every case, however, this significance is a function of vagueness. Properties like **being intelligent**, **being alive**, or **being tall** are not natural dimension determinables – but the predicates 'Being Intelligent', 'Being Alive' and 'Being Tall' are vague. But in the case that a property is significant* but not vague, then that property must be a natural dimension or sub-dimension determinable.

As I argue in chapters two and three, 'Phenomenally Conscious' is not vague. This implies that if the concept is not defective – if it denotes a property – then the fact that **phenomenal consciousness** is Significant* implies that **phenomenal consciousness** is a natural dimension (or sub-dimension) determinable. Strictly speaking, it does not quite imply this: strictly what it

implies is that the difference between anything that is a clear instance of **phenomenal consciousness** and anything that is a clear instance of **no phenomenal consciousness** is a difference of some natural dimension (or sub-dimension) determinable. But in such a case the most conservative theory is that the property in question is itself the determinable over which its instances all vary from its non-instances (or alternatively that it is a Boolean compound of such determinables: a Cartesian Product of dimension determinables).⁸⁸ I conclude that since 'Phenomenally Conscious' is not vague, and since phenomenal consciousness is significant*, phenomenal consciousness is a natural dimension determinable.

§5.3: NORMATIVE SIGNIFICANCE

In the previous section I consider a sort of *descriptive* significance that a property can have. I do not present any direct argument for the claim that **phenomenal consciousness** is significant*, though I do suggest that the claim is very intuitive and that indeed many properties are significant*. I now turn my attention to ways that a property can be *normatively* significant. I will

⁸⁸ The alternative would be to say that different kinds of difference between **consciousness** and **non-consciousness** involve differences over different dimension determinables.

consider a few ways that **phenomenal consciousness** is normatively significant and argue that these imply that it is significant*.

Phenomenal consciousness matters both morally and epistemically.⁸⁹

Many debates about the ethical significance of animals hinge on whether animals are in fact conscious.⁹⁰ We have here the idea both that consciousness is a precondition for being a moral patient, and the idea that **consciousness** has some distinctive sort of intrinsic value. Certainly, it is hard to resist the idea that **pain** is bad, or the idea that **love** is valuable in itself.

One might hold that though these specific states are morally important, there is no importance to **consciousness** as such. But it seems, to the contrary, that **consciousness** as such figures critically in the explanation of why each of these more specific states matter in the way they do. It is difficult to identify any less general feature that all of the states that seem to matter in this way have in common, and it is difficult to think of a conscious state that clearly does not matter at all – there is something normatively serious about extinguishing a conscious light, however dim it may be. Here the suggestion is not only that the *capacity* for experiencing matters morally, and so

⁸⁹ But again for a limited defense of the denial of this claim see Lee forthcoming. Lee argues that though there may be a special sort of normative significance associated with **consciousness**, there may be some analogous sort of thing, *schormative schmignificance*, associated with a certain sort of non-conscious being; for example, a non-conscious superficial functional isomorph of a conscious being. But so long as Lee does not deny that consciousness is normatively significant, he does not deny the central claim I make in this section.

⁹⁰ Cf Thomas Huxley 1874.

extinguishing that capacity matters morally, but also that whether or not something is occurrently experiencing matters morally (or if we want to talk about states rather than experiencers, whether or not a state is an experiential state matters morally)

Consciousness also seems to be epistemically significant: according to some, the justification for our epistemic attitudes may depend on our being conscious in forming them.⁹¹ Here again, one might hold that it is only specific states that are epistemically important. In this instance there is a natural candidate for the more specific feature common to all conscious states that matter epistemically – being a perceptual state. However, on a very plausible theory of **consciousness**, the Act-Object view, all conscious states are metaphysically of a kind with perceptual states, and so potentially suited to play the relevant sort of justificatory role. Also, it is plausible that some non-perceptual states play a justificatory role, for example, my headache may justify me in believing that I have a headache.⁹²

How do we get from these normative observations to a metaphysical conclusion? The answer is that normative facts not only supervene on descriptive facts, but also a constitutive presupposition of normative inquiry is

⁹¹ Cf Declan Smithies 2006 and Lee *supra*.

⁹²Cf Pryor 2000.

that normative facts are not brute facts.⁹³ Supervenience alone does not guarantee us much, but the constitutive presupposition of normative inquiry seems to be that if there is a major normative difference between two cases then there had better also be a major descriptive difference between those cases. It may sometimes be difficult to say what makes a normative difference a major normative difference, and so in practice this principle may not always give us a useful decision procedure. But we may take it that the difference between having and lacking the normative status that consciousness bestows is a major normative difference. As such, it must supervene on a major descriptive difference.

Phenomenal Consciousness is not the only property that is normatively significant. **Life** may be normatively significant, as may be **Intelligence**. But as we saw in the previous section, there is no difficulty in saying that those properties are significant*. We may say this even though we do not believe in *Elan Vital* or Intellectual Entelechy, because 'Life' and 'Intelligence' are both vague, and it is easy for a property to be significant* if the predicate that canonically expresses it is vague (so to speak⁹⁴).

⁹³ Cf Dworkin 2011, Nagel 1986.

⁹⁴ Remember that strictly speaking, Significance* is a characteristic of predicates rather than properties; I do not presuppose here that there are vague properties. But it is useful to talk that way. Cf Schiffer 2003.

So we have that if a property is normatively significant then it is significant*. But as we saw in the previous section, if a property is significant* and non-vague, then it is a natural dimension or sub-dimension determinable. **Phenomenal Consciousness** is normatively significant, and I have argued that 'Phenomenally Conscious' is non-vague. It follows that **Phenomenal Consciousness** is a natural dimension or sub-dimension determinable.⁹⁵

§5.4: IF PHENOMENAL CONSCIOUSNESS IS A NATURAL DIMENSION (OR SUB-DIMENSION) DETERMINABLE THEN PROPERTY DUALISM OR PANPSYCHISM IS TRUE

With which material properties might we identify the property **phenomenal consciousness**, if **phenomenal consciousness** is a natural dimension or sub-dimension determinable? Which material properties specify an entire (sub)dimension of similarity and difference? This is an elite category. **Spatiality** seems to be such a property, as do **temporality** and **causality**, though many take the ultimate natural category to be **spatiotemporality**

⁹⁵ Another route to this conclusion would evaluate whether something could be an indeterminate case of the normative status bestowed by the relevant property. If we find, for example, that nothing can be an indeterminate case of the normative status associated with **being alive**, then the significance* of **being alive** that comes due to the vagueness of 'Life' may not be adequate to countenance the normative significance of life – we might then have reason to rethink Vitalism. Someone who was unpersuaded by my argument that 'Phenomenally Conscious' is not vague might nevertheless be persuaded by the argument that there are no borderline cases of 'Having the Intrinsic Value that Consciousness bestows' or 'Having the capacity for epistemic justification', and this might be another route to our conclusion.

rather than **spatiality** and **temporality**, or anyway they take **spatiality** and **temporality** to be sub-dimensions of **spatiotemporality** (or they take **spatiotemporality** to subsume into some yet more general space, e.g. some multi-dimensional phase space).

We might identify **phenomenal consciousness** with a sub-dimension rather than a dimension. This widens our candidate pool a little bit. We may consider **being the sort of thing that can collapse a wave function**, or **being three dimensional**. But this is not much of an improvement. We may also allow that **phenomenal consciousness** is actually some sort of Boolean compound, or Cartesian product, of natural dimensions – for example **being three dimensional and being the sort of thing that can collapse a wave function** but this again will not widen our candidate pool very much.

It is important to not confuse the claim that **phenomenal consciousness** is a natural (sub) dimension determinable with the claim that **consciousness** is some kind of causal threshold – the neuron that completes the circuit. There are theories on which **phenomenal consciousness** involves such a causal or structural threshold, for example the 40 Hz theory, or the Orch-Or Theory, but these theories do imply that **phenomenal consciousness** is not a natural dimension (sub) determinable. **40 Hz** does not differ over any

dimension or sub-dimension of real quality space from **39 Hz**.⁹⁶ The Orch-Or theory says that **consciousness** arises in virtue of some sort of synchronicity between rates of neuronal firing, and the decoherence time of quantum processing inside neuronal cell walls. No property along these lines is even remotely close to being a dimension of quality space. I take it that if we are looking to physics, our best bet would be some feature associated with the power to collapse wave functions. This is arguably a basic determinable dimension or sub-dimension of **causal power**, and the collapse of wave functions does seem to have something to do with **phenomenal consciousness**.⁹⁷ But there are several reasons why such an identification would be problematic. For one thing, on some respectable theories, like GRW⁹⁸, very small systems can collapse wave functions, so the proposal might amount to something in the vicinity of Panpsychism.⁹⁹ For another thing, even

⁹⁶The 40 Hz and Orch-Or theories are in fact both couched in vague language, but we can imagine perfectly precise modifications. For more on Gamma Wave (40 Hz) theory, see Crick and Koch 1990. For more on Orch-Or see Penrose 1989 and Hameroff 1998. For an argument that Orch-Or is empirically inadequate see Tegmark 2000.

⁹⁷Another example might be Tonini's Phi. If we hold that any amount of Phi corresponds to some amount of consciousness, then a version of Panpsychism follows, but there is also a question of whether Phi really describes a proper dimension of quality space. On the other hand, the suggestion that there is some non-zero value of phi that is the minimal value giving rise to consciousness is in violation of the principle that **phenomenal consciousness** is a natural dimension determinable.

⁹⁸ For more on the Ghirardi-Rimini-Weber interpretation of quantum mechanics see for example Ghirardi, Rimini, Weber 1985 and Albert 1993 .

⁹⁹ Panpsychism also follows without assuming GRW if we conclude, as we perhaps should, that the relevant sub-dimension is something like **being potentially causally involved in the collapse of a wave function**. All more selective properties in the vicinity may be determinates of this one. But this is a very easy property to instantiate, barring some appropriate natural joint in the vicinity of **causal involvement**.

if it does not amount to Panpsychism it might run us afoul of the mental problem of the many since if one complex physical system has the power to collapse wave functions then presumably so does any almost entirely overlapping complex physical system. Finally, an identification leaves you with two mysteries: on the one hand the Explanatory Gap (why is **phenomenal consciousness** identical to the ability to collapse wave functions, rather than something else?), and on the other the question of why exactly these things, rather than other things, have the ability to collapse wave functions. It is no explanation to say that they do so because they are conscious: identities do not explain in this way. Property Dualism seems the more sensible route at this stage: we avoid the explanatory gap, and we get an *explanation* of the physical laws (these things have the ability *because* they are conscious) – and in addition we find something nomologically useful for consciousness to do. But in any case, that **phenomenal consciousness** is a natural (sub) dimension determinable means far more than that **phenomenal consciousness** is some kind of causal threshold phenomenon – so if we seek to identify it with any sort of physical phenomenon, we seem to have to do so in a way that makes something in the vicinity of Panpsychism true.

It is important not to confuse the intuition that **phenomenal consciousness** is a natural dimension determinable with the easier-to-countenance thought that conscious things are distinguished from *nearby* non-

conscious things.¹⁰⁰ One way to succumb to this temptation is to take something's being a causal threshold to make it an adequate reduction base for **phenomenal consciousness**, but another way is to take something's marking a sort of *metaphysical* threshold to make it an adequate reduction base. It is tempting to countenance the difference between conscious individuals and non-conscious individuals by combining some sort of requirement of physical complexity, with some extra metaphysical component. For example, as we discussed in section one, you might say that mereological composition is restricted, and that **phenomenal consciousness** is physical condition **C**, but being the sort of mereologically structured entity that you have to be to instantiate **C** already makes you special: most arbitrary collections of particles do not compose a sum. Or you might be a universalist about mereological composition, but take there to be a distinction between the persistence conditions of true substances, which can survive the change of parts, and mere junk composites, which cannot. Then you might say that consciousness is the physical condition **C** for true substances – in other words, it is the composite property of **being a true substance and being C**.

¹⁰⁰We might characterize the claim that **phenomenal consciousness** is a threshold property as the claim that **phenomenal consciousness** is *locally* significant: clear cases of 'Phenomenally Conscious' exhibit some major difference from all *spatiotemporally nearby* clear cases of 'Not Phenomenally Conscious'.

But our data is not merely that conscious things stand apart from non-conscious things, it is that **consciousness** is its own (sub) dimension in some objective quality space. Simply taking mereological composites to be sparse on the ground does not make an ordinary physical property like **C** into a dimension of quality space – nor the composite property **being C and being mereologically structured**. The property of **being mereologically structured** is perhaps its own dimension of quality space, and it is accordingly open to us to say that being conscious is being mereologically structured. Similar remarks go for **being a true substance**. But in either case, we inherit the worst of both worlds. **Phenomenal consciousness** gets identified with something we understand in an entirely different way, giving rise to the Explanatory Gap, and this something else is either strongly emergent and ungrounded in any ordinary material features (and so no real gain from the standpoint of parsimony over Property Dualism), or something in the vicinity of Panpsychism is true.

A final question is whether there might be hidden dimensions of objective similarity and difference: whether, for example, the property **C** in our previous examples might just *turn out* to be a natural dimension determinable, even though it does not seem to be, and our canonical physical concepts of this apparently ordinary physical property do not disclose that it is.

This suggestion is different from the more plausible suggestion that some high-level or complex properties might nevertheless be natural. That suggestion, charitably construed, is just the suggestion that determinable properties (region determinables or dimension determinables) can be natural, and as I have argued, something in the vicinity of this thought has to be correct. The present suggestion compares better with the suggestion that it turns out that **grue** is more natural than **green** or **blue**. But it is actually less plausible even than that suggestion.

It is even less plausible because, when it comes to determinable properties, we *individuate* those properties in terms of the similarity relations they impose on the relevant class of determinates. It is not an option to say that the fact that **red** is more similar to **orange** than it is to **green**, is the same fact as the fact that **the sound of a trombone** is more similar to **the sound of a sackbut** than it is to **the sound of a theorbo**. Likewise, it is no use pointing to some property which we identify as a region determinable in some ordinary physical or functional class of properties, and then identifying it with some other determinable grounding some entirely different set of similarity or difference relations. Neither may we identify the similarity relations of **phenomenal consciousness**, which must have the structure of a dimension, with some already existent class of similarity relations that do not have a full dimensional or sub-dimensional structure.

One final hope would be to concede that **consciousness** is in the relevant sense a primitive property: a *sui generis* set of similarity and difference relations with a full (sub) dimensional structure between natural determinates, but nevertheless hold out that the natural determinates it relates are physical. I take it to be plausible that if a collection of determinates is physical, then the determinables of those determinates should also be physical. I do not see any knock down argument for this conclusion. But what would be the gain of such an identification? It would amount to holding that there is some full dimensional similarity structure that we have not identified in any more scientifically acceptable terms among some realm of physical determinates. This will of necessity amount to saying that there are objective similarities and differences between physical properties that their canonical physical descriptions do not disclose. It is not at all obvious that the theoretical virtue of such a partial dualism outweigh the virtues of a more straightforward dualism.

§5.5: THE NATURAL DIMENSION ARGUMENT VERSUS THE CONCEIVABILITY ARGUMENT

I have given several arguments from a conceptual premise – the non-vagueness of the concept ‘Phenomenally Conscious’ – to a metaphysical

conclusion – Property Dualism or Panpsychism. From a distance, this sounds a lot like a conceivability argument. Is not my argument really just the usual conceivability argument in disguise?

The conceivability argument involves the critical epistemic-ontological linking principle that everything conceivable is possible. There is no principle in my argument playing that role. I appeal to a critical epistemic-semantic linking principle in my argument that ‘Phenomenally Conscious’ is not vague, but to get to the metaphysical conclusion I add the metaphysical premise that **phenomenal consciousness** is significant* .

There are questions about how this metaphysical premise is justified. Here, there may be an epistemic-ontological linking principle in play – the move from the premise that some property *seems* to be significant* to the conclusion that it *is* significant*. As we saw, significance* basically comes for free if the concept canonically expressing the property in question is vague. Since a concept in general only seems vague if it is vague, this might give us a restricted defense of the principle. There are also more general reasons to accept the principle. It may be that many of our objective similarity judgments are mandated by the concepts with which we make them, so that those concepts would be defective if the properties they expressed did not make those judgments true. This would not quite give us the principle that if a property seems significant* then it is significant*, but it would give us that if a

property seems significant* then either it is or the canonical concept of that property is defective. Other reasons to accept this linking principle have to do with a general anti-skeptical approach to metaphysics. How do we know that **grue** is less natural than **green** or **blue**? One austere response is to say that actually these beliefs are not justified, and only science can teach us anything about objective similarity – our prescientific intuitions at best only show us the structure of objective similarity relations. This response may be too austere, however, since it is hard to see how to do science unless you bring with you a basic capacity to identify similarities and differences in terms of spatiotemporal structure, causal structure, etc.

So there is a way to justify my claim that **phenomenal** consciousness is significant* in terms of a general epistemic-ontological linking principle. But this is a very different principle from the sort that adherents of the classical Conceivability argument use (that sort of principle says something about the link between the Conceivability of some proposition or state of affairs and its Possibility). Rejecting that sort of principle does nothing to motivate a rejection of the one I have in mind. And such a principle may not be the only way to defend the premise that **phenomenal consciousness** is significant*. My argument in section three from the normative significance of **phenomenal consciousness** to its significance* hinges on special features of **phenomenal**

consciousness that make it significant* even if the general linking principle is false.

I conclude by briefly sketching a different argument for the disjunction of Property Dualism or Panpsychism, from the sorts of principles I have invoked, but that looks more explicitly like a conceivability argument.

This argument makes use of the notion of a Canonical Concept of some property – a concept that seems to disclose the nature or essence of the property in question, or anyway some aspect of its nature or essence. Consider the linking principle: If property **P** is objectively similar to property **Q**, then for any canonical concept ‘P’ of **P** there must be some canonical concept ‘Q’ of **Q** that gives us a positive conception of that similarity.

In chapter three I argue that ‘Phenomenally Conscious’ is *non-excludable* – that we have no positive conception of what it is to not be phenomenally conscious; only a negative conception. I exploit the non-excludability of this concept to argue for its non-vagueness. But here we argue directly from non-excludability: Suppose that **phenomenal consciousness** is similar to but distinct from some physical property **p** (nothing that is **p** is phenomenally conscious). But then the non-excludability of ‘Phenomenal Consciousness’ strongly suggests that there will be no canonical concept of **p** that gives us a positive conception of how **p** is objectively similar to **phenomenal consciousness**. I am not sure what reason we have to accept the

linking principle of this argument, independent of our reasons to accept the linking principle of the standard conceivability argument. This linking principle is certainly closer to the classical conceivability-possibility linking principle than is the linking principle telling us that if a property seems to be significant* then it is significant*. However, it is not obvious that rejection of the classical linking principle entails rejection of the one I consider in this section.

If you reject the classical linking principle of the conceivability argument you reject the claim that all identity statements are *a priori*, and you also reject the claim that all identity statements are *a priori* entailed by the microphysical facts, or by the ontologically fundamental facts. But the present linking principle does not imply otherwise. There is no requirement here that every canonical concept of **p** and **q** disclose **p**'s metaphysical relation to **q**, but only that one of them does. So if **p** = **q** the principle is satisfied by 'P is P'. So I take it to be an open question whether this argument gets us anywhere. But in any case, it is very different from the argument at the center of this chapter. Though a linking principle is useful there, it may not be necessary, and in any event it is a linking principle of a very different order from the linking principle of the classical conceivability argument, or of the variant I have just sketched here.

BIBLIOGRAPHY

- Albert, D. Z. (1993). *Quantum Mechanics and Experience*. Cambridge, Mass.: Harvard University Press
- Alston, W. (1964). *Philosophy of Language*. Englewood Cliffs, N.J.: Prentice-Hall.
- Antony, M. (2006a). Vagueness and the Metaphysics of Consciousness. *Philosophical Studies* 128: 515-538.
- (2006b). Are Our Concepts **Conscious State** and **Conscious Creature** Vague? *Synthese*.
- (2006c). Papineau on the Vagueness of Phenomenal Concepts. *Dialectica* 60 (4), 475 - 483
- Armstrong, David (2004). *Truth and Truthmakers*, Cambridge: Cambridge University Press
- (1997). *A World of States of Affairs*. Oxford University Press.
- (1988). Can a Naturalist Believe in Universals? *Science in Reflection*, Edna Ullman - Margalit, editor. Dordrecht: Kluwer Academic Publishers.
- (1978). *Universals and Scientific Realism*. Cambridge: Cambridge University Press.
- Barnes, E. & Williams, J.R.G. (2009). Vague Parts and Vague Identity. *Pacific Philosophical Quarterly* 90 (2): 176-187
- Bechtel, W. and Richardson, R. (1998). Vitalism. In E. Craig (Ed.), *Routledge Encyclopedia of Philosophy*. London: Routledge.
- Bealer, G. (1994). Mental properties. *Journal of Philosophy* 91:185-208.
- Bennett, K. (Forthcoming). By Our Bootstraps. *Philosophical Perspectives*
- Berkeley, George (1713). *Three Dialogues between Hylas and Philonous*.
- Bermudez, J. (2004) Vagueness, phenomenal concepts and mind-brain identity. *Analysis* 64: 134-139.
- Block, N. (Forthcoming) Functional Reduction. in Horgan, T., Sosa, D., and Sabates, M. eds., *A Festschrift for Jaegwon Kim*

- (2006) Max Black's Objection to Mind-Body Identity. in *Oxford Studies in Metaphysics*, II, edited by Dean Zimmerman with replies by John Perry and Stephen White, Oxford University Press
- (2003), Do Causal Powers Drain Away?, *Philosophy and Phenomenological Research* 67: 133-50.
- (2002) The Harder Problem of Consciousness. *Journal of Philosophy* XCIX: 391-425
- Block, N. and Stalnaker, R. (1999). Conceptual Analysis, Dualism, and the Explanatory Gap. *Philosophical Review* 108:1-46.
- Boghossian, P. (2010) Williamson on the A Priori and the Analytic. in Bernecker and Pritchard (eds.): *A Companion to Epistemology*. London: Routledge
- (1990a). The Status of Content. *Philosophical Review* 99: 157-84
- (1990b) Naturalizing Content. in Barry Loewer and Georges Rey (eds.): *Meaning in Mind: Fodor and his Critics*, Oxford: Basil Blackwell, pp. 65-86
- (1989) The Rule Following Considerations. *Mind*, 98, 507-549
- Boghossian, P. and Velleman, D. (1991) Physicalist Theories of Color. *Philosophical Review*, vol. 100, no 1. Pp. 67 - 106
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*,
Harvard University Press (Cambridge)
- Brewer, B. (2007). Perception and its Objects. *Philosophical Studies* 132: 87-97.
- Bricker, P. (2006). The Relation between General and Particular: Entailment vs. Supervenience. *Oxford Studies in Metaphysics*, vol II, Dean Zimmerman, ed.
- Broad, C. D. (1925). *The Mind and its Place in Nature*. London: Routledge and Kegan Paul
- Brogaard, B. (2010) Degrees of Consciousness. Presented at the SPAWN conference on Metaphysics, Syracuse University, July 24-28, 2010
- Burge, T. (2010). *Origins of Objectivity*. Oxford University Press
- Byrne, A. and Hilbert, D. (2006). Color Primitivism. In Schumacher, R. (ed.) *Perception and the Status of Secondary Qualities*. Kluwer.

- Campbell, J. (2002). *Reference and Consciousness*. Oxford University Press.
- Campbell, K. (1981). *The Metaphysics of Abstract Particulars*. *Midwest Studies in Philosophy*, vol iv: The foundations of analytical philosophy. Edited by P. French, Theodore E. Uehling, Jr. and H. Wettstein. University of Minn. Press.
- (1970). *Body and Mind*. London: Doubleday.
- Carruthers, P. (2000). *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press
- Casati, R. and Varzi, A. (1999). *Parts and Places. The Structures of Spatial Representation*, Cambridge MA, MIT Press, 1999
- Chalmers, D.; Manley, D.; Wasserman, R.; eds. (2009). *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.
- Chalmers, D. (Forthcoming) *Constructing The World*
- (Forthcoming) The Two Dimensional Argument Against Materialism.
- (2006a) Perception and the Fall from Eden. in Hawthorne and Gendler (eds.) *Perceptual Experience*. Oxford University Press.
- (2006b) *Phenomenal Concepts and the Explanatory Gap* in (T. Alter & S. Walter, eds) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press.
- (2004). The Foundations of Two-Dimensional Semantics. In M. Garcia - Caprintero and J. Macia (eds), *Two-Dimensional Semantics: Foundations and Applications*. Oxford: Oxford University Press.
- (2002a). The Content and Epistemology of Phenomenal Belief. In Q. Smith and A. Jokic (eds.), *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press. <http://consc.net/papers/belief.html>.
- (2002b). Does Conceivability entail Possibility? In T. Gendler and J. Hawthorne (eds.), *Conceivability and Possibility*. Oxford: Oxford University Press.
- (2002c). Consciousness and its Place in Nature. In D. Chalmers, ed *Philosophy of Mind*:

- Classical and Contemporary Readings*. Oxford, 2002
- (1999). Materialism and the Metaphysics of Modality. *Philosophy and Phenomenological Research* 59:473-93. <http://consc.net/papers/modality.html>.
- (1997). Moving Forward on the Problem of Consciousness. *Journal of Consciousness Studies* 4:3-46. Reprinted in Shear 1997.
- (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chalmers, D. and Jackson, F. (2001). Conceptual Analysis and Reductive Explanation. *Philosophical Review* 110:315-61. <http://consc.net/papers/analysis.html>.
- Churchland, P. S. (1997). The Hornswoggle Problem. In Shear 1997.
- Churchland, P. M. (1996). The Rediscovery of Light. *Journal of Philosophy* 93:211-28.
- (1984). *Matter and Consciousness*. Cambridge, MA: MIT Press. .
- Clapp, L. (2001). Disjunctive Properties: Multiple realizations. *Journal of Philosophy*, 98, 111-136.
- Clark, A. (2000). A Case where Access Implies Qualia? *Analysis* 60:30-38.
- Crick, F., & Koch, C. (1990). Towards a Neurobiological Theory of Consciousness. *Seminars in the Neurosciences* v.2, 263-275
- Curd, P. (2011) Presocratic Philosophy. *The Stanford Encyclopedia of Philosophy (Fall 2011 Edition)*. Edward N. Zalta (ed.) URL=<http://plato.stanford.edu/archives/fall2011/entries/presocratics/>
- Dasgupta, S. (Forthcoming). On The Plurality of Ground.
- Davidson, D. (1991). Three Varieties of Knowledge. in A. Phillips Griffiths (ed.), *A.J.Ayer Memorial Essays: Royal Institute of Philosophy Supplement*, 30, Cambridge: Cambridge University Press
- (1970). Mental Events. In L. Foster and J. Swanson, editors, *Experience and Theory*. Amherst: Massachusetts University Press.
- Davis, W. (2003). *Meaning, Expression and Thought*. Cambridge University Press.
- Demos, R. (1917). A Discussion of a Certain Type of Negative Proposition. *Mind* n.s. 26, 188-196

- Dennett, D. C. (2001). The Fantasy of First-Person Science.
<http://ase.tufts.edu/cogstud/papers/chalmersdeb3dft.htm>.
- (1998). Animal Consciousness: What Matters and Why. In D. Dennett (Ed.), *Brainchildren*. Cambridge: MIT Press.
- (1996). Facing Backward on the Problem of Consciousness. *Journal of Consciousness Studies* 3:4-6.
- (1991). *Consciousness Explained*. Boston: Little, Brown & Company.
- Descartes, R. (1641). *Meditations on First Philosophy*, E. S. Haldane and G. R. T. Ross, trans. In *The Philosophical Works of Descartes*, Vol. 1. Cambridge: Cambridge University Press, 1970.
- Deutsch, M.(2005) Intentionalism and Intransitivity. *Synthese* 144: 1-21
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, Mass.: MIT Press.
- Dworkin, R. (2011) *Justice for Hedgehogs*. Cambridge, Mass.: Belknap Press of Harvard University Press
- (2006), *Justice in Robes*. Harvard University Press, Cambridge, MA
- Ehring, D. (2003). Part-Whole Physicalism and Mental Causation. *Synthese* 136:359–388.
- Ellis, B. (1999). ‘The Really Big Questions. *Metascience* 8: 63–73.
- Evans, G. (1979). Reference and Contingency. *The Monist* 62:161-89.
- Evans, G. (1978). Can There Be Vague Objects? *Analysis* 38.
- Feigl, H. (1958), The Mental and the Physical. *Minnesota Studies in the Philosophy of Science*, Volume II (Concepts, Theories, and the Mind-Body Problem): 370–498.
- Field, H., (1994a). Deflationist Views of Meaning and Content. *Mind*, 103 (411): 249–84.
- (1994b). Disquotational Truth and Factually Defective Discourse. *Philosophical Review*, 103 (3): 405–52.
- Fine, K. (2009). The Question of Ontology. In Chalmers, D., Manley, D. and Wasserman, R. (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.
- (2001). The Question of Realism. *Philosophers' Imprint* 1 (2):1-30.

- (1994). Essence and Modality. *Philosophical Perspectives* 8:1-16.
- (1975) Vagueness, Truth and Logic. *Synthese* 54: 235-59
- Fodor, J. (1998). *Concepts*. New York: Oxford University Press.
- Foster, J. (1991). *The Immaterial Self: A Defence of the Cartesian Dualist Conception of the Mind*. Oxford: Oxford University Press.
- Funkhouser, E. (2006). The Determinable–Determinate Relation. *Nous* 40:548–569.
- Gaifman, H. (2010). Vagueness, Tolerance and Contextual Logic. *Synthese*, 174: 5-46
- Gärdenfors, P. (2000). *Conceptual Spaces*. Cambridge: MIT Press
- Garfinkel, A. (1981). *Forms of Explanation*, New Haven, CT: Yale University Press.
- Geach, P. T. (1980) *Reference and Generality*. 3rd edn. Ithaca: Cornell University Press.
- Gillett, Carl & Rives, Bradley (2005). The Nonexistence of Determinables: or, a World of Absolute Determinates as Default Hypothesis. *Noûs* 39 (3):483–504.
- Ghirardi, G.C., Rimini, A., and Weber, T. (1985). A Model for a Unified Quantum Description of Macroscopic and Microscopic Systems. *Quantum Probability and Applications*, L. Accardi et al. (eds), Springer, Berlin.
- Goff, P. (forthcoming). Vagueness and Consciousness.
- Goodman, N. (1955). *Fact, Fiction, & Forecast*. Cambridge, MA: Harvard University Press.
- Griffin, D. R. (1998). Unsnarling the World-Knot: Consciousness, Freedom, and the Mind-Body Problem. Berkeley: University of California Press.
- Hale, B. and Hoffman, A. (2010). *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press.
- Hale, B. and Wright, C. (2001) *The Reason's Proper Study: Essays Towards A Neo-Fregean Philosophy of Mathematics*, Oxford: Oxford University Press
- Hameroff, S.R. (1998). [Quantum](#) Computation in Brain Microtubules: The Penrose-Hameroff Orch-Or Model of Consciousness. *Philosophical Transactions Royal Society London (A)* **356**: 1869–1896

- Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives* 4:31-52.
- Hegel, W.F. (1977). *Phenomenology of Spirit*, trans. A. V. Miller, Oxford University Press.
- Heil, J. (2003). Levels of Reality. *Ratio (new series)* 16:205–221.
- Hellie, B. (2005). Noise and Perceptual Indiscriminability. *Mind* 114: 481–508.
- Hewitt, S. (2008). *Normative Qualia and a Robust Moral Realism*. Doctoral Dissertation. NYU.
- Higginbotham, J. (1998). Conceptual Competence. *Philosophical Issues*. 9:149-162.
- Hill, C. S. (1997). Imaginability, Conceivability, Possibility, and the Mind-Body Problem. *Philosophical Studies* 87:61-85.
- Hodgson, D. (1991). *The Mind Matters: Consciousness and Choice in a Quantum World*. Oxford: Oxford University Press.
- Horgan, Terence E. (1989). Mental Quausation. *Philosophical Perspectives* 3:47–74.
- Horwich, P., (1998a). *Meaning*. Oxford: Clarendon.
- (1998b). *Truth*. Oxford: Blackwell (first edition 1990)
- Huxley, T. (1874). On the Hypothesis that Animals are Automata, and its History. *Fortnightly Review* 95:555-80. Reprinted in *Collected Essays*. London, 1893.
- Jackson, F. (2000). *From Metaphysics To Ethics*. Oxford University Press.
- (1994). Finding the Mind in the Natural World. In R. Casati, B. Smith, and G. White (eds.), *Philosophy and the Cognitive Sciences*. Vienna: Holder-Pichler-Tempsky.
- (1982). Epiphenomenal Qualia. *Philosophical Quarterly* 32:127-136.
- (1979). A Note on Physicalism and Heat. *Australasian Journal of Philosophy* 58:26-34.
- James, W. (1890). *The Principles of Psychology*. Reprint: Cambridge, MA: Harvard University Press.

- (1904). Does 'Consciousness' Exist?, *Journal of Philosophy, Psychology and Scientific Methods*, 1(18).
- Johansson, I. (2000). Determinables as Universals. *The Monist* 1, 2000, pp. 101-121
- Johnson, W. E. (1921). *Logic*, Part I, Cambridge: Cambridge University Press.
- Johnston, M. (1992). How to Speak of the Colors. *Philosophical Studies* 68 (3): 221-263
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*. New York: Oxford University Press.
- Katz, J. (1972). *Semantic Theory*. New York: Harper & Row.
- Kim, J. (1992). Multiple Realization and the Metaphysics of Reduction. *Philosophy and Phenomenological Research* 52:1-26. Reprinted in Kim 1993a, pp. 309-35.
- Kirk, R. (1974). Zombies vs Materialists. *Proceedings of the Aristotelian Society* (Supplementary Volume) 48:135-52.
- Kripke, S.A. (1982). *Wittgenstein on Rules and Private Language*, Oxford: Blackwell
- (1980). *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.
- Lee, G. (Forthcoming). On Being an Alien. *Festschrift for Ned Block*
- Lenin, V.I. (1909). *Materializm i Empirio-Kritisizm*. Translated as *Materialism and Empirio-Criticism*, Moscow: Foreign Languages Publishing House, 1959.
- Levine, J. (2006). Phenomenal Concepts and the Materialist Constraint. In (T. Alter & S. Walter, eds) *Phenomenal Concepts and Phenomenal Knowledge: Essays on Consciousness and Physicalism*
- (2000). *Purple Haze: The Puzzle of Conscious Experience*. Cambridge, Mass.: MIT Press.
- (1983). Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly* 64:354-61.
- Lewis, D. (1994). Reduction of Mind. In S. Guttenplan (ed.), *Companion to the Philosophy of Mind*. Oxford: Blackwell.
- (1993) Many, but Almost One. in *Ontology, Causality and Mind: Essays in Honour of D. M. Armstrong*, Bacon, John (ed) New York: Cambridge University Press
- (1991). *Parts of Classes*. Oxford: Blackwell Publishers

- (1988). What Experience Teaches. *Proceedings of the Russellian Society* (University of Sydney).
- (1988). Evans Misunderstood. *Analysis* 48.
- (1986). *On the Plurality of Worlds*. Wiley-Blackwell.
- (1984). Putnam's Paradox. *Australasian Journal of Philosophy*. 62:221-236.
- (1983a). New Work for a Theory of Universals. *Australasian Journal of Philosophy*, 61: 343-377
- (1983b). Extrinsic Properties. *Philosophical Studies*, 44: 197-200
- Lewis, D. and Langton, R. (1998). Defining Intrinsic. *Philosophy and Phenomenological Research*, 58: 333-345
- Loar, B. (1990). Phenomenal States. *Philosophical Perspectives* 4:81-108
- Lockwood, M. (1993). The Grain Problem. In H. Robinson (ed.), *Objections to Physicalism*. Oxford: Oxford University Press.
- (1989). *Mind, Brain, and the Quantum*. Oxford: Oxford University Press.
- Lopez de Sa, D. (Forthcoming). Arguments from Vagueness against Identifications.
- Lowe, E.J. (2001). Ontic Indeterminacy of Identity Unscathed. *Analysis* 61 (271): 241-245
- (1996). *Subjects of Experience*. Cambridge: Cambridge University Press.
- (1995) The Problem of the Many and the Vagueness of Constitution. *Analysis* 55: 179-82
- Lycan, W.G. (1996). *Consciousness and Experience*. Harvard, Mass.: MIT Press.
- Mach, E. (1886). *Die Analyse der Empfindungen und das Verhältnis des Physischen zum Psychischen*. Fifth edition translated as *The Analysis of Sensations and the Relation of Physical to the Psychical*, New York: Dover. 1959
- (1894), *Populärwissenschaftliche Vorlesungen*. Reprint of the fifth edition of 1923, Vienna: Böhlau Verlag, 1923. *Translated as Popular Scientific Lectures*, La Salle: The Open Court Publishing Company, 1943.
- (1905), *Erkenntnis und Irrtum*. Fifth edition translated as *Knowledge and Error*, Dordrecht: Reidel, 1976.

- Marksonian, N. (1998) Brutal Composition. *Philosophical Studies* 92: 211-49
- Martin, M. (2004). The Limits of Self-Awareness. *Philosophical Studies* 120 (1-3):37-89.
- Maxwell, G. (1979). Rigid Designators and Mind-Brain Identity. *Minnesota Studies in the Philosophy of Science* 9:365-403.
- (1976). Scientific Results and the Mind-Brain Issue: Some Afterthoughts. in Globus, G., 1976, 329–358.
- Maxwell, N. (1968). Understanding Sensations. *Australasian Journal of Philosophy* 46:127-45.
- McDaniel, K. (2009). Ways of Being. In Chalmers et al. (2009). 290–319.
- McGee, V. and McLaughlin, B. (2000) The Lessons of the Many. *Philosophical Topics* 28: 129-51.
- McGinn, C. (1996). *The Character of Mind*. (2nd ed.). Oxford: Oxford University Press.
- (1989). Can we Solve the Mind-Body Problem? *Mind* 98:349-66.
- Merricks, T. (1998) Against the Doctrine of Microphysical Supervenience. *Mind* 107: 59-71
- Millikan, R. (2000). *On Clear and Confused Ideas*. Cambridge: Cambridge University Press
- Mills, E. (1996). Interactionism and Overdetermination. *American Philosophical Quarterly* 33:105-115.
- Moore, G.E. (1903). *Principia Ethica*. Cambridge: Cambridge University Press
- Mulligan, K, Simons, P., and Smith, B., (1984) Truthmakers. *Philosophy and Phenomenological Research* (44).
- Nagel, T (1986). *The View from Nowhere*. Oxford: Oxford University Press.
- (1979). Panpsychism. in Nagel's *Mortal Questions*, Cambridge: Cambridge University Press.
- (1974). What is it Like to Be a Bat? *Philosophical Review* 83:435-50.
- (1971). Brain Bisection and the Unity of Consciousness. *Synthese*, 22: 396–413.

- Neander, K. (2004) Teleological Theories of Content. *The Stanford Encyclopedia of*
- Nordby, K. (1990). Vision in a Complete Achromat: A personal Account. In R. Hess, L, Sharpe, and K. Nordby (eds.), *Night Vision: Basic, Clinical, and Applied Aspects*. Cambridge:Cambridge University Press.
- Olson, E. (1995). Why I Have No Hands. *Theoria* 61: 182-197
- Papineau, D. (2006) Phenomenal and Perceptual Concepts. In T. Alter and S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge*. New York: Oxford University Press
- (2002). *Thinking About Consciousness*. Oxford: Oxford University Press.
- (1993a). *Philosophical Naturalism*. Oxford: Blackwell.
- (1993b). Physicalism, Consciousness, and the Antipathetic Fallacy. *Australasian Journal of Philosophy* 71:169-83.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Peacocke, C. (2008). *Truly Understood*, Oxford University Press.
- Penrose, R. (1989). *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford University Press
- Perry, J. (2001). *Knowledge, Possibility, and Consciousness*. Cambridge, Mass.: MIT Press.
- Popper, K. and Eccles, J. (1977). *The Self and Its Brain: An Argument for Interactionism*. Berlin, Heidelberg: Springer.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs*, 34 (4): 517–549.
- Quine, W. V. (1951). Two Dogmas of Empiricism. *Philosophical Review* 60:20-43.
- Raven, M. (Forthcoming). In Defence of Ground. *Australasian Journal of Philosophy*.
- Rey, G. (1995). Toward a Projectivist Account of Conscious Experience. In T. Metzinger (ed.), *Conscious Experience*. Paderborn: Ferdinand Schoningh.
- Robinson, W. S. (1988). *Brains and People: An Essay on Mentality and its Causal Conditions*. Philadelphia: Temple University Press.
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In Hale and Hoffman (2010). 109–136.
- Rosen, G & Smith, N. (2004): Worldly Indeterminacy: A Rough Guide. *Lewisian Themes*. Oxford University Press.

- Rosenthal, D. M. (1997). A Theory of Consciousness. In N. Block, O. Flanagan, and G. Güzeldere (eds.), *The Nature of Consciousness*. Cambridge, Mass.: MIT Press.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson and Co.
- Russell, B. (1927). *The Analysis of Matter*. London: Kegan Paul.
- (1924). Logical Atomism. in J.H. Muirhead (ed.), 1924. Reprinted in Russell, 1956, 321–343.
- (1923). Vagueness. *Australasian Journal of Philosophy and Psychology* **1**, 84-92.
- (1922). Physics and Perception. *Mind*, 31, 478–485.
- (1921). *The Analysis of Mind*. Reprinted: London: George Allen & Unwin.
- (1918). *The Philosophy of Logical Atomism*.
- Sainsbury, M. (1990). Concepts Without Boundaries. in Keefe, R. and Smith, P. (eds.) *Vagueness: A Reader*, MIT Press.
- Sayre, K. (1976). *Cybernetics and the Philosophy of Mind*, Atlantic Highlands: Humanities Press
- Schaffer, J. (2004). Two Conceptions of Sparse Properties. *Pacific Philosophical Quarterly* 85:92-102.
- (2009). On What Grounds What. In Chalmers et al. (2009), 347—383.
- Schiffer, S. (2003). *The Things We Mean*. Oxford; OUP
- (2000). Vagueness and Partial Belief. *Philosophical Issues* 10: 220-257
- (1998). Two Issues of Vagueness. *The Monist* 88: 193-214
- Schopenhauer, A. (1859). *Die Welt als Wille und Vorstellung (The World as Will and Representation)* [third edition]
- Searle, J. (1992). *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.
- Sebastian, M (Forthcoming). Borderline Experiences One Cannot Undergo.
- Sellars, W. (1997). *Empiricism and the Philosophy of Mind*, Robert Brandom (ed.), Harvard University Press.; Cambridge, MA
- (1981). Is Consciousness Physical? *The Monist* 64:66-90.

- Shear, J. ed. (1997). *Explaining Consciousness: The Hard Problem*. Cambridge, Mass.: MIT Press.
- Shoemaker, S. (2007). *Physical Realization*. Oxford University Press.
- (2001), Realization and Mental Causation. in Gillett and Loewer 2001, pp. 74-98. Reprinted in S. Shoemaker, (2003), *Identity, Cause, and Mind*, Expanded Edition, Oxford: Clarendon Press, pp. 427-51.
- (1975). Functionalism and Qualia. *Philosophical Studies* 27:291-315.
- Sider, T. (Forthcoming). *Writing the Book of the World*. Oxford University Press
- (2003). Maximality and Microphysical Supervenience. *Philosophy and Phenomenological Research* 66: 139-149
- (2001). *Four Dimensionalism*. Oxford University Press
- (1995). Sparseness, Immanence, and Naturalness. *Nous* 29: 360-77.
- Simon, J. (forthcoming). Goff on Vagueness and Consciousness.
- Smart, J. J. C. (1959). Sensations and Brain Processes. *Philosophical Review* 68:141-56.
- Smith, B. and Simon, J. (2007) Truthmaker Explanations. *Truthmakers and Metaphysics* Ontos-Verlag. Reprinted in translation as (2011). Explications Verifactrice. *Philosophiques*.
- Smithies, Declan (2006). *Rationality and the Subject's Point of View*. Doctoral Dissertation. NYU.
- Soames, S. (1998). *Understanding Truth*. Oxford University Press.
- Stalnaker, R. (1978). Assertion. In P. Cole (ed.), *Syntax and Semantics: Pragmatics, Vol. 9*. New York: Academic Press.
- Stapp, H. (1993). *Mind, Matter, and Quantum Mechanics*. Berlin: Springer-Verlag.
- Stazicker, J. (Forthcoming). Attention, Visual Knowledge and Psychophysics.
- Stoljar, D. (2005). Physicalism and Phenomenal Concepts. *Mind and Language*
- (2001). Two Conceptions of the Physical. *Philosophy and Phenomenological Research* 62:253-81.

- Strawson, G. (2008). *Real Materialism and Other Essays*. Oxford University Press.
- (2000). Realistic Materialist Monism. In S. Hameroff, A. Kaszniak, and D. Chalmers (eds.). *Toward a Science of Consciousness III*. Cambridge, Mass.: MIT Press.
- (1994). *Mental Reality*. Cambridge, MA: MIT Press.
- Strevens, M. (2004). The Causal and Unification Accounts of Explanation Unified-Causally. *Nous*, 38: 154–76.
- Swinburne, R. (1986). *The Evolution of the Soul*. Oxford: Oxford University Press.
- Tegmark, M. (2000). Importance of Quantum Decoherence in Brain Processes. *Physical Review E* 61: 4194–4206
- Thompson, B. (2010). The Spatial Content of Experience. *Philosophy and Phenomenological Research* 81 (1):146-184.
- Tye, M. (2003). A Theory of Phenomenal Concepts. In (A. O'Hear, eds) *Minds and Persons*. Cambridge University Press
- (2000). *Consciousness, Color, and Content*. Cambridge, MA.: MIT Press.
- (1996). Is Consciousness Vague or Arbitrary? *Philosophy and Phenomenological Research* 56(3), 679-685.
- (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, Mass.: MIT Press.
- Unger, P. (2008). *All The Power in the World*. Oxford University Press.
- (2004). Mental Problem of the Many. In Zimmerman, D. (ed.) *Oxford Studies in Metaphysics*, Vol. 1. Oxford: Clarendon Press.
- (1988). Conscious Beings in a Gradual World. *Midwest Studies in Philosophy* 12, 287-333.
- Van Cleve, James (1990). Mind-Dust or Magic? Panpsychism versus Emergence. *Philosophical Perspectives*. 4:215-226.
- Van Gulick, R. (1993). Understanding the Phenomenal Mind: Are we all just Armadillos? In
- M. Davies and G. Humphreys (eds.) *Consciousness: Philosophical and Psychological Aspects*. Oxford: Blackwell.

- Van Inwagen, Peter (1990) *Material Beings*. Ithaca: Cornell University Press
- Varzi, A and Collins, J. (2000) Unsharpenable Vagueness. *Phil Topics*
- Weatherson, B. (2008). Vagueness As Indeterminacy.
- (2003). Many Many Problems. *Philosophical Quarterly*
- White, S. (1986). Curse of the Qualia. *Synthese* 68:333-68.
- Wigner, E. P. (1961). Remarks on the Mind-Body Question. In I. J. Good (ed.), *The Scientist Speculates*. London: Basic Books.
- Williamson, T. (2009). *The Philosophy of Philosophy*. Blackwell.
- (1997). Reply to Commentators. *Philosophy and Phenomenological Research* 57(4), 945-953.
- (1994). *Vagueness*. London: Routledge.
- Wilson, J. (Forthcoming). Fundamental Determinables. *Philosophers Imprint*.
- (2011). Non-Reductive Realization and the Powers-Based Subset Strategy. *The Monist*. 94: 121-154.
- Wittgenstein, L. (1953). *Philosophical Investigations*. G.E.M. Anscombe and R. Rhees (eds.), G.E.M. Anscombe (trans.), Oxford: Blackwell.
- (1922). *Tractatus Logico-Philosophicus*.
- Woodfield, A. (1991). Conceptions. *Mind* 100(4): 547-572
- Woodward, J. and Hitchcock, C. (2003). Explanatory Generalizations, Part I: A Counterfactual Account. *Nous* 3:1-24.
- Wright, C. (2004). Warrant for Nothing (and Foundations for Free?). *Arist Soc Supplement*
- (2001). On Being in a Quandary: Relativism, Vagueness, Logical Revisionism. *Mind* 110: 45-98
- (1975). On the Coherence of Vague Predicates. *Synthese* 30, 325-365.
- Yablo, S., (1992). Mental Causation. *Philosophical Review* 101: 245-80

Zimmerman, D. (2011). From Experience to Experienter. in Baker, M.C. and Goetz, S. (eds.) *The Soul Hypothesis*, London: Continuum Publishing.