

This electronic thesis or dissertation has been downloaded from the King's Research Portal at <https://kclpure.kcl.ac.uk/portal/>



An Argument for, and an Account of, Genuine Metaphysical Vagueness

Hughes, Robin

Awarding institution:
King's College London

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without proper acknowledgement.

END USER LICENCE AGREEMENT



Unless another licence is stated on the immediately following page this work is licensed

under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International

licence. <https://creativecommons.org/licenses/by-nc-nd/4.0/>

You are free to copy, distribute and transmit the work

Under the following conditions:

- Attribution: You must attribute the work in the manner specified by the author (but not in any way that suggests that they endorse you or your use of the work).
- Non Commercial: You may not use this work for commercial purposes.
- No Derivative Works - You may not alter, transform, or build upon this work.

Any of these conditions can be waived if you receive permission from the author. Your fair dealings and other rights are in no way affected by the above.

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

An Argument for, and an Account of, Genuine Metaphysical Vagueness

Robin Hughes

King's College London

Submitted for the degree of PhD in Philosophy

Table of Contents

Abstract	3
Acknowledgements	4
0. Introduction	5
0.1 What is vagueness?	9
0.2 Categories of vagueness	28
0.3 Connections between different types of vagueness	34
0.4 Formulating desiderata.....	36
0.5 Structure of the thesis.....	39
Part A	41
1. Supervaluationism	42
1.1 Supervaluationism makes true the statement ‘there is no sharp cut-off’	44
1.2 Supervaluationism cannot handle all instances of vagueness	48
1.3 Problems with assertion in cases of vagueness	51
1.4 Problems with the ontology of supervaluationism	52
1.5 Failures of reference and correspondence truth	53
1.6 Higher order vagueness	56
2. Epistemicism	60
2.1 How are the sharp boundaries determined?	65
2.2 How can language users know the meanings of vague terms?	68
2.3 The implausibility of extreme sensitivity to meaning.....	70
2.4 Why do we believe there are no sharp boundaries?	73
Part B	78
3. A general defence of metaphysical vagueness	79
3.1 General arguments against metaphysical vagueness.....	79
3.2 The Evans/Salmon argument from identity	81
3.3 Sider’s argument for unrestricted composition	85
3.4 Williamson and Sorensen’s argument for classical logic	86
4. Previous accounts of metaphysical vagueness	89
4.1 What is metaphysical vagueness?	90
4.2 Specific theories	93
4.3 General theories.....	108
5. My account of metaphysical vagueness	126
5.1 What are the desiderata?	127
5.2 Vague states of affairs	129
5.3 Truth-functional background.....	139
5.4 Can we reason in cases of vagueness?	150
5.5 What is the sorites solution offered by this?	154
6. Conclusion	159
6.1 Evaluation of my account against desiderata.....	162
6.2 General considerations.....	165
6.3 Areas for further research	167
Bibliography	171

Abstract

Metaphysical vagueness, the idea that the world itself, independent of our thoughts and representations of it, is fuzzy or indeterminate in some way, is often taken to be incoherent, unconvincing, unmotivated, or all three. I argue that this is not the case, and put forward an account of metaphysical vagueness in which the vagueness is taken to be fundamental both ideologically and ontologically.

I begin by clearing up what is meant by vagueness more generally, and coming to a neutral characterisation of it which can be accepted by all parties to the debate. I then address the general objections to metaphysical vagueness, arguing that it is coherent as an idea, and that the supposed arguments showing it to lead to contradiction do not achieve their aim. I remove the burden of proof from the proponent of metaphysical vagueness, arguing that there are actually general intuitions in favour of it, and that there is thus *prima facie* reason to consider whether there is an enlightening account of it on offer.

I consider the two most prominent alternative accounts of vagueness; first, that it is linguistic in nature, represented by supervaluational theories on which our language is underdetermined in cases of vagueness, but theoretically could be made more precise. I then consider the epistemic theory of vagueness, on which it is held to be characterised by a specific type of ignorance. I find that both of these theories face serious objections.

I then move on to discuss some previously offered theories of metaphysical vagueness. I begin with various theories which attempt to argue for the existence of specific types of metaphysical vagueness, for instance the existence of vague objects. None are found to be totally convincing, and I argue that we should look for a more general theory, given the widespread nature of vagueness. I then examine three 'general' theories of metaphysical vagueness, but again none are found to be both convincing and genuinely respectful of what I see as the fundamental nature of metaphysical vagueness.

Finally, I propose my own positive account of metaphysical vagueness, on which it consists in the obtaining of vague states of affairs; that is, those which contain a vague constituent. It is held to be a fundamental fact about the world that in some cases the world is vague, and we cannot explain or reduce this vagueness to any other concepts. We need to keep the concept of vagueness in our fundamental description of the world. I explain how this proposal can be accommodated into consistent theories of logic and reasoning, and how it helps us to solve the sorites paradox.

Acknowledgements

I have many people to thank for their help and support during the creation of this thesis.

Firstly, thanks must go to my supervisors; Chris Hughes, Eliot Michaelson, Carlo Nicolai, and Jessica Leech. Chris especially continued to give me comments and advice even after his retirement (congratulations!), and Jess gave invaluable help in working through the necessary corrections. Thanks also to John Callanan for essential support and guidance in the home straight.

My research was funded by an AHRC Research Studentship via the London Arts and Humanities Partnership between 2016 and 2019, and I gratefully acknowledge their support.

I should also thank my current employers and managers, for allowing me the time to finish writing the thesis after I had started full-time work.

Finally, and most importantly, I must thank my wife, my cat, and my family, for a lifetime of support (and cajolement when needed!).

0. Introduction

Consider a series of people lined up in height order; they range from exactly one metre fifty centimetres tall at one end, to exactly two metres tall at the other, with only a one millimetre difference in height between adjacent people. Starting with the shortest person, someone asks you to identify whether each individual in front of you is tall or not tall. You consider the person who is one metre fifty centimetres tall and confidently assert that this person is not tall. This continues for a while, until at some point you reach people who are of sufficient height that your confidence in their not being tall falters. Perhaps you still assert that they are not tall, but less confidently. Later still, you refrain from asserting that the people in front of you are not tall, but you are not ready to assert that they are tall, either. Towards the end of the series, you are back on firmer footing; when asked to classify the height of someone who is one metre and ninety centimetres, you confidently assert that they are tall, and so on until the end of the line. There are borderline cases of 'tall'; individuals who are not obviously tall, but neither are they obviously not tall.

Now imagine that a troublemaking philosopher asks whether you agree with the statement that 'a difference in height of one millimetre can change whether someone is tall or not'. You do not agree; such a small change cannot intuitively cause anyone to make that leap. Now the philosopher points to your classification of the shortest person in the series; the person is definitely not tall. In line with your intuitive belief that one millimetre is too small a measure to have any impact on whether someone can be classified as tall or not tall, you conclude that the next person in the line is also not tall. And so too the next person And so on, with height changes of one millimetre each time, until you are forced to conclude that the two metres tall person is also not tall. But earlier, before the philosopher got involved, you had confidently asserted that this last person in the series was definitely tall! Something has gone wrong.

This type of paradox (often called the sorites paradox from the Greek 'soros' meaning 'heap')¹ has been the motivation for a vast amount of literature on the topic of vagueness. For it is vagueness that has led us into difficulty here; the vagueness as to who can be truly categorised as tall, who cannot, and who is such that we simply struggle to make any decision at all. When considering concepts such as 'tall', we have both the intuition that a small change cannot cause a jump from being not tall to being tall, and the knowledge that small changes can aggregate to a large change.

In the case of 'tall' it is plausible to locate the vagueness somewhere in our concepts or language; when the idea of being tall emerged into everyday use, it simply was not defined well enough to allow us to point at the first man in the line who qualified. But there is theoretically nothing stopping us from sharpening the concept up and removing the vagueness; we could define someone as being tall if they, say, are at least six feet tall. We might do some violence to our natural understanding of 'tall', but we have eliminated the vagueness and blocked the undesirable conclusion of the paradoxical argument.

Now consider the case of separate sperm and egg cells, which then fuse and over time develop into a foetus, which then develops further and is given birth to as a baby, thus becoming a living, breathing, human. Where at the start of this process there was not a person, at the end there definitely is. But at what point in the process did the personhood emerge? Unlike in the case of 'tall' (as outlined above), this does not (at least at first glance) appear to be a case of us being lazy or inconsistent with our concepts or our language. While the term 'person' may in fact be under-defined, this does not feel like the source of the issue. People are real things which exist in the world, independent of language or representation, and yet there still appears to be vagueness in the question of when exactly a person sprang into existence; vagueness which we are unable to easily define out of existence.

This latter type of case is often the motivation for positing the existence of metaphysical vagueness; vagueness that exists independent of thought and representation, or in

¹ It is often first attributed to Eubulides, writing in the fourth century B.C.E; see Jon Moline, 'Aristotle, Eubulides and the Sorites', *Mind*, 78.322 (Jul., 1969), pp. 393-407.

Barnes' formulation, 'vagueness in what there is as opposed to our descriptions or knowledge of what there is'.²

Historically, many philosophers have been uncomfortable with the idea of metaphysical vagueness and dismissed it as even a possibility, arguing that vagueness simply must be linguistic or epistemic; in Dummett's words, 'the notion that things might actually be vague, as well as being vaguely described, is not properly intelligible'.³

This, then, is the question which I will seek to answer in this thesis:

Is there genuine metaphysical vagueness?

I will answer in the affirmative, by two methods. In Part A I argue that there is good reason to accept the existence of genuine metaphysical because the leading theories which deny this thesis (supervaluationism and epistemicism) cannot provide fully convincing accounts of the phenomena of vagueness, and in particular the phenomena of vagueness which are seemingly metaphysically located.

There is a compelling argument to be made that vagueness, as well as potentially arising in the realms of language, reasoning, perception, or knowledge, must have a place in our fundamental metaphysical description of the world itself. An explanation of the world in which its fundamental makeup is fully precise and determinate will not be an accurate description; there will be elements of reality which simply are indeterminate, and a description which rests entirely on determinate concepts will necessarily mis-represent them by over-specifying. There is thus prima facie motivation for avoiding this mis-representation by accepting the existence of genuine metaphysical vagueness, which these theories are unable to account for.

In Part B I argue that, having accepted that there is a theoretical space for an account of genuine metaphysical vagueness, we should choose an account which accommodates states of affairs which are indeterminate by their very nature, and which, in cases of metaphysical vagueness, actually obtain. While others who also affirm that we should

² Elizabeth Barnes, 'Ontic Vagueness: A Guide for the Perplexed', *Noûs*, 44.4 (2010), p. 601.

³ Michael Dummett, 'Wang's Paradox', *Synthese*, 30 (1975), p. 314. See also Bertrand Russell, 'Vagueness', *Australasian Journal of Psychology and Philosophy*, 1.2 (1923), pp. 84-85. This is discussed further in Chapter 3 of this thesis.

accept metaphysical vagueness believe that it can be explained by the fact that various determinate states of affairs indeterminately obtain,⁴ I will argue that in order to truly build a theory of genuine metaphysical vagueness, we need to be willing to accept indeterminacy in states of affairs themselves. I will explain how this framework functions, and how it best conforms with the observed phenomena of vagueness.

First, though, it is worth giving a brief overview of the current state of the debate on vagueness, and in particular how vagueness is characterised, in order to show the significance of the question of whether metaphysical vagueness is plausible, or even necessary.

⁴ Known as 'Unsettledness views'; see Elizabeth Barnes and Ross Cameron, 'Are There Indeterminate States of Affairs? No', in *Current Controversies in Metaphysics* ed. Elizabeth Barnes (Abingdon: Routledge, 2017), pp. 120-132, and a full discussion in Chapter 5 of this thesis.

Classifying Vagueness

0.1 What is vagueness?

Before we can begin an investigation into vagueness, its manifestations and causes, we must be clear about the terms of discussion. That is, we must decide what will fall under the term 'vague', in order that we can formulate a theory which will adequately deal with the phenomena. If we don't know what are the central cases which must be counted as instances of vagueness, we will not know whether putative accounts are able to explain that vagueness.

Vagueness in the philosophical sense is hard to pin down. While it is easy to get a feeling for the phenomena we want to target (being bald is a vague matter, whereas being exactly 60 years old is not), providing necessary and sufficient conditions for something's being vague is harder than might initially be supposed. In this section I will discuss various attempts to do just this. It should be noted that at this stage we are merely working at the level of identifying, not analysing. That is, we are attempting to identify which phenomena we want to be able to capture in a theory of vagueness, not taking any stand as to what that theory of vagueness should be. As such, a plausible answer to the question of what does it mean for something to be vague should be as neutral as is possible between accounts of what generates vagueness. This may be an impossible task, as offering any characterisation of vagueness will surely involve taking a stand on some issue, even unintentionally, which some theorists may take issue with. However, I think it is a worthwhile aim to bear in mind. It is a similar approach to that taken by Greenough, who aims to provide what he calls a 'minimal theory of vagueness', one which can 'provide an uncontroversial definition of vagueness, a definition which isolates the constitution of vagueness from a perspective which is as neutral as possible on matters logical and philosophical'.⁵ In fact I will be slightly less ambitious than Greenough and settle short of offering a full definition of vagueness; an elucidation of key cases and characteristics will suffice here, as I will go on to examine, and offer, full theories of vagueness later in the thesis.

⁵ Patrick Greenough, 'Vagueness: A Minimal Theory', *Mind*, 112.446 (Apr., 2003), p. 237.

To take an example from non-vagueness literature, I am following Leech's approach to identifying metaphysical necessity – substitute 'vagueness' for 'metaphysical necessity' in the below passage:

[M]y arguments here are based on what seems to me to be the fairly uncontroversial observation that there is a great diversity of metaphysical accounts of metaphysical necessity in the literature, accompanied by a widely held assumption that there is some kind of genuine disagreement between these accounts, and hence a shared target notion the metaphysics of which is at issue. My aim is to draw out an understanding of that target notion that leaves room for – and thereby also makes sense of – that genuine disagreement.⁶

It is also worth noting here the two aspects of any answer to the thesis's primary research question, one of which I will not remain neutral on, but one of which I will:

- 1) Are there any cases of genuine metaphysical vagueness?
- 2) Are all cases of vagueness to be accounted for as metaphysical vagueness?

As I have stated, I will answer affirmatively, to 1), that we should accept the existence of genuine metaphysical vagueness, and disagree with those who take it to be unmotivated or incoherent. However, 2) poses the further question of whether we should accept *only* metaphysical vagueness; that is, does an account of metaphysical vagueness provide a comprehensive account of the phenomena of vagueness, such that competing theories, which locate vagueness elsewhere, are superfluous? On this, I reserve judgement. It may well be the case that, as well as the instances of metaphysical vagueness which I will explore, there are instances of vagueness which arise from our use of language, or the limitations of our knowledge; in which case, semantic or epistemic accounts will still have an important part to play in our overall understanding of vagueness. I do not wish to argue for an 'imperial' theory of metaphysical vagueness, under which all vagueness is necessarily metaphysical, but instead argue that it should be accepted that at least some cases of genuine metaphysical vagueness exist (even if alongside other types).

There are three main characterisations offered, each of which is closely related to our pre-theoretic intuitions regarding the phenomena of vagueness:

⁶ Jessica Leech, *Thinking of Necessity* (Oxford: Oxford University Press, 2023), p. 134.

- (1) Something is vague if and only if it is susceptible to sorites-style reasoning to an absurd conclusion;
- (2) Something is vague if and only if it gives rise to borderline cases, in some way or another;
- (3) Something is vague if and only if it is tolerant, where tolerant means accepting of small changes with no change in status.

I will discuss each of them in turn, before concluding that the most appealing characterisation is Greenough's, which takes vagueness to arise where there exists epistemic tolerance.

I will then introduce the three main locations where vagueness has been posited to reside: in our language (semantic vagueness); in our knowledge, or lack of knowledge, of some certain facts (epistemic vagueness); and finally in the world itself, independent of language and knowledge (metaphysical vagueness). The next route of investigation will be to discuss whether there are any obvious links between the various phenomena of vagueness. This will help us in our desire to determine a set of desiderata for any convincing account of vagueness, and metaphysical vagueness in particular. If it seems plausible, or even obvious, that some types of vagueness will entail others in at least some cases, then any adequate account will need to explain why this is. Firstly, I will investigate links within the three categories; that is, links between, for example, different types of metaphysical vagueness. Secondly, I will ask whether there are any links to be found between categories.

Once all this is complete, we should have a clear (or as clear as possible) map of the terrain of vagueness, and an idea of why the question of whether to accept genuine metaphysical vagueness is an important one to answer. We can then continue with our investigation, and lay out the desiderata for any convincing account of vagueness, and metaphysical vagueness in particular.

0.1.1 Everyday vs philosophical vagueness

Before we go into any more detail, I believe it is useful to note that the concept of vagueness which I will be investigating in this thesis, and which has occupied philosophers for many years, is not exactly the same as the notion of vagueness as used in everyday conversation. This may seem so obvious as to not merit attention, but I think it is useful as a jumping off point. Drawing attention to the differences between everyday vagueness and philosophical vagueness will help us to identify the characteristics of the philosophical notion we wish to give an account of.

Khatchadourian points out that, in 1962 at least, ‘the *Oxford English Dictionary* defines the word ‘vague’, as applied to words and phrases, as “not precise or exact in meaning”.’⁷ The current Merriam-Webster primary definition of vague is ‘not clearly expressed; not having a precise definition’ (interestingly, one of the additional definitions is ‘not sharply outlined’; this appears to be more in line with the philosophical notion of vagueness, but it is unclear how widespread its use is).⁸ As we can see, and as Khatchadourian notes, the definitions of this everyday notion of vagueness are essentially negative; it is the lack of some sort of precision or exactitude, rather than the presence of some other feature. As we will discuss later, perhaps this points to an important lesson that should be heeded by any investigator of vagueness; plausibly, precision or exactitude can be lacking in many ways and for many reasons, and thus vagueness (in some sense or other) can come from many sources. In addition, it is perhaps worth noting that the everyday use of the term is not exclusively linguistic. It is easy to imagine a scenario where I ask a stranger for directions to the bank, and they (not having much time to spare) give a wave in a general direction before carrying on with their day. In relating this story, I would be justified in saying that the stranger had made ‘a vague gesture’ as to where the bank was. That is, non-linguistic actions appear to be capable of lacking precision or exactitude in the way that makes the term ‘vague’ applicable to them.

⁷ Haig Khatchadourian, ‘Vagueness’, *The Philosophical Quarterly*, 12.47 (Apr., 1962), p. 138.

⁸ Merriam-Webster, ‘Vague’, <<http://www.merriam-webster.com/dictionary/vague>>.

Khatchadourian then sets out to discover what is the exactitude that vague expressions lack. He begins by picking out the essential features of a non-vague expression. The first is that ‘there exists a set of relatively fixed and well defined linguistic rules which conventionally govern its usage’; the second is that a non-vague expression ‘does convey some definite ideal content’.⁹ Conversely, a vague expression has no such fixed rules governing it; moreover, it ‘cannot properly be said to express some one or more fixed concept or idea’.¹⁰

Now, this characterisation of a vague expression is acceptable as far as it goes, but the examples Khatchadourian gives show the limitations of it. He gives a number of philosophical terms and phrases which fall under this definition of vagueness; for instance, ‘correspond with’ as used by a correspondence theory of truth, or ‘ingress’ as used by Plato. What he does not give are examples of everyday expressions which count as vague on this definition. Perhaps the best candidate is something like ‘about’, as used in the sentence ‘There were about 35,000 people at the match’. Plausibly, there are no fixed rules governing its use (could it be used if there were 37,000 people at the match? 38,000?) and it does not express a fixed concept. However, in another sense, there are fixed rules; it can be used if the number of people at the match is saliently similar to 35,000.

There are other examples which bring out the impreciseness more clearly; if I ask a friend when he thinks World War III will break out and he replies ‘Soon’, this is extremely vague. ‘Soon’ can mean almost any time period when situated in the right context. For an astronomer ‘soon’ could mean in the next million years; for a geologist ‘soon’ could mean in the next 10,000 years; for a Southern train announcer, ‘soon’ could mean in the next hour. If my friend is a pessimist, he could believe war will occur in the next few months; if he is a believer in historical cycles, ‘soon’ could mean in the next twenty years.

However, there are other uses which would ordinarily be classed as vague, which are governed by well-defined rules. For instance, if my brother had answered that my

⁹ Khatchadourian, ‘Vagueness’, p. 140.

¹⁰ *Ibid.*, p. 141.

parents were arriving at the airport ‘between 1pm and 5pm’, his phrase, taken by itself, would have been precise and exact. There are well defined linguistic rules for it, and it precisely applies to a certain four-hour period. In the context of our conversation, though, and my need for more granular knowledge, it is still vague (and unhelpful). The notion of unhelpfulness is, I believe, important to everyday vagueness; the key feature in all examples is some level of un informativity. I need to know a particular time to aim to be at the airport. My brother’s answer is uninformative whether or not it is precise; giving a specific range of four hours is no help to me, as I need to know more exactly when I need to leave the house.

There are of course other examples than time; the phrase ‘Brexit means Brexit’ is (disconcertingly, infuriatingly) vague, because it conveys no information. Taken strictly and literally, it is a necessarily true statement, governed by defined linguistic rules. As used to answer any question, it is extremely uninformative. These examples perhaps show that (at least certain) violations of Grice’s maxims of quantity or relevance are referred to in everyday speech as instances of vagueness.¹¹

I believe the problem is that Khatchadourian is unduly focusing on individual expressions, when the everyday notion of vagueness is more properly applied to larger linguistic items, such as clauses or sentences. Examples of language which would classically be termed as vague are things such as replies to questions, or explanations, or even whole stories. In addition, I think Khatchadourian, in attempting to sharply pin down the nature of everyday linguistic vagueness, has already steered into technical philosophical waters. Plausibly, the only feature common to all uses which fall under the everyday notion of vagueness is indeed a lack of exactitude; but we cannot pin down which form of exactitude is lacking.

When a friend asks me where another friend is, and I reply, ‘Somewhere in England’, I am being vague because (a) the location I give is nowhere near exact enough to be informative in the conversation. It is not the case that there is an inexactitude of meaning; ‘Somewhere in England’ expresses a definite content, and gives rise to

¹¹ See H. P. Grice, ‘Logic and Conversation’, in *Syntax and semantics Vol. 3: Speech acts* eds. P. Cole, and J. Morgan (New York: Academic Press, 1975), pp. 41-58.

determinate truth values. On the other hand, when the friend asks how many people were at the party last night, and I say, ‘I reckon about 30’, I am being vague because (b) I am deliberately avoiding committing myself to a specific number. I am not being deliberately uninformative, but the nature of ‘about’ means that it is unclear what exactly is the range of guests included in it. Another possible form of inexactitude is when I, discussing the poor recent form of a football goalkeeper, observe that ‘he is not very tall’, and my friend disagrees by saying that the goalkeeper is six feet tall. In this case the inexactitude is due to (c) the context-dependence of the term ‘tall’; I could have been saying that the goalkeeper is not tall relative to other professional goalkeepers, or not tall relative to the general population (of which country?). Still other times, vagueness is used as almost a synonym for (d) ambiguity in natural language.

Philosophers have tended to focus on examples similar to (b), rather than (a), (c), or (d). Keefe and Smith, for instance, preface their discussion of vagueness by dispensing with ‘underspecificity’ (their term for un informativity), context-dependence, and ambiguity, while being careful to note that vague terms may be, and indeed often are, context-dependent and possibly ambiguous.¹² This seems understandable; ambiguity is extensively discussed elsewhere, as is deliberate un informativity and contextual dependence. But it should not mean that the fact that these cases are often called vague, in ordinary talk, should be forgotten. It is an interesting fact that many sorts of imprecision are grouped under the banner of vagueness, and we should use that as a guide in our investigations.

That is to say, in a long-winded way, that I agree with Raffman when she says that ‘vagueness is a form of unclarity’ – in the interesting case above, the phrase ‘about 30’ is uninformative because it is unclear as to the exact range of numbers which would satisfy it.¹³ In the everyday sense, perhaps this is all we can say. The philosophical investigation into vagueness is an attempt to investigate one aspect of unclarity, which does not fall under previously explained ideas of ambiguity, un informativity, and so on.

¹² Rosanna Keefe and Peter Smith, ‘Theories of Vagueness’, in *Vagueness: A Reader* eds. Rosanna Keefe and Peter Smith (Cambridge: MIT Press, 1997), pp. 5-6.

¹³ Diana Raffman, *Unruly Words: A Study of Vague Language* (Oxford: Oxford University Press, 2014), p. 2.

Raffman opines that this is ‘an unclarity about the boundaries of things’;¹⁴ Keefe and Smith similarly take philosophically interesting vagueness to have ‘to do with borderline cases or with the lack of sharp boundaries’.¹⁵ This proposal will be discussed in due course.

But, as Khatchadourian notes, the philosopher, while disregarding many forms of everyday vagueness, welcomes in other examples which would not count to the layman. For instance, ‘red’ is a paradigmatically vague concept for a philosopher, but would usually not be picked up on in conversation. van Rooij states that ‘no linguistic expression whose meaning involves perception and categorization can be entirely free of vagueness.’¹⁶ Russell went even further, famously saying that he wanted ‘to prove that all language is vague’.¹⁷ Of course, if all language was vague in the everyday sense, then ‘the distinction between ‘vague’ and ‘not vague’ or ‘precise’ words would not have arisen in ordinary language.’¹⁸ Thus the everyday notion will apply to instances which the philosophical does not, and *vice versa*. In defining philosophical vagueness, then, we must move away from questions of usage, unformativity, and pragmatics, and search for deeper features which may sometimes surface in these ways.

0.1.2 Sorites-susceptibility

Perhaps the most obvious route for defining vagueness is in terms of sorites-style paradoxes. After all, if called upon to explain philosophical debates about vagueness to a member of the public, the easiest way is to run through a sorites argument on, say, the property of being bald. This is of the following form:

1. A man with 0 hairs on his head is bald.

¹⁴ Ibid.

¹⁵ Keefe and Smith, ‘Theories of Vagueness’, p. 5.

¹⁶ Robert van Rooij, ‘Vagueness and Linguistics’, *Vagueness: A Guide* ed. G. Ronzitti (Dordrecht: Springer, 2011), p. 124.

¹⁷ Bertrand Russell, ‘Vagueness’, *Australasian Journal of Psychology and Philosophy*, 1:2 (1923), p. 84.

¹⁸ Khatchadourian, ‘Vagueness’, p. 138.

2. If a man with 0 hairs on his head is bald, then a man with 1 hair on his head is bald.
3. If a man with 1 hair on his head is bald, then a man with 2 hairs on his head is bald.

:

:

:

10,001. If a man with 9,999 hairs on his head is bald, then a man with 10,000 hairs on his head is bald.

Therefore: A man with 10,000 hairs on his head is bald.

Each premise of the argument seems plausible, yet together they entail the absurd conclusion that a man with an abundance of hair on his head (the argument can run to any number you wish) is classified as bald.

Such arguments seem inextricably bound up with the notion of vagueness under investigation; the concepts which can play the role of generating the arguments appear to be precisely the concepts we wish to discuss. Why not, then, merely define vagueness as being susceptible to sorites reasoning; or, to put it another way, as being able to generate a sorites paradox? This would appear a very neat and intuitive characterisation.

There are problems, however. As Eklund points out, it is not clear what being susceptible to sorites reasoning means; it cannot mean 'that there are *sound* versions of the sorites paradoxes in which this predicate is crucially used, for the sorites reasoning cannot be taken as sound on anyone's account'.¹⁹ Presumably, Eklund says, it must mean that it is possible for a sorites argument to be made using the concept which has intuitive pull. It is hard to see, though, how this intuitive pull is a helpful

¹⁹ Matti Eklund, 'Characterizing Vagueness', *Philosophy Compass*, 2.6 (2007), p. 902.

notion to bring in. Firstly, it seems a rather shaky and unquantifiable concept on which to base a definition; what percentage of respondents need to assent to the argument before it is classed as intuitive, and how strong does their assent need to be? Secondly, it is surely a contingent psychological fact regarding us that we find some chains of reasoning compelling and others not, whereas plausibly it is necessary that some concepts are vague (I am not committed to it being the case that some concepts are necessarily vague, but a neutral characterisation should not rule it out altogether, by equating vagueness with such contingent psychological factors).

In addition, as Greenough puts it, there 'is a strong sense in which a sentence is sorites-susceptible *because* it is vague, and not vice versa'.²⁰ Defining vagueness as sorites-susceptibility appears to be a bad case of putting the cart before the horse. Our investigation could well be thought of as an attempt to identify the feature of a concept which is responsible for its being sorites-susceptible; obviously, it will not do to conclude that the offending feature is the concept's sorites-susceptible nature.

Another objection, raised by Weatherson, is that sorites-susceptibility does not cover all the apparent cases of vagueness. He gives the example of 'has few children for an academic'; intuitively, this is vague, since we are not sure whether having two children is few for an academic, or one. But we are not tempted by the premises necessary to generate a sorites paradox, which would be of the form, 'if an academic with n children has few children, then an academic with $n+1$ children has few children'. We may not be sure whether it is the step from one to two, or two to three which fails, but we are sure that one of them does, and so have no temptation to draw the paradoxical conclusion that an academic with six children has few children.²¹ There is no intuitive pull behind the more general claim that one child cannot make the difference between having few children and not having few children; one child is a big enough difference within this context.

²⁰ Patrick Greenough, 'Vagueness: A Minimal Theory', *Mind*, 112.446 (Apr., 2003), p. 244.

²¹ Brian Weatherson, 'Vagueness as Indeterminacy', in *Cuts and Clouds: Vagueness, its Nature, and its Logic* eds. Richard Dietz and Sebastiano Moruzzi (Oxford: Oxford University Press, 2010), pp. 80-81.

Sorites-susceptibility, then, seems like a particularly unpromising avenue for characterising vagueness. It is hard to pin down exactly what it means for something to be sorites-susceptible; even if we could, it does not appear to cover all instances of vagueness, and moreover sorites-susceptibility is plausibly a consequence of vagueness (sometimes), rather than a cause.

0.1.3 Borderline cases

Perhaps the characterisation of vagueness in terms of the existence of borderline cases is more promising, then. It is just as intuitive as an explanation of vagueness as sorites-susceptibility; everyone accepts that predicates such as ‘bald’ and ‘red’ have cases in which they neither clearly apply nor clearly do not apply. In fact, it appears to be a more fundamental notion, the possession of which by a concept has some explanatory power, and allows the production of sorites paradoxes. It makes sense to say that the obtaining of borderline cases is what gives the intermediate premises of sorites paradoxes their plausibility.

Weatherson has tried to revivify this once-popular theory, by arguing that it is better-placed to accommodate the data on vagueness than other characterisations. For example, the case of ‘few children for an academic’, which we saw caused problems for the sorites-susceptibility approach, can easily be handled. One and two children are plausibly borderline cases of few children for an academic; the granularity of the distinctions made by this concept poses no problem.

There are other problems for the borderline cases account, though. Greenough creates an example where we stipulate that ‘the open sentence “ x is an oldster” is determinately true of every person sixty-eight years of age and over, determinately false of those persons sixty-five years of age and under, and neither determinately true nor determinately false of the remainder’.²² Eklund similarly supposes that we define ‘nice’ as applying to n if $n < 13$, and not applying to n if $n > 15$.²³ In both of these cases, if what

²² Greenough, ‘Vagueness: A Minimal Theory’, p. 245.

²³ Eklund, ‘Characterizing Vagueness’, p. 898.

is meant by a borderline case is that the concept does not determinately apply, yet does not determinately not apply, then the concepts admit of borderline cases. Someone who is aged sixty-seven is a borderline case of 'oldster', and 14 is a borderline case of 'nice'. But neither case seems like it should come under the classification of vagueness. 'Oldster' and 'nice' are precisely, if incompletely, defined; they bear none of the hallmarks of vagueness. Perhaps this intuition stems from the fact that both are artificially created, and artificially incomplete, concepts; both could easily be extended, and made complete, without offending against their nature, which does not seem to be the case with more paradigmatically vague instances.

Weatherson simply denies this fact, and argues that, for the sake of the other benefits the borderline cases characterisation brings, we should bite the bullet and classify 'nice' and 'oldster' as vague. He admits that there is an intuition that they are not, but believes that 'if the cost of respecting that intuition is that we misclassify several other terms, we should reject the intuition'.²⁴ In effect this is a cost-benefit analysis; Weatherson argues that competing characterisations of vagueness have various problems, which are worse than merely violating the intuition regarding 'nice' and 'oldster'. For our current purposes, such a manoeuvre is unnecessary; we are merely focused on explicating the central and typical cases of vagueness, so as to understand what it will be essential for an account of vagueness to explain. The more recherche, or atypical cases, can be discussed when we get into the details of particular accounts, and an ability to deal with them could be seen as a bonus.

Perhaps there is another way around the problem, though. Greenough entertains the possibility that we need a more nuanced notion of a borderline case, one that can distinguish between examples like 'oldster' and examples like 'bald'. This is achieved by invoking the idea of higher-order vagueness. Higher-order vagueness is probably best explained through example. It is given that there are men who are determinately bald, men who are determinately not bald, and men for whom it is indeterminate whether they are bald. But there may be a further level of indeterminacy; just as the addition of one hair does not change a man from bald to not, the thought is that one hair does not

²⁴ Weatherson, 'Vagueness as Indeterminacy', p. 85.

turn a determinately bald man to an indeterminately bald man. That is, there must be a further group of men for whom it is indeterminately indeterminate whether they are bald.

The idea, then, is that we could define a vague term or expression as one which gives ‘rise to borderline cases to the borderline cases, and in turn borderline cases to those borderline cases, and so on’.²⁵ This would disallow examples such as ‘oldster’ and ‘nice’, since they only generate one region of borderline cases, which has determinate boundaries. But the question is whether this informal characterisation can be cashed out in detail. Greenough discusses Hyde’s proposal that the concept of a borderline case is actually ambiguous between the type of ‘precise’ borderline cases generated by ‘oldster’, and the type of (higher-order vagueness including, genuinely vague) borderline case generated by ‘bald’:

It is not necessary to explicitly state any extra conditions in one's characterisation of vagueness to ensure compatibility with the phenomenon [of higher-order vagueness] unless one thinks that the notion of a border case is precise-unless-explicitly-qualified – and this is simply false!²⁶

The idea is that we can then characterise vagueness as giving rise to **borderline cases**, where **borderline cases** are borderline cases of the higher-order vagueness variety.²⁷

This is problematic for two reasons. Firstly, it is questionable whether ‘borderline cases’ truly is ambiguous as Hyde claims. Secondly, and more importantly, it is ‘not yet a settled question whether any respectable theory of vagueness should indeed entail that vague terms are higher-order vague’.²⁸ Many have argued that higher-order vagueness simply does not exist, or is at least highly problematic and does not behave as the informal explication would have it.²⁹ It would be highly undesirable for our supposedly neutral initial characterisation of vagueness to take a strong stand on such a controversial issue.

²⁵ Greenough, ‘Vagueness: A Minimal Theory’, p. 246.

²⁶ Dominic Hyde, ‘Why Higher-Order Vagueness is a Pseudo-Problem’, *Mind*, 103.409 (Jan., 1994), p. 40.

²⁷ Greenough, ‘Vagueness: A Minimal Theory’, p. 246.

²⁸ *Ibid.*, p. 247.

²⁹ See Crispin Wright, ‘Is Higher-Order Vagueness Coherent?’, *Analysis*, 52.3 (1992), pp. 129-139; John Burgess, ‘The Sorites Paradox and Higher-Order Vagueness’, *Synthese*, 85.3 (1990), pp. 417-474; and Section 10 of Greenough, ‘Vagueness: A Minimal Theory’.

It may seem like a more plausible way to flesh out the intuition behind characterising vagueness as the existence of borderline cases is to invoke the idea of sharp boundaries, or more specifically the lack thereof. In brief, the argument would be that the difference between ‘oldster’ and ‘bald’ is that, while both have borderline cases, one draws sharp boundaries between the determinate cases and the indeterminate cases, while the other does not. It is precisely this lack of sharpness which is characteristic of vagueness.

Wright argues that Frege often ran the two ideas together, as in the following passage:

A definition of a concept (a possible predicate) must be complete; it has to determine unambiguously for every object whether it falls under the concept or not (whether the predicate can be applied to it truly). Thus, there must be no object for which, after the definition, it remains doubtful whether it falls under the concept, even though it may not always be possible, for us humans, with our deficient knowledge, to decide the question. Figuratively, we can also express it like this: a concept must have sharp boundaries.³⁰

This conflation is plausible, says Wright, and has lingered on, but is wrong; ‘there seems no reason why having borderline-cases should imply blurred boundaries...[the existence of borderline cases] is consistent with the obtaining of a perfectly sharp distinction between cases for which it is defined and cases for which it is not’ (as we saw with ‘oldster’ and ‘nice’).³¹

On this reading, the real intuition behind the positing of borderline cases as characteristic of vagueness is what Shapiro terms the ‘no sharp boundaries’ thesis; the idea that ‘there should be *no* relevant sharp borders anywhere in the series’.³² Raffman expresses the same idea; ‘no sharp line can be drawn between (clear) red and *any* other category – even a “borderline” category’.³³ Wright calls it the ‘seamlessness intuition’; ‘unless we have an indefinite hierarchy of kinds of borderline case, it seems there will have to be sharp boundaries’.³⁴

³⁰ Gottlob Frege, *Basic Laws of Arithmetic: Derived using concept-script* (Oxford: Oxford University Press, 2013), transl. and ed. P. Ebert and M. Rossberg, § 56.

³¹ Crispin Wright, ‘On the Coherence of Vague Predicates’, *Synthese*, 30.3/4 (Apr. – May, 1975), p. 329.

³² Stewart Shapiro, *Vagueness in Context* (Oxford: Oxford University Press, 2006), p. 130.

³³ Diana Raffman, ‘Vagueness Without Paradox’, *The Philosophical Review*, 103.1 (Jan., 1994), p. 41 fn. 1.

³⁴ Crispin Wright, ‘The Illusion of Higher-Order Vagueness’, in *Cuts and Clouds: Vagueness, its Nature, and its Logic* eds. Richard Dietz and Sebastiano Moruzzi (Oxford: Oxford University Press, 2010), p. 527.

But this thesis carries its own problems. As Eklund notes, there are actually two possibilities when saying that vague expressions do not draw sharp boundaries; either they do draw boundaries which are non-sharp, or they draw no boundaries at all. If the first option is taken, then the question arises of what exactly it means to draw a non-sharp boundary; it appears that 'to say that the predicate draws non-sharp boundaries is just another way of saying that it draws vague boundaries'.³⁵ Obviously when attempting to give a characterisation of vagueness, drawing on the notion of vagueness is less than ideal. Is Eklund too quick to equate non-sharp boundaries with vague boundaries? There might perhaps seem to be a precise way of defining non-sharpness, for instance along the lines of saying that a non-sharp boundary is one which does not move from a predicate to its opposite in one step, but moves through a series of intermediate statuses. This line of thought, though, quickly brings us back into the thicket of questions regarding higher-order vagueness. Raffman's formulation above explicitly outlaws sharp boundaries between any categories, where sharp is taken to mean changing status between consecutive items in the series. Taken as such, sharp boundaries come to mean, in effect, any boundaries.

So we move to the second option, holding that vague expressions draw no boundaries. This is the route Sainsbury, for example, takes; 'The right way to characterize the vagueness of a predicate is by the fact that it classifies without drawing boundaries'.³⁶ Eklund criticises this option on the basis of its basic unclarity. What can it mean to say that 'bald' draws no boundaries, when it is a datum that it determinately applies to some men, while determinately not applying to others; in doing so, 'does it not effect some sort of boundary?'.³⁷

Sainsbury's response is ostensive; we have all seen, or can imagine seeing, the colour spectrum. While 'we can discern no boundaries between the different colours...[t]his does nothing to impede the classificatory process'.³⁸ This is initially plausible, but does

³⁵ Eklund, 'Characterizing Vagueness', p. 901.

³⁶ Mark Sainsbury, 'Is There Higher-Order Vagueness?', *The Philosophical Quarterly*, 41.163 (Apr. 1991), p. 179; see also Mark Sainsbury, 'Concepts Without Boundaries', in *Vagueness: A Reader* eds. Rosanna Keefe and Peter Smith (Cambridge: MIT Press, 1997), pp. 251-264.

³⁷ Eklund, 'Characterizing Vagueness', p. 902.

³⁸ Sainsbury, 'Concepts Without Boundaries', p. 258.

not fare well upon closer inspection. Sainsbury says that the different colours ‘stand out as clearly different, yet there are no sharp divisions’; the spectrum has ‘bands, but no bounds’.³⁹ As we have seen, a lack of sharp divisions does not entail a lack of divisions altogether. What Sainsbury calls a band appears to be what others would term a vague, or non-sharp boundary; true boundarylessness is thus not achieved.⁴⁰ Sainsbury attempts to explain more by drawing an analogy with magnetic poles, with instances of vague predicates being pulled closer to them or their contraries, but this is unpersuasive without further explicit explanation.

In any case, even if it is granted that Sainsbury is right, and the idea of a boundaryless concept is coherent or even plausible (or that non-sharp boundaries can be explicated in a way not reliant on the notion of vagueness), there is an overarching problem for a characterisation of vagueness as lack of sharp boundaries. This problem arises from the nature of the project of characterisation which we are undertaking. We are aiming for a neutral delineation of vagueness, one that everyone can agree upon as a basis for identifying the various instances of vagueness. Only later in the thesis will I attempt to offer a substantive theory of what metaphysical vagueness consists in, such that its instances fall under this characterisation. This is the same objective as Greenough, who is attempting to deliver what he calls ‘a minimal theory of vagueness’, to act as a neutral starting point for further debates.⁴¹ Taking this as the aim, then, the no sharp boundaries thesis is unacceptable, taking as it does a substantive stand on the nature of vagueness. A characterisation of vagueness as lack of sharp boundaries would immediately rule out epistemicist theories of vagueness, for instance, which hold that there are sharp boundaries, but they are essentially unknowable for some reason.

³⁹ Ibid.

⁴⁰ Above, ‘non-sharp boundaries’ were dismissed as being characteristic of vagueness due to their involvement of vagueness itself, not dismissed from the possibility of existing.

⁴¹ Greenough, ‘Vagueness: A Minimal Theory’, p. 446.

0.1.4 Tolerance

But perhaps there is a related idea which can be more useful. It may quite rightly be pointed out that focussing on a lack of sharp boundaries is confusing cause and effect, in the same manner that focussing on sorites-susceptibility was. That is, it seems plausible that concepts lack sharp boundaries because they are vague, rather than being vague because they lack sharp boundaries. However, there is another notion which can explain why a ‘no sharp boundaries’ account was tempting; this is tolerance, the ‘notion of a degree of change too small to make any difference, as it were’.⁴² The idea then is that something is vague if it is tolerant, and if it is tolerant then it will necessarily not draw sharp boundaries.

But what is tolerance more formally? Wright gives the following definition:

F is tolerant with respect to Φ if [and only if] there is some positive degree of change in respect of Φ insufficient ever to affect the justice with which F is applied to a particular case.⁴³

This notion is a generalisation of the intuition behind the intermediate premises of the sorites paradox. Using the example of ‘bald’ from 0.1.2, the generalisation would be ‘If a man with n hairs on his head is bald, then a man with $n+1$ hairs on his head is bald’; this is true because one hair is a degree of change small enough that it is insufficient to change whether the term ‘bald’ should be applied in a particular case. It also necessarily leads to a denial of sharp boundaries, since a sharp boundary just is a point where a small change does indeed ‘affect the justice with which F is applied’.

Tolerance is not susceptible to the problem of concepts such as ‘oldster’ and ‘nice’. ‘Oldster’ would not qualify as vague under a tolerance characterisation, since there is no degree of change insufficient ever to affect its application; it only requires one millisecond for someone to turn 65, and thus change whether ‘oldster’ applies to them. It does, however, suffer from some of the other previously discussed problems. Firstly, just as no sorites argument can truly be said to be sound, no term can truly said to be tolerant. As mentioned above, tolerance licenses the intermediate premises of the

⁴² Wright, ‘On the Coherence of Vague Predicates’, p. 333.

⁴³ Ibid., p. 334.

sorites argument and leads to paradox. If 'bald' really is tolerant, then through a series of changes of one hair, a man with bountiful locks can truly be said to be bald.

Secondly, the existence of tolerance is a substantive claim about the nature of vagueness, just as the no sharp boundaries thesis was. If in order to be vague something has to exhibit tolerance, then epistemicist accounts of vagueness are ruled out off the bat; for an epistemicist things are never actually tolerant, but merely appear so. In developing a neutral characterisation, this is undesirable. Of course, if there is no plausible candidate for a such a neutral account, then we may well have to judge which substantive issues we must take a stand on in order to proceed.

Luckily, though, there is such a candidate, based indeed on tolerance itself. Greenough notes that 'all partisans to the dispute can or do, for whatever reason, accept that vague terms draw no known boundary across their associated dimension of comparison'.⁴⁴ Since this is the case, it is reasonable to move to a notion of *epistemic* tolerance; that is, tolerance whereby the small change is never enough to make a *known* difference to application, rather than an actual difference. The proposal then becomes:

F is *epistemically* tolerant with respect to Φ if and only if any small positive degree of change with respect to Φ is insufficient to make a known difference to the correctness of applying F to a particular case.⁴⁵

Since it is the case that tolerance entails epistemic tolerance – after all, if a change makes no difference to the correctness of applying F, it necessarily makes no known difference to the correctness of applying F – an epistemic tolerance characterisation can handle all the cases which a genuine tolerance characterisation did.

Paradigmatically vague concepts such as 'bald' will still be classified as vague, since a change of one hair cannot make a known difference to its application. Equally, it does not fall prey to the counterexample of partially defined concepts; a change of one minute does make a known difference to the application of 'oldster', since the conditions of applications for that term, and its boundaries, are known.

⁴⁴ Greenough, 'Vagueness: A Minimal Theory', p. 257.

⁴⁵ *Ibid*, p. 258.

Can it also deal with Weatherson's problematic case of 'having few children for an academic'? This case was originally raised for a sorites-susceptibility characterisation, but would appear to be equally applicable. It might appear that moving from one child to two, or two to three, may be enough to make a known difference between the applicability of the concept, and this is the smallest possible change available. As such, there is no positive change which makes no known difference to application, and yet the concept seems intuitively to be 'vague'; epistemic tolerance has miscategorised.

There are a couple of replies available, I believe. One is at the level of intuitions. I find it hard to fully accept that moving from one to two children, or more plausibly two to three, can make a *known* difference to the correctness of the application of 'having few children for an academic'. To be sure, it could be known that moving from one to three makes a difference, and suspected that one of the two steps is enough by itself, but not known. There is at least enough of a doubt there. Moreover, the concept of 'having few children for an academic' is not a paradigmatic case of vagueness; the intuition that any characterisation should categorise it as vague is nowhere near as strong as for, say, 'old'. Thus if it is the case that epistemic tolerance miscategorises it, this is not a grave objection (particularly since no other characterisation we have discussed is any better placed to deal with the example).

The second reply is from Eklund, and plays on the fact that, given another context, 'having few children for an academic' could certainly be classified as vague through epistemic tolerance.⁴⁶ In a world where every family has between ten and twenty children, the difference of one child between, say, fifteen and sixteen will make no known difference. Plausibly it is the existence of this possibility which licenses us to classify 'having few children for an academic' as vague in the current context; we are not just interested in actual epistemic tolerance, but possible epistemic tolerance.

So overall, then, Greenough's characterisation of vagueness as epistemic tolerance fares better than the others on offer. In general it is preferable to begin with a theory-neutral account. Greenough sees this as an admirable aim in its own right, while I am

⁴⁶ Eklund, 'Characterizing Vagueness', p. 903.

using it as a staging post on my path to a substantive theory of metaphysical vagueness. This is for two principal reasons. Firstly, I think that, all else being equal, it is preferable to be working with mutually agreed terms and definitions until as late as possible in your enquiry; this can be seen as a form of anti-revisionism, or at least an underlying resistance to unnecessary revisionism. Secondly, I wish to leave open the possibility that not all instances of vagueness are generated by the same underlying type of phenomena. If our characterisation of vagueness took a stand on what it actually consists in, or what generates it, then this would be ruled out. To put it another way, it should be an open question as to whether all of the things which we characterise as interesting, philosophical, vagueness have the same fundamental nature (and later I will argue that they in fact do not). To this end, a neutral, epistemically-based characterisation is preferable at this stage.

0.2 Categories of vagueness

Now that we have a delineation of the phenomenon of vagueness that is of interest to philosophers, we can outline the ways in which philosophers have tried to explain it. As mentioned earlier, they are generally grouped into three categories: vagueness in our language (semantic vagueness); vagueness in our knowledge, or lack of knowledge, of some certain facts (epistemic vagueness); and finally vagueness in the world itself, independent of language and knowledge (metaphysical vagueness). For the time being, I will merely register them as they are commonly given; later on I will critically discuss them as theories.

0.2.1 Semantic

Semantic or linguistic vagueness is the most widely discussed of the three main categories of vagueness, and the only one whose existence is accepted by (almost) all parties to the debate. Presumably this is because it is closest to the everyday meaning of the term 'vague'.

The cases which are generally taken to be semantically vague, and which most of the literature focuses on, are those such as ‘tall’, ‘heap’, ‘bald’, and ‘red’. These are the sort of terms Lewis had in mind when he talked of semantic indecision (as well as potentially many others).⁴⁷ In the most general terms, the idea behind semantic vagueness is that there is something in the meaning of the above terms that gives rise to epistemic tolerance (to use the hopefully neutral characterisation of vagueness arrived at earlier). Some propositions such as ‘Man *x* is bald’, will, due to the semantics of ‘bald’, appear to lack a definite truth-value, and we will be hard-pushed to come down on whether we would wish to assert the statement, or to agree with it when asserted by another. Specifically, to distinguish these statements from other instances where a truth value is lacking, the vagueness arises from the fact that there is something about the semantics of ‘bald’ which means that while it clearly applies in some cases, and clearly does not in others, there are cases where a series of small changes in the person being described will lead to us not being able to tell whether it applies or not. There are, of course, different explanations on offer of what exactly is happening in these cases, and why – we do not need to enter into those discussions here.

Sorensen has drawn a further distinction between what he calls *product* and *process vagueness*. Briefly, a *product vague statement* is ‘a statement that expresses a vague proposition’; a *process vague statement* is ‘a statement that vaguely expresses a proposition’.⁴⁸ The more familiar case, according to Sorensen, is product vagueness without process vagueness, as in the following example: ‘A six feet tall man is tall’. There is a specific proposition expressed by this statement, and that proposition exhibits the kind of interesting vagueness mentioned above.

An example Sorensen gives of process vagueness is the following: ‘Men are taller than women’.⁴⁹ When this statement is uttered, it seems clear that it does not express the proposition that all men are taller than all women; but it is unclear as to what proposition it is intended to express. Examples of possible intended propositions are ‘Most men are taller than most women’, and ‘The average height of men exceeds the

⁴⁷ See David Lewis, *On the Plurality of Worlds* (Oxford: Blackwell, 1986), p. 212.

⁴⁸ Roy Sorensen, ‘Process Vagueness’, *Linguistics and Philosophy*, 13.5 (Oct., 1990), p. 589.

⁴⁹ *Ibid.*, p. 591.

average height of women.⁵⁰ This sort of vagueness can, I believe, explain some everyday instances of vagueness which might be classified as uninformative. For instance, imagine I am in a group where everyone is introducing themselves by explaining where they were born, grew up, went to school and university, and where they live; when it is my turn, I simply state, 'I'm from London'. I have been vague, because it is unclear what proposition I have expressed; plausibly it could be 'I was born in London', or 'I grew up in London', or 'I live in London', or any number of combinations.

Sorensen also introduces the notion of a *process borderline statement*. This is a statement *S* which is a borderline signifier of a proposition *p*, and where 'knowing whether *S* signifies *p* is a necessary condition for knowing *S* to be true or knowing *S* to be false'.⁵¹ An example given is the following: 'Andrew Johnson escaped impeachment by one vote'. It is disputed whether impeachment is the process of being indicted, or being convicted. According to Sorensen, the statement above is thus a borderline case of ambiguity. Moreover, it cannot be decided to be true or false without this vagueness being dissolved; if impeachment is indictment, it is false, whereas if impeachment is conviction it is true.

The notion of process vagueness seems to be applicable to many cases which fall under the general notion of unclarity which we take to be the defining feature of everyday vagueness.⁵² It is product vagueness, however, which seems to be aligned with the philosophically interesting notion(s) of vagueness which will be the subject of the rest of this thesis.

0.2.2 Epistemic

Another category of vagueness which is proposed and discussed is epistemic vagueness. Epistemic vagueness is something of an outlier compared to linguistic and

⁵⁰ Ibid., pp. 591-592.

⁵¹ Ibid., p. 591.

⁵² For more on the details of process vagueness, see the rest of Sorensen's article, and its application with regards to David Tuggy, 'Ambiguity, Polysemy, and Vagueness', *Cognitive Linguistics*, 4.3 (1993), pp. 273-290.

metaphysical vagueness, in that it states that there is a fact of the matter about an apparently borderline case of some term. According to an epistemicist, it is either true or false that some man is bald; the vagueness consists in the fact we are unable to know whether it is true or false. Of course, there are many things which we cannot know, and even many things which are true or false but unknowable (the Twin Primes conjecture, perhaps). The important feature in cases of vagueness is the reason that we cannot know.

For Williamson, this is the fact that in cases of inexact knowledge, the mechanism for judging something to be the case is unreliable. In his example, when I am in a stadium and judging how many people are there, I may state that there are m people. I may even be right. But I would still have judged there to be m people if there were actually $m+1$ or $m-1$ people; my method for judging is not sensitive to the truth, and so is unreliable. I would not claim to know how many people are in the stadium, even if I guessed right. Similarly, in cases of vagueness, our method for judging is unreliable. There may well be a fact that, say, a man with 4,000 hairs on his head is bald, while a man with 4,001 is not, and we may use 'bald' correctly of someone with 3,999 hairs. We might just as well have used 'bald' of a man with 4,002 hairs, however, in which case we would have used it incorrectly. Our method of judgement is unreliable, and so can never furnish us with knowledge.⁵³

0.2.3 Metaphysical

Let us now turn to metaphysical vagueness, which is the primary focus of this thesis. In general, it is the idea that there may be vagueness genuinely existing in reality, independent of any thought or representation.

Perhaps the most commonly discussed case is mereological vagueness. That is, the idea that it can be vague whether some thing is a part of some other thing. To motivate this notion, consider a wispy cloud; plausibly, for some water droplet on the very edge

⁵³ Timothy Williamson, *Vagueness* (New York: Routledge, 1996), pp. 230-231.

of the formation, it is vague whether that particular droplet is actually a part of the cloud, or is simply too far away from the droplets in the main body of the cloud. Or consider 'some objects - this rock on the summit, that rock in the penumbra, and lots of other stuff - which are candidates for composing Kilimanjaro. The rock on the summit is definitely a part of Kilimanjaro, while the rock in the penumbra is not definitely a part of it'.⁵⁴ Plausibly, for some rock located between the flat plain and the summit, it will be vague whether that rock is a part of Kilimanjaro or part of the plain.

A separate type of metaphysical vagueness is existential vagueness. This is the notion that it can be vague whether an object exists or not. If it is supposed that a living being is something over and above the matter that makes it up, then plausibly, since there is no specific point of death, it will be vague at some point whether the living being exists. Perhaps another case could be a small group of water droplets hanging in the air, roughly in formation but not very closely aligned; plausibly, it could be vague as to whether a cloud exists there.

Another type of metaphysical vagueness which is often mentioned is compositional vagueness; that is, vagueness as to whether some group of things composes a larger thing. The cases which motivate this are often the same as the motivating cases for existential vagueness; plausibly, in the cloud example, it is vague whether the cloud exists precisely because it is vague whether the droplets compose a larger object, namely a cloud. For a case where existential vagueness does not seem to be present in the same way, consider Lewis's example of the outback. He supposes that it could be the case that there are multiple collections of matter (rocks, land, scrub, etc.) which are equal candidates for composing the outback; for each specific collection, it is vague whether they compose a larger object (of course, not according to Lewis; for him each collection is merely a candidate to be the referent of the term 'outback').⁵⁵ However, it is not vague as to whether the outback exists; it determinately does.

Yet another proposed form of metaphysical vagueness is vague identity. Perhaps the most plausible case of this involves a teleporter. Suppose that you step through the

⁵⁴ Chad Carmichael, 'Vague Composition Without Vague Existence', *Noûs*, 45.2 (Jun., 2011), p. 316.

⁵⁵ Lewis, *On the Plurality of Worlds*, pp. 212–13.

teleporter, which malfunctions and spits out two of 'you' at separate locations. We might want to say that in this case, it is vague as to whether you are identical with one, other, or both of the beings who have been spat out. Or, to utilise Lewis's outback case again, you might want to say that it is vague as to whether the outback is identical with any one of the candidate collections of matter which vaguely compose it.

A final type of metaphysical vagueness sometimes mentioned is the vagueness of the open future. In some sense, it may be said that statements regarding the future are vague, if the future truly is open. Moreover, this is due to the inherent unsettledness of the future itself. The future lacks clarity in the sense that it is undetermined. However, this does not come under the characterisation we are working with, of (possible) epistemic tolerance. While it may well be the case that the open future is an example of metaphysical indeterminacy, it is not epistemically tolerant in any obvious way. Any small change in the possible future scenario will always leave it as indeterminate, because there is nothing determinate about the future. This is not a problem, as it is intuitively fairly clear that the open future is of a different kind to the other cases of metaphysical vagueness; this is further evidence that epistemic tolerance is a good characterisation of the notion we wish to investigate. Incidentally, it also provides evidence that vagueness cannot simply be identified with indeterminacy.

It appears that there may be links between the various sub-types of metaphysical vagueness. For instance, it seems hard to believe that mereological vagueness could be entirely unconnected to compositional vagueness, assuming that both exist. Presumably, if it is vague whether some thing is a part of another thing, then the composition of the larger thing will be vague in some way. Additionally, it seems plausible that there will be cases where if it is vague whether some x s constitute a larger y , then it is vague whether that larger y exists (for this to be the case, there would have to be no other candidate group of z s to compose the larger y , nor another candidate larger object y' which the x s may compose).

Relatedly, it is famously argued by Evans (and Salmon) that the existence of vague objects necessarily leads to vagueness about identity.⁵⁶ Of course, Evans takes that to be a *reductio* of vague objects, since he finds vague identity to be incoherent. Those debates will be examined in section 3.1.1. For now it is enough to note that there is at least *prima facie* reason to think there is a link between vagueness in objects (whether that be compositional, mereological, or existential) and vagueness in identity.⁵⁷

0.3 Connections between different types of vagueness

The purpose of this chapter is to clear the ground, and give ourselves an idea of what it is that any convincing account of vagueness needs to accomplish. Now that we have a characterisation of what makes something interestingly vague, and have identified apparent examples of it, we should look at whether there are any obvious links between different types of vagueness. If there are, then any convincing account will need to explain what those links consist in; or, at the least, explain why we were led to believe that there were such links.

For the purposes of this investigation, any links between sub-types of linguistic or epistemic vagueness are of secondary importance. But we should still pay attention to possible connections between metaphysical vagueness and the other two categories. For instance, if there really is existential vagueness, this might lead to instances of linguistic vagueness, since it will be vague whether a name designating the vaguely existing object truly refers or not. Depending on your semantic theory, this might raise various issues. More generally, it may be the case that metaphysical vagueness will entail semantic indeterminacy in many instances. After all, the guiding thought behind semantic vagueness is often that the concept has not been completely defined; if the world truly is vague (in whatever way), there may simply be no way to completely define it, even if all the requisite semantic work is done. Depending on your semantic theory,

⁵⁶ Gareth Evans, 'Can There Be Vague Objects?', *Analysis*, 38.4 (Oct., 1978), p. 208; Nathan Salmon, *Reference and Essence* (Oxford: Blackwell, 1982).

⁵⁷ Katherine Hawley, for example gives an argument that vagueness in existence will entail vagueness in identity, if the principle of extensionality is retained; see Katherine Hawley, 'Vagueness and Existence', *Proceedings of the Aristotelian Society*, 102 (2002), p. 133.

and how exactly you cash out metaphysical vagueness, this may lead to widespread semantic vagueness, or introduce a new distinction between metaphysically-generated semantic vagueness and semantic vagueness originating in our thinking.

Similarly, it appears that in many cases metaphysical vagueness will entail at least a form of epistemic vagueness. If there simply is no fact of the matter for some circumstance, because there is a fundamental unsettledness, then there will be a necessary ignorance regarding that fact. It is an open question whether this would fall under previous theories of epistemic vagueness, or require a new distinction, but it is a question which merits discussion.

0.3.1 How to tell which type of vagueness is which?

This leads us to the problem of identifying the origin of the vagueness. So far we have been taking things at face value; if it appears to be metaphysical, we classify it as such. But when we go further and attempt to not just identify things as apparently vague, but really explain what is going on, we will need to separate out the metaphysical from the semantic. As we have just seen, they seem to be intricately linked in many ways. Moreover, ordinary language seems to be riddled with vagueness, so how can we ever be sure that what we're talking about is vague simply because of this, and not in and of itself? This is a methodological problem, without a neutral answer which will be accepted by everyone.

The general technique to be followed is to sharpen the language and concepts we are using as much as possible, and see if vagueness still remains. For instance, if we can use thought experiments which take out explicit representation or expression, this will be helpful. Obviously a semantic imperialist (someone who believes that all vagueness is fundamentally semantic) will say that any thought experiment is necessarily mediated by our concepts, which are inherently vague, and so we cannot penetrate through to the metaphysical level. If semantic vagueness is taken to be very widespread, as it is by many, this will be a major problem. On one level I agree, in that we are always using vague concepts. But I disagree with the assumption that this

means we can discount the existence of metaphysical vagueness. Firstly, if the objection is correct, then it means we definitely can't rule out metaphysical vagueness, because that is equally as unjustified as ruling it in. If we cannot penetrate through the haze of semantic vagueness, then no conclusion can be reached. Secondly, proponents of semantic vagueness typically do not take its existence to preclude metaphysical speculation in areas other than vagueness. Thus it is ad hoc to argue that it does so in precisely in the area of metaphysical vagueness. Overall, I think that with enough awareness of what sort of vagueness is exhibited by the semantics (more details will be discussed at the apposite moments), we can compensate for it, and recognise when the vagueness is not of that type.

0.4 Formulating desiderata

By now, we can appreciate the importance of the question motivating this thesis: should we accept the existence of genuine metaphysical vagueness? As we have seen, vagueness appears to be a fairly widespread phenomenon, and moreover one which is puzzling, challenging many of our assumptions about how to think and reason about the world. A convincing account of vagueness, how it functions, and how we should navigate it, is therefore a valuable aim. In addition, a fair number of the instances of vagueness we identified above seem, at least on their face, to be located in the world. We are thus forced to address the question of whether this is as it seems; is the world itself vague in some way? If so, how can this be explained, and what does it mean for the rest of our metaphysics, our systems of reasoning, and so on? If not, why does it appear to be? What is going on in other realms (semantic or epistemic, for example) which leads to the mistaken idea that the world itself is vague? Getting clear answers to these questions will help illuminate many enduring philosophical problems.

But what do we mean by a 'clear answer'? It will be useful to draw up some desiderata for any convincing account of metaphysical vagueness. So far we have reached the following point: vagueness, in the interesting philosophical sense, is characterised as (possible) epistemic tolerance. This is the notion which will be under investigation in

this thesis. There are various phenomena which exhibit epistemic tolerance. Broadly they fall under three categories; metaphysical, semantic, and epistemic. Within those categories there are various sub-types, which interact with each other in various ways. In addition, there is interaction between the categories, notably with metaphysical vagueness apparently leading to the other two forms in some way.

The first desideratum for an illuminating answer to the question under discussion, and the most important, is that it must explain why the various metaphysical phenomena which exhibit epistemic tolerance occur. This needn't necessarily be a unified answer; plausibly, there are different underlying causes for the surface instances of tolerance. A non-unified account, though, must offer an explanation of why these separate causes produce a similar epistemic effect.

Secondly, a convincing account must explain why some of the instances of metaphysical vagueness interact with each other in the manner that they do. That is, an account of compositional vagueness (vagueness as to whether some things compose another thing) must also give an explanation of its relation to existential vagueness (vagueness as to whether an object exists); when does one lead to the other, and why. Of course, it might be the case that not all the apparent connections are actual, in which case an explanation is needed of why we have the intuition that they are.

Thirdly, and relatedly, a convincing account must explain why metaphysical vagueness can lead to, or impact upon, semantic and epistemic instances of vagueness. This explanation could take many forms. As I have stressed throughout, our characterisation of vagueness is theory-neutral and superficial; an explanation of the phenomena is consistent with denying that there is anything fundamental occurring, while explaining how the surface phenomena arise.

These three desiderata are roughly in order of importance, but they are all serious considerations, which a rigorous account must satisfy in some way or another. However, those are just the criteria specific to a theory of metaphysical vagueness. Of course, there are many general specifications which a convincing account will satisfy. For instance, one might think that ontological parsimony is a worthwhile goal, and a

strong plus-point for any theory exhibiting it (though there is always the question of what ontological parsimony actually boils down to). Others might value theoretical parsimony; an account which invokes the fewest notions (or perhaps fundamental notions) while offering sufficient explanatory power is to be applauded, on this view. Then there are considerations of revisionism; in particular, it is often objected that metaphysical vagueness necessitates moving to non-classical logics, which many are loathe to do. On the flipside of anti-revisionism is the goal of respecting the motivating intuitions and factors driving the discussion; that is, making a genuine attempt to explain the phenomena under examination, rather than merely explaining them away.

These general methodological issues will be dealt with throughout the thesis. It is not necessary to rank them in order; generally, there is a weighing-up to be done for each theory. What something gains in ontological parsimony, it can often lose in the theoretical kind. The question of logical revision will receive an in-depth discussion later on, as it is seen by many to be one of the most critical questions facing theories of metaphysical vagueness, and vagueness in general.

Now, some may object that all of the above desiderata are in fact sideshows to the one actual objective we should be going for in our account – that it be a true reflection of the metaphysics of the world. The thought would be that in any metaphysical investigation, the aim is to accurately describe the world, and while this might be achieved by an elegant theory, on other occasions the world might not be so obliging and the correct account may not be able to be superficially sleek.

I have some sympathy with this view. After all, I believe that (at least parts of) the world genuinely is metaphysically vague, and being told that a rival, say, semantic imperialist theory is actually more parsimonious and requires less revision to classical logic will not change my mind. However, I think the three desiderata I set out above, which are specific to this investigation, capture this intuition. That is, I think that the account of vagueness which best explains the observed phenomena will get closest to the truth of the matter. Conversely, any theory which does not accord with the metaphysical truth will inevitably lead to counter-examples and issues, perhaps only solvable with additional complications or ad hoc reasoning, which will remove any supposed

theoretical advantage. Consider the example of pre-Copernican geocentric models of planetary motion, which were forced to build a complex system of epicycles to account for the observed phenomena. In a similar way, apparently elegant accounts of (non-metaphysical) vagueness may be shown to be unsatisfactory through the process of carefully considering their implications when confronted with instances of (apparent) metaphysical vagueness.

0.5 Structure of the thesis

We have arrived at the first staging-post of our investigation, then. I have set the terms of discussion, and defined our primary objective: we are searching for a theory of metaphysical vagueness which can fulfil the desiderata outlined above. I will proceed in the following manner.

In Part A, I will first critique the two dominant theories which answer the research question in the negative; that is, those which argue that we should not accept the existence of genuine metaphysical vagueness, because there are other explanations of the phenomena of vagueness available which they take to be superior. First, in Chapter 1, I will discuss supervaluationism, the primary account which seeks to locate vagueness in semantics, and argues that it is the fact that there are multiple acceptable precisifications of vague terms which leads to the apparent contradictions in instances of vagueness. Then, in Chapter 2, I move on to epistemicism, the idea that vagueness is actually an artefact of limitations of the way we obtain and claim knowledge in certain circumstances. In both cases, I will show that while each has its benefits, and can offer on one level a concise explanation of the widespread issue of vagueness, they both suffer from some problems of consistency, and neither are convincing enough to outweigh the negative fact that they don't truly engage with metaphysical vagueness, or offer an account of why it (mistakenly, on their view) appears to exist.

In Part B, I move on to metaphysical vagueness directly. I begin in Chapter 3 with a brief defence of the concept of metaphysical vagueness in general, in the face of general arguments purporting to show its incoherence or implausibility. In Chapter 4 I then look

at previously offered accounts which do answer yes to research question – that is to say, theories of metaphysical vagueness which accept that genuine metaphysical vagueness exists. These range from quite specific theories about, for example, the possibility of mereological simples being vaguely located at certain points or regions, to general accounts that attempt to provide a plausible model for talking about metaphysical vagueness however it may manifest. I conclude that while most of these accounts offer interesting perspectives on metaphysical vagueness and its various instances, none are able to offer a convincing general account that respects the fact that such vagueness is present in the fundamental make-up of the world (and therefore should also be fundamental in our account of the world).

In Chapter 5 I propose my account of genuine metaphysical vagueness, which is an amalgam of Jessica Wilson's framework of the (determinate) obtaining of vague states of affairs with an extension of Michael Tye's account of metaphysical vagueness. Tye's view, which I find to be on the right lines of what a fleshed out, convincing positive answer to the research question would look like, is focused on vague objects – I expand this to become a wider account of metaphysical vagueness and its various manifestations, and posit that genuine metaphysical vagueness consists in the obtaining of indeterminate states of affairs, which are states of affairs containing genuinely vague entities.

I conclude by reiterating the benefits of answering the research question affirmatively: when it is possible to give a convincing, coherent account of metaphysical vagueness which supports pre-theoretic intuitions and observations of how the world is, it is simpler to accept it, rather than deny it for technical, theoretical reasons to do with preserving classical logic, or some such. The most straightforward, accurate description of the world should include genuine metaphysical vagueness.

Part A

Metaphysical vagueness is often taken to be unmotivated, on top of being implausible and/or incoherent. The basic idea is that everyone agrees that there is vagueness inherent in the concepts we use, whether it be located in language (the majority), or the boundaries of our knowledge (the minority). Since this is the case, there is no need to posit a separate, further type of vagueness (metaphysical vagueness); everything can be explained through the vagueness of concepts. I will attempt to undermine this thought by pointing out that the explanations given by linguistic and epistemic vagueness are not as clean and powerful as often supposed. Once we dig into certain examples of vagueness, and the specifics of the linguistic or epistemic accounts, we find that there is often more complexity than first imagined.

This complexity opens up a space for metaphysical vagueness to do important explanatory work. Depending on your other metaphysical commitments, the admittance of metaphysical vagueness may well provide a simpler overall explanation of the phenomenon of vagueness. This will not be, and indeed does not attempt to be, a knock-down argument. Instead, it challenges the abrupt assertion that metaphysical vagueness is unnecessary; perhaps it is, but this needs to be the subject of investigation, and a cost-benefit analysis. I will begin with a discussion of linguistic supervaluationism, as it is the most widely accepted model of semantic vagueness, and its proponents take it to be widely applicable.

1. Supervaluationism

Supervaluationism can be thought of as a sort of truth-value gap theory. That is, a supervaluationist accepts that unlike in the extreme cases of hirsuteness, for example, which everyone agrees are definitely true (or definitely false), the borderline cases simply do not have a definite truth-value. They are neither definitely true nor definitely false, because the semantic concept of ‘being bald’ is insufficiently detailed to determinately apply or not apply. However, this is not the end of the story.

The key supervaluationist insight is to look not just at how our language actually assigns meaning to those terms, but the possible ways in which it *could* assign meaning to them. In our actual language, vague terms are incomplete, in that their extension and anti-extension are not exhaustive. But it is at least possible to imagine a language where vague terms are no longer vague, and have been made fully precise; for every person, they are either tall or not tall. This hypothetical (section of) language is what’s called a ‘precisification’; a theoretical way of making ‘tall’ precise. What we are interested in are the admissible precisifications; that is, the precisifications which do not disagree with any of our everyday judgments regarding the actual use of ‘tall’. Any precisification which, for instance, classified the five feet tall man as tall contradicts our actual use of ‘tall’, and as such is inadmissible as a hypothetical precisification. As Keefe puts it, ‘Though our practices do not determine a precise extension to “tall”, they do determine a (vague) range within which the precise extension would have to be if there were one’.⁵⁸

Now that this basic machinery is in place, we can outline the supervaluationist account. In order to work out the truth value of statements in actual language, we look at the values they are assigned in each admissible precisification, and arrive at a ‘supervaluation’. Definitely true statements, such as ‘A seven feet tall man is tall’, are true because they are true on all admissible precisifications. Likewise, ‘A five feet tall man is tall’ is false because it is false on every admissible precisification. In borderline cases, the statements will be true on some precisifications and false on others; this is

⁵⁸ Rosanna Keefe, *Theories of Vagueness* (New York: Cambridge University Press, 2000), p. 153.

why they are not assigned a truth value. It is in this way that the borderline region is represented according to supervaluationism.

One important motivating factor for this account is its ability to handle what Fine calls ‘penumbral connections’. Fine points out that traditional three-valued systems of truth struggle to cope with assigning the right values to various statements which concern relations between elements in the borderline region. For instance, let’s say that Bob is 5’11”. On a three-valued account, both the statements (P) ‘Bob is tall’ and (Q) ‘Bob is not tall’ are indefinite. Since, in general, the conjunction of two indefinite statements is itself indefinite, the three-value theorist is committed to saying that P & Q (‘Bob is tall’ and ‘Bob is not tall’) is indefinite. This is an unappealing conclusion; even though Bob is in the borderline region (or penumbra), it is still a fact that he cannot be both tall and not tall, so P & Q should be assigned the value of false, not indefinite. The three-valued theorist could, of course, accept that P & Q is false, but only at the cost of turning ‘&’ into a non-truth-functional notion. Fine also takes this to be an issue for degree-theoretical systems, which assign P & Q anything more than a 0 degree of truth.⁵⁹

However, supervaluationism does have a method of explaining the values we think should be assigned to these statements regarding penumbral connections. In the case of P & Q, according to supervaluationism, in order to calculate its truth value, we need to evaluate it at every admissible precisification. That is, we need to see whether P & Q is true or false with regards to every admissible way of making the extension of ‘tall’ complete and precise. Now, we can readily see that on every such precisification, either Bob will be classified as tall or classified as not tall, and never both. After all, a precisification which categorises Bob as both tall and not tall will not be admissible, because it clashes with our understanding of the term tall (and a putative precisification which classifies him as neither is not complete, and so is not truly a precisification). Thus, on some admissible precisifications, P will be true, and Q will be false. On the admissible precisifications where Q is true, P will be false. Therefore, whichever admissible precisification we look at, P & Q will always come out as false; and so, given our supervaluationist rules, P & Q is false (simpliciter).

⁵⁹ Kit Fine, ‘Vagueness, Truth and Logic’, *Synthese*, 30.3/4 (Apr. – May, 1975), pp. 269-270.

This is an important benefit to the supervaluationist account. Firstly, as Keefe points out with regards to sentences such as ‘anyone taller than a tall person is also tall’ and ‘no one who is tall is also short’, ‘These general statements certainly seem to be true’.⁶⁰ Being able to straightforwardly respect the intuition regarding penumbral truths is a *prima facie* advantage. Fine, in fact, goes even further regarding the importance of penumbral truths:

If language is like a tree, then penumbral connection is the seed from which the tree grows. For it provides an initial repository of truths that are to be retained throughout all growth. Some of the connections are internal. They concern the different borderline cases of a given predicate: if Herbert is to be bald, then so is the man with fewer hairs on his head. But many other of the connections are external. They concern the common borderline cases of different predicates: if the blob is to be red, it is not to be pink; if ceremonies are to be games, then so are rituals; if sociology is to be a science, then so is psychology. Thus penumbral connection results in a web that stretches across the whole of language.⁶¹

If Fine is right, then the failure of other theories to adequately explain the necessary truth values of penumbral connections is a serious flaw. However, supervaluationism is not without problems of its own, as we shall now see. Indeed, the first is related to the notion of penumbral connections themselves.

1.1 Supervaluationism makes true the statement ‘there is a sharp cut-off’

Consider the statement, regarding our 5’11” man, ‘Bob is either tall or not tall’. Generally, most people would want this statement to be given the value of true; after all, it is an instance of the classical theorem of excluded middle. The intuition would be that there are only two applicable predications available here, tall or not-tall, and so even though Bob is a borderline case which we are unsure about, we can be sure that he is one or the other. Supervaluation can in fact respect this intuition, and validate this instance of excluded middle as true. The statement is true on every admissible precisification, as on every precisification Bob is either predicated as tall or not-tall (and only one of these). Therefore ‘Bob is either tall or not tall’ is true. Indeed,

⁶⁰ Keefe, *Theories of Vagueness*, p. 162.

⁶¹ Fine, ‘Vagueness, Truth and Logic’, pp. 275-6.

supervaluationism validates all other theorems of classical logic, as Keefe notes: ‘if a sentence is logically true according to classical logic, then it will be true on every complete specification (since they each obey classical logic) and thus, by supervaluational principles, it will count as true simpliciter’.⁶²

Now, it might be pointed out here that maintaining all classical theorems is not a desirable trait in a theory of vagueness; it is traditional, classical, reasoning which, when faced with instances of vagueness, has led to paradox. The idea is that it is precisely these sorts of cases where classical theorems might lead us astray; ‘Bob is either tall or not tall’ appears to demand a level of specificity regarding Bob’s classification which is exactly the matter under question. Bob is a borderline case, where we do not know whether to say that Bob is tall, not-tall, both, or neither. Asserting that ‘x is tall or not tall’ is true for every man x will seemingly lead us to the conclusion that everyone fits neatly into one category or the other; the conclusion we want to avoid in cases of vague predicates. Keefe has an intriguing response to this objection:

A statement which is similar to the trivial ‘either *a* is red or not’ but which would be informative (because sometimes false) is ‘either *a* is definitely red or definitely not-red’, and it could be that the two are sometimes confused: if someone were to assert the former, then, on the assumption that they are obeying the Gricean rule of being informative, it would be reasonable to take them to mean the latter. And the fact that the former is *never* informative could explain why it is so common for our judgements of both sentences to be dictated by our judgements of ‘either *a* is definitely red or definitely not-red’ and why we thus consider ‘either *a* is red or not’ to be not true.⁶³

This argument has merit. In everyday language, the ‘definitely’ operator is usually hidden; even in a case where various borderline shades are being considered, and then a paradigm shade is offered, the normal response is something like, ‘now *this* is red’. So it is understandable that in the case of ‘Bob is either tall or not-tall’, we might interpret that as having the same assertive force as an everyday ascription of being (definitely) tall. This leads to the idea that the statement is saying that Bob is either assertibly, definitely, tall, or assertibly, definitely, not-tall. Of course, this interpretation is not true, and is not given the value of truth by supervaluationism.

⁶² Keefe, *Theories of Vagueness*, p. 164.

⁶³ *Ibid.*

One issue with this explanation, though, is that it is not clear that analogous Gricean-adjacent interpretations can be developed for instances of other classical theorems which are also problematic.

There is another issue with supervaluationism's treatment of 'Bob is either tall or not tall'. While it gives the right value for the statement, it appears that it only does so at the cost of treating 'or' non-truth-functionally; the truth of this disjunction does not depend on the truth of its disjuncts. Neither 'Bob is tall' nor 'Bob is not tall' are true, since different precisifications will assign different values to them; likewise, neither of them are false. We therefore have a true disjunction containing two disjuncts which are neither true nor false. Now consider the (somewhat unusual) statement 'Bob is either tall or tall'. As we have seen, 'Bob is tall' is neither true nor false, giving us a disjunction containing two disjuncts which are neither true nor false. However, this disjunction is itself neither true nor false; there will be some precisifications on which it is true, and some on which it is false. Thus, according to the supervaluationist system, we cannot deduce the truth-value of a disjunction from merely the truth-values of its disjuncts.

A similar issue arises for existential statements such as 'there is an n such that any man with a height of n millimetres is not tall, and any man with a height of at least $n+1$ millimetres is tall'. This statement comes out as true according to supervaluationism, since every admissible precisification draws such a sharp cut-off at some point in the height spectrum. However, there is no n such that it is true *of that* n that any man of height n is not tall and any man of height $n+1$ is tall; each precisification has a different n which fulfils the role, and so no n makes this true on every precisification. Thus we have a true existential statement, with no true instance.

Fine calls what is going on in both these sorts of cases a 'truth-value shift'; 'A v -A holds in virtue of a truth that shifts from disjunct to disjunct for different complete specifications'.⁶⁴ He does not seem particularly concerned about it, for I believe the same reason that Lewis accepted such shifts. After contemplating the

⁶⁴ Fine, 'Vagueness, Truth, and Logic', p. 286.

objection that it is ‘peculiar’ that supervaluationism makes true the existence of a cut-off, but doesn’t make any particular cut-off true, Lewis replies; ‘So it is. But once you know the reason why, you can learn to accept it’.⁶⁵ Once we are in the realm of vagueness, some part of our previous thinking will have to change, since we have landed ourselves in various paradoxical predicaments; if that thing is the existence of uninstantiated true existentials, or true disjunctions with indeterminate disjuncts, then it is a small price to pay. Keefe, along similar lines, advocates for ‘the indirect argument that we should accept the phenomenon [of truth-value shifts] because of its role in an altogether successful theory of vagueness’.⁶⁶

Even accepting that truth-value shifts are an acceptable compromise, though, there is another problem with the existential statement. Supervaluationism makes ‘there is an n such that any man with a height of n millimetres is not tall, and any man with a height of at least $n+1$ millimetres is tall’ true (even though it has no true instance); but surely, in cases of vagueness, we do not want that statement to be true. After all, it is stating that there is a sharp cut-off between tall and not-tall men, which is precisely what we want to avoid (epistemicists might disagree; that will be discussed in the following section). Similarly, the supervaluationist is committed to the falsity of the corresponding universal sentence, ‘for every n , it is not the case that a man of n millimetres is tall while a man of $n-1$ millimetres is not tall’. This statement is a denial of any sharp boundary, and, as we saw in Chapter 1, is a guiding intuition in the area of vagueness; it is also false in every admissible precisification, and so is false. How can the supervaluationist explain this violation of intuitions?

Keefe admits that this commitment is ‘one of the least appealing aspects of the theory’ for her.⁶⁷ But she attempts to mitigate the depth of the violation, by drawing on the distinction between the truth of the general, quantified statements, and the truth of specific instances of those statements. As we have seen, while supervaluationism is committed to the truth of the former, it denies the latter. Keefe points out that ‘typical

⁶⁵ Lewis, ‘Many, But Almost One’, in *Papers in metaphysics and epistemology* (Cambridge: Cambridge University Press, 1999), p. 173.

⁶⁶ Keefe, *Theories of Vagueness*, p. 182.

⁶⁷ *Ibid.*, p. 183.

arguments for tolerance... are in fact arguments against the possibility of a true instance of Fx_i & $\neg Fx_{i+1}$.⁶⁸ Supervaluationism is able to accept all arguments against true instances, as in accordance with its valuations.

It might still feel counter-intuitive to deny all true instances, while avowing the truth of the universal statement. Keefe argues that this uneasiness is mistaken, but has an explanation; 'Our belief that there is no true instance of the quantification gets confused with a belief that the quantified statement is not true'.⁶⁹ The key intuition in cases of vagueness, according to this view, is that there are no true instances of a cut-off. We mistake that for a belief in the generalised statement because we are used to reasoning in non-vague environments, where such an inference would be valid. Keefe compares it to cases of moral duty, where it may be true that someone ought to do X, while being false that there is someone of whom it is true that they ought to do X. This feels somewhat underwhelming; if the explanation requires suspending quantification rules in vague contexts, is this worthwhile? However, as Keefe notes, 'admitting *some* counter-intuitive consequence in connection with the sorites paradox is a disadvantage which will plague *any* theory of vagueness', so I will leave the objection here for the reader to judge its power.⁷⁰

1.2 Supervaluationism cannot handle all instances of vagueness

Another potential problem for supervaluationism is the fact that it does not appear able to easily cover all the proposed instances of vagueness. Take Mt. Everest, for example. It seems as if it is a vague object in some sense at least, since there are rocks, land, etc. which are borderline cases of belonging to Everest. According to a supervaluationist, this is because there is indeterminacy about the meaning of 'Mt. Everest' in some way. The mountain itself, and its boundaries, are fully precise and determinate. The question can then be raised, why is there the indeterminacy in language? It seems odd that there

⁶⁸ Keefe, *Theories of Vagueness*, pp. 184-185.

⁶⁹ *Ibid.*, p. 185.

⁷⁰ *Ibid.*, p. 183.

is such a disconnect between language and the thing it seeks to represent (and that this disconnect is so widespread as vagueness is).

Of course, the supervenientist could respond that the indeterminacy in language is a result of vagueness in our concepts. The above problem is obviated, because there is no explanation needed as to why vague language arose directly from a precise world. Instead, there is the straightforward explanation that a vague language derived from, and represents, vague concepts. This reply is interesting, because it can be partly accepted by even the believer in metaphysical vagueness themselves, while also failing to address the underlying objection.

On the one hand, a proponent of metaphysical vagueness could (and, I will argue later, should) accept that a large number of our concepts are vague, and the language we use inherits this vagueness. There may be disagreements about what exactly this vagueness consists in, but that is unimportant for current purposes. So far, the supervenientist and believer in metaphysical vagueness are in agreement. The difference is in why they take our concepts to be vague, and it is this which explains why the supervenientist reply is somewhat lacking. For the metaphysician, an explanation is easy to construct. If you believe that the world itself is vague, then it stands to reason that concepts meant to represent that world will contain some sort of vagueness too.⁷¹ We form concepts in order to represent, form judgements about, and guide behaviour within, the world. If the world is vague, precise concepts will not be able to perform these roles as well as they could if they were vague themselves. The supervenientist has no such obvious account available. According to them, the world is precise; why, then, are our concepts, which aim to represent the world accurately, vague? This, of course, is not an insurmountable problem; it is, though, an awkward consequence of supervenientism which requires extra explanation.

Furthermore, supervenientism represents our vague expressions as being vague simply because there is no decision between which candidate precise expression is actually meant. This just seems implausible, given how our language actually seems to

⁷¹ The intricacies of the relationship between the vagueness of the world and the vagueness of our concepts will vary, depending on your treatment of metaphysical vagueness (and concept-building), as I discuss in Ch. 5.

work. It seems to be essential to vague expressions that they exhibit some form of indeterminacy; according to supervenience, this could be eradicated simply by deciding which precise meaning to plump for.

To take an example; according to the supervenience, we could eliminate the vagueness of the term 'tall' by stipulating that any man over the height of 5'11" is tall, whereas men who are 5'11" or under are not tall.⁷² This would indeed remove the possibility of there being borderline cases of 'tall'. However, it would remove some expressive power from the term, rendering it less useful. For instance, in many cases you would not be sure of someone's exact height, but sure that they are unusually tall. Of course, this may well be because you are also not sure of what the exact comparison class is, and so the appropriate classification for 'tall'; this is just another form of vagueness inherent in some of our terms, vagueness regarding the salient comparison class. I believe that were this to be the case, there would be a new term coined, which we may call 'tall-ish', to be used in a vague manner.

This is not a counter-example to the supervenience picture. After all, they only say that for every vague term, it is contingent whether that term is vague. The fact that, were the term to be made precise, a new vague term would arise does not invalidate this. The new term ('tall-ish') would be contingently vague as well, and susceptible to precisification.

The point is that we have a need to communicate using vague terms. The supervenience is open to saying that vague terms are only contingently vague, but it seems odd to then admit that if they were precise, other (contingently) vague terms would come into existence. If language were fully precisified, it would not have the expressive power we require. This may well be because the language has to represent concepts which are themselves vague; a precise language will not be up to this task. However, as we have seen, positing vagueness in concepts does not preclude the existence of metaphysical vagueness. Perhaps an easier answer is that there is an

⁷² Of course, in order to eliminate all the vagueness of 'tall', we would have to stipulate tallness conditions for adults, children, basketball players, etc.

inherent indeterminacy in language, which occurs since language is describing an indeterminate world (via indeterminate concepts).

Again, this is not meant to definitively incapacitate the supervenientist. However, it places the burden of argument on them; they need to provide an explanation of why their account entails either the possibility of a fully precise language, when such a language would not seem up to the needs of its users, or that vague terms are only contingently vague, but every usable language necessarily contains vague terms.

1.3 Problems with assertion in cases of vagueness

There is a related problem with regards to supervenientism departing from common sense ideas of meaning, raised by Wright and Shapiro. In paradigm instances of borderline cases, it is the case that users of language are struggling to decide which way to come down on the issue. When faced with a paint catalogue and asked to categorise a certain shade intermediate between paradigm red and paradigm orange, the home-decorator will try to think which paradigm it is closest to, perhaps by comparing. They will then reach a conclusion, albeit perhaps tentatively. Their partner may well disagree with them. As Wright says, ‘the indeterminacy will be initially manifest not in (relatively confident) verdicts of indeterminacy but in (hesitant) differences of opinion... about a polar verdict, which we have no idea how to settle’.⁷³ Shapiro takes it as a datum that ‘in at least some situations, a speaker is free to assert Pa and free to assert $\neg Pa$, without offending against the meanings of the terms’; he calls this the *open-texture* thesis.⁷⁴ According to supervenientism, however, they are offending against the meanings of the terms. While both the concepts of red and orange could admissibly be precisified to include the shade, neither of them definitively do in our actual language. Therefore, assuming that a language-user should attempt to make true statements, in this case the user is wrong to plump for one or the other. Of course, they have not said something false, according to the supervenientist; but they have not said something true either. The supervenientist could well opt to move to an account of assertion where we

⁷³ Crispin Wright, ‘On Being in a Quandary’, *Mind*, 110.437 (Jan., 2001), p. 70.

⁷⁴ Shapiro, *Vagueness in Context*, p. 10.

merely aim to avoid falsity; then the user will have done nothing objectionable, since ascribing either red or orange is not definitely false. However, this is a modification of the general assumption regarding assertion which should be acknowledged, and may carry other negative consequences.

1.4 Problems with the ontology of supervenience

Another problem for supervenience is that it has the consequence of implying an ontology which offends against common sense (a charge which is often levelled at proponents of metaphysical vagueness). A supervenience account of Mt. Everest, for example, will say that there are multitude of ways in which various rocks, areas of land, plants, etc., are arranged, each differing in at least one detail from all others. Each of these ways is an acceptable precisification of our vague term 'Mt. Everest'; that is, each of these arrangements is a candidate for being the precise object referred to by 'Mt. Everest'. In actual fact, our language has not made a decision about which arrangement is correct. This leaves us with an extremely large number of arrangements of matter in the vicinity of where we ordinarily take Everest to be. Initially, you might think that supervenience is thus positing the existence of millions of mountains where previously we thought there was only one, and arguing that our term 'Mt. Everest' declines to pick out one in particular.⁷⁵ This is extremely implausible.

However, this is not how supervenience generally flesh out the account. More usually, these candidate Everests are treated purely nominally; they are not themselves mountains, existing in the world. They are possible arrangements of various bits of matter into a non-natural kind (Mt. Everest). Each precisification of the vague term 'Mt. Everest' picks out a different arrangement; however, on each precisification, only one arrangement is picked out as actually existing. Thus, it is determinately true that there is only one mountain actually existing in the vicinity of Everest, since this is true on every precisification. It appears that the supervenience has circumvented the problem. But

⁷⁵ This is a variant on the 'Problem of the Many', introduced by Peter Unger, 'The Problem of the Many', *Midwest Studies in Philosophy*, 5.1 (Sept., 1980), pp. 411–67, and P. T. Geach, *Reference and Generality* (Ithaca: Cornell University Press, 1980), and famously discussed by many including Lewis in 'Many, But Almost One'.

have they really? This is simply an example of the penumbral truths objection discussed previously, where the supervenientist is committed to a true existential statement with no true instance. In particular, it is the case that there is no particular candidate Everest-collection which can be said to determinately exist. We merely have numerous collections of matter, treated nominally, and for each of them it is indeterminate whether they compose a mountain. The supervenientist can claim that one and only one mountain does exist in the vicinity, but this might seem like something of a technical, hollow victory. The proponent of metaphysical vagueness, on the other hand, can easily and simply say that there is exactly one mountain in the vicinity of Everest, and that mountain is a (metaphysically) vague object.

Before the supervenientist can use the argument that metaphysical vagueness is implausible and unnecessary, an explanation needs to be given of why these results are not problematic. At this point, some may respond that the ‘technical victory’ I (somewhat disparagingly) referred to above is all that is required. After all, common sense intuitions are often unreliable when we get into matters such as fundamental ontology, and as such are not worth ruining a perfectly coherent theory for. This point is taken, but only as far as it goes. For one, supervenientists, and critics of metaphysical vagueness generally, often rely on intuitions themselves when they assert that vagueness in the world is simply implausible. In this specific case, though, the argument is not that common sense views preclude supervenientism due to its consequences in the case of Everest. It is that, given the other problems for supervenientism, it does not hold a decisive theoretical advantage; it provides no economy of explanation. Thus if metaphysical vagueness is plausible theoretically, as I will argue it is, it can only be a point in its favour if it coheres better with widely held beliefs.

1.5 Failures of reference and correspondence truth

Relatedly, the widespread nature of semantic vagueness will, according to the supervenientist, mean that almost all instances of supposed reference are failures (or at least misguided). McGee and McLaughlin spell out this problem in detail, by drawing

a distinction between the disquotational notion of reference and the correspondence notion of reference. Disquotational reference can be ‘implicitly defined by sentences of the following forms’:

‘Kilimanjaro’ refers, in the language we actually now speak as we actually now use it, to Kilimanjaro, provided it exists, and to nothing else.⁷⁶

The correspondence notion of reference is the further notion ‘according to which, if a singular term refers, the activities of the speakers of language, together with various causal connections between the speakers and their environment, pick out a unique individual as the referent of the term’.⁷⁷ Supervenience can handle disquotational reference; the disquotational reference schema is fulfilled on every admissible specification since on every admissible precisification there is a unique Kilimanjaro which exists and is referred to by that term. However, correspondence reference is not achieved, as there is no unique individual which is picked out on every precisification. The causal connections to the world cannot even get started, since there is no specific part of the world which the term determinately refers to. Correspondence reference is an important relation which many would like to retain; supervenience implies that whenever a vague term is used, it does not occur. Since vagueness is widespread in natural language, this is a significant price to pay. The supervenience theorist is led to accepting the conclusion that when speakers refer to objects which may have vague boundaries (plausibly almost every macro-sized object will satisfy this condition), there is actually no specific, unique, object which they are referring to. Such terms become, in at least some sense, empty and circular, with no tangible link to the world. When a supervenience theorist is asked what they are referring to when they use the name ‘Kilimanjaro’, their only available response is ‘the thing that we call Kilimanjaro’, rather than being able to point at an object.

This seems to be at least a *prima facie* disadvantage. The proponent of metaphysical vagueness has no such problem with correspondence reference. In general terms, though there may be some wrinkles depending on exactly which account of metaphysical vagueness is being followed, the picture is as follows. When users of a

⁷⁶ Vann McGee and Brian P. McLaughlin, ‘The Lessons of the Many’, *Philosophical Topics*, 28.1 (Spring 2000), p. 139.

⁷⁷ *Ibid.*

language utter a term such as ‘Kilimanjaro’, this action, in conjunction with similar actions by other users, and various connections between the users and the world around them, means that the term succeeds in referring to a unique (and vague) object. The term ‘Kilimanjaro’ (as used in that language community in that world) corresponds to a unique, actually existing, object – the mountain Kilimanjaro. The fact that the mountain may be a vague object in whatever way the theory in question cashes this out (by having certain parts vaguely, perhaps, or by having vague boundaries) does not cause an issue. Correspondence reference does not depend on objects being precise, but it does depend on there being a unique object which is able to, via the network of connections, be referred to by the term, and therefore correspond to it.

Why is correspondence truth something worth having? That is, what advantages does a theory gain by being able to explain and retain correspondence reference, other than the nice feeling of satisfaction and common sense?⁷⁸ Well for one, correspondence reference allows us to explain what has changed in the counterfactual situation where language (or geography, or human migration patterns) developed differently and ‘Kilimanjaro’ was used to name a different mountain to the one it does now.

For the sake of simplicity, let’s say that in this scenario that other aspects of English are unchanged, and the only difference from our actual practices is that ‘Kilimanjaro’ is used to refer to a different (perhaps nearby) mountain to the actual Kilimanjaro (and only used to do so – it is not also used to refer to our Kilimanjaro). How does the supervenient, relying only on disquotational reference, explain what is different between this scenario and our actual language use? In both the actual world and the counterfactual world, the disquotational sentence above (‘Kilimanjaro’ refers, in the language we actually now speak as we actually now use it, to Kilimanjaro, provided it exists, and to nothing else) is satisfied – in both worlds, (on every admissible precisification) there is a unique Kilimanjaro which is being referred to. Of course, they are different mountains in the different worlds, but this cannot be appealed to – as the supervenient would not want to say that when people in our world use ‘Kilimanjaro’ they are talking past each other due to the fact that there might be different boundaries

⁷⁸ ‘When I use a naming term I refer to a specific object in the world’ is a good approximation of how non-philosophers would describe that part of language use, I assume.

of Kilimanjaro on different precisifications. There does not seem to be a way for disquotational reference to distinguish between the two situations.

Correspondence reference has a ready-made explanation for the situation; in our world, the term 'Kilimanjaro' corresponds with the actual mountain Kilimanjaro, and that is what the term refers to. In the counterfactual situation, however, the term refers to a different mountain; the practice of the language users and the web of causal connections have made it correspond to a different, unique object. I think the ability to give a clear explanation of what is different, reference-wise, in these two situations is a clear bonus for correspondence reference over disquotational reference. Unfortunately, it is one that supervenientists cannot make use of, given their need to hedge their bets with regards to reference, so to speak, due to the fact that in each precisification a slightly different object is picked out by 'Kilimanjaro'.

1.6 Higher order vagueness

There is also a slightly more technical issue; supervenientism appears to suffer from a problem with higher order vagueness. Supervenientism was created as a theory to deal with the implausibility of sharp cut-offs when dealing with vague expressions; however, it ends up creating its own sharp cut-offs. This can be seen when we note that within a specification, everything is either true or false. Equally, everything is either true or false in all specifications (and thus determinate in its truth value), or not. What has been created is a sharp delineation between statements which are determinately true (true in every admissible precisification), those which are indeterminate in truth value, and those which are determinately false (false in every admissible precisification). The sharp cut-off between truth and falsity has been discarded, only to be replaced by two brand new sharp cut-offs.

Now the supervenientist can try to get around this by having it be uncertain whether something is an admissible specification. If it is indeterminate as to whether some precisification is admissible or not, then it could be indeterminate as to whether some statement is true on every admissible precisification or not. That is, it is indeterminate as to whether it is determinate. This is the route that Keefe takes; the notion of an

admissible specification 'is vague...There could thus be a borderline case admissible specification'.⁷⁹ Cases where a sentence p is false in such a borderline admissible precisification, but true in all definitely admissible precisifications, will be a borderline case of 'supertrue'; the supervaluational schema will not determine it to be supertrue, but will also not determine it to lack a truth value, leaving its truth-value unsettled.

We have got rid of the previously problematic cut-offs between truth, falsity, and indeterminacy; unfortunately we have simply replaced them with new cut-offs between (super)truth, (super)falsity, neither (super)true nor (super>false, and unsettled. We could introduce yet more cut-offs, but as Sainsbury memorably put it, 'you do not improve upon a bad idea by iterating it'.⁸⁰

Even if the iteration is not seen as problematic in itself, it raises an issue for the idea of supertruth. Given this modification of supervaluationism, maintaining that (determinate) truth is truth on all specifications 'no longer seems viable... for there is no longer a unique set of "all specifications"'.⁸¹ Supervaluationism loses all (useful) determinate truth. Keefe argues that this is too hasty a conclusion; this 'vagueness... does not, however, eliminate definite truths, since there will be clear cases of sentences true on all admissible specifications, which will be definitely true'.⁸²

Asher, Dever, and Pappas offer a modification of supervaluationism, as usually construed, which they believe can alleviate the problem of higher-order vagueness. Interestingly, they take it to be a positive of their theory that it removes supertruth from the pedestal it is often placed on. Their underlying idea is that 'supervaluation theory loses sight of the fact that supertruth is a modal notion, akin to actual truth'; supertruth, therefore, 'represents a way of being true, rather than a fundamental notion of truth'.⁸³ They thus argue that we should build the structure of supervaluationism with equal regard to truth at a specification, taking that as basic, much as genuine modal theories take truth at a world to be the basic notion which modal truths derive from. This allows

⁷⁹ Keefe, *Theories of Vagueness*, p. 203.

⁸⁰ R. M. Sainsbury, 'Concepts Without Boundaries', in *Vagueness: A Reader* eds. Rosanna Keefe and Peter Smith (Cambridge: MIT Press, 1997), p. 255.

⁸¹ Keefe and Smith, 'Theories of Vagueness', p. 35.

⁸² Keefe, *Theories of Vagueness*, p. 203.

⁸³ Nicholas Asher, Josh Dever, and Chris Pappas, 'Supervaluations Debugged', *Mind*, 118.472 (Oct., 2009), p. 916.

them to construct what they call ‘a genuinely structured notion of higher-order vagueness’ within a supervaluationist structure; their system can admit of ‘infinitely many non-redundant levels of higher-order vagueness’.⁸⁴

But this modification fails as a response to the objection above for two reasons. Firstly, they take the problem of higher-order vagueness to be that according to supervaluationism, there can be no levels of vagueness; according to various logical principles, the hierarchy collapses such that either every member of the series satisfies the predicate, or none does. Considerations of higher-order vagueness force even the original distinction between definite truth, falsity, and indeterminacy to be eroded. Their solution is to show that we can retain a hierarchy of levels. They have shown that supervaluationism can indeed retain many cut-offs, between various levels of indeterminacy; these cut-offs, though, are what we are currently trying to avoid.

Secondly, in downgrading the status of supertruth, and taking truth at a precisification to be basic, they bring to the fore one of the core problems with supervaluationism. Supervaluationism was brought in to explain the widespread phenomena of indeterminacy. Its general approach is that this indeterminacy does exist in our language and concepts as we use them; however, hypothetically this could be eradicated, were we to undergo a process of precisification. These precisifications, though, are really only plausible as hypothetical models. In actual language, many statements containing vague terms are indeterminate; if they weren’t, these philosophical theories wouldn’t be needed. Thus the level we should be most often working on is the one on which indeterminacy occurs; that of supertruth.

The Asher, Dever, and Pappas modification goes against this underlying idea by highlighting truth at a precisification, which contains no indeterminacy. The whole root of this problem is that supervaluationism doesn’t seem to have a fundamental place in its theory for vagueness; vagueness always supervenes on precise things, which cannot easily model vagueness in all aspects. It is akin to a carpet slightly too large for the room

⁸⁴ Ibid, p. 923: there is a detailed technical discussion of how exactly this is structured on pp. 917-923, but it is not particularly relevant for my discussion here.

it is in: whichever area you smooth down, the bunching will always appear somewhere else. A fundamental metaphysical vagueness might provide a simpler account.

We have seen, then, that supervaluationism is a well fleshed-out, robust theory of vagueness. It provides a way to explain what is going on in these apparently paradoxical cases, taking seriously the notion of a borderline case. By introducing the somewhat theoretical notion of an admissible precisification, and quantifying over the multiplicity of these, it respects the unsettled nature of the borderline region while maintaining the theorems of classical logic. However, it does have some drawbacks. There are places where it abuts against general intuition, and some technical issues which somewhat belie the smooth outward appearance of the theory.

These are not knockdown arguments, and I hope I have not presented them as such. Vagueness is a large and confusing area; it is almost certainly impossible to construct an account of it which explains all the phenomena while respecting every other intuition or commitment one might have. Every putative theory, then, must set out its stall, admit its shortcomings, and be open to a cost/benefit analysis. Proponents of supervaluationism believe the costs outlined above are a price worth paying for the elegant solution on offer. I disagree; in particular, I do not think that supervaluationism offers such a clear and convincing account of vagueness that we are licensed to answer this thesis' research question in the negative, which we would be if we were in possession of a theory which can explain all the apparent instances of metaphysical vagueness. Instead, we should continue our search.

2. Epistemicism

Epistemicism offers a clear answer to the problem of vagueness, which is to deny that there is any vagueness (in a sense). While others take the lesson of the sorites paradox to be that there is no specific point at which an attribution of a vague term moves from true to false, and seek to explain how this is the case, epistemicists reply that there is in fact a sharp cut-off. There is a time at which the addition of one hair will turn a previously bald man into a non-bald man. This part of their account needs no more explanation. They do, however, have some work to do. For instance, there is the question of why the boundaries occur at the point at which they do. In addition, an account is needed of why we are ignorant of the location of this boundary. That is, they need to explain the special type of ignorance which occurs in cases of vagueness. Furthermore, it would be favourable to have an account of why we seem inclined to deny the existence of such boundaries when questioned.

Epistemicists do have answers to these questions. The answers, though, reveal a hidden complexity to the account, which on the face of it has an appealing simplicity. In this section I will mainly be discussing the work of Williamson, as he has produced the most sustained and wide-ranging defence of epistemicism. At certain points, however, I will refer to other epistemicist accounts, when Williamson's replies are unappealing.

At its most basic, the epistemicist theory is just that we are ignorant in cases of vagueness, even though there are objective, precise facts of the matter. That is, the phenomenon of vagueness is an epistemological one; the world itself, and our semantics, are perfectly precise, we just do not have the epistemic capability to ascertain all the facts thus determined. Williamson motivates this idea of ignorance clearly:

No one knows whether I am thin. I am not clearly thin; I am not clearly not thin. The word 'thin' is too vague to enable an utterance of 'TW is thin' to be recognized as true or false, however accurately my waist is measured and the result compared with vital statistics with the rest of the population. I am a borderline case for 'thin'... Suppose that an utterance of 'TW is thin' is either true or false. Then since we do not know that TW is thin and do not know that TW is not thin, we are ignorant of something. Either 'TW is thin' expresses an unknown truth, or 'TW is not thin' does. We do not even have an idea how to find out whether TW is thin, given my actual measurements and

those of the rest of the population. Arguably, we cannot know in the circumstances that TW is thin or that TW is not thin; in that sense, we are necessarily ignorant of something.⁸⁵

One important thing to note here is Williamson's assumption that 'an utterance of "TW is thin" is either true or false'; for Williamson, bivalence is an important principle to uphold. 'If one abandons bivalence for vague utterances, one pays a high price'; he argues that denying bivalence necessarily leads to a denial of classical logic, which is problematic since 'classical semantics and logic are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains'.⁸⁶ For the epistemicist, there is a fact of the matter as to how many millimetres high a man must be before he counts as 'tall', and every statement regarding such questions is determinately true or false.

This is where the epistemicist earns the famous incredulous stare. It is the combination of deep, perhaps necessary, ignorance with the postulation of determinate facts nevertheless. If Williamson is not clearly thin, and not clearly not thin, and we cannot know whether he is thin or not, what basis is there to claim that there is a fact out there which we are ignorant of? Perhaps the question is best phrased in a slightly different manner. In other, everyday cases of ignorance, we can generally understand what is preventing us from gaining knowledge, and what is (at least in principle) required for us to break through that barrier. In cases of vagueness, the same is not true; although there is supposedly a fact out there, we have no idea how to ascertain that fact. Why is this the case? As Williamson himself acknowledges,

Although we cannot know whether the term applies in a borderline case, we can know whether it applies in many cases that are not borderline. The epistemic view may reasonably be expected to explain why the methods successfully used to acquire knowledge in the latter cases fail in the former.⁸⁷

⁸⁵ Williamson, *Vagueness*, p. 185.

⁸⁶ *Ibid*, p. 186.

⁸⁷ *Ibid*, p. 216.

This requires some more detail about Williamson's explanation of our ignorance in cases of vagueness. He begins by introducing a concept he calls inexact knowledge, by way of example:

Vision gives knowledge about the height of a tree, hearing about the loudness of a noise, touch about the temperature of a surface, smell about the age of an egg, taste about the constituents of a drink. Memory gives knowledge about the length of a walk, testimony about the physical characteristics of a criminal. The list could of course be continued indefinitely. In each case, the knowledge is inexact. One sees roughly but not exactly how many books a room contains, for example: it is certainly more than two hundred and less than twenty thousand, but one does not know the exact number. Yet there need be no relevant vagueness in the number. The inexactness was in the knowledge, not in the object about which it was acquired.⁸⁸

Inexact knowledge, judging by the examples given, arises when our powers of discrimination are not up to the task of divining the precise fact under discussion. However, they are not so weak as to leave us entirely in the dark; we can make a judgement that we are certain the answer will be within certain (perhaps generous) boundaries. This qualified ignorance does not entail a lack of precision in the thing we are judging.

It also does not entail that we are forbidden from forming beliefs, making statements, or reasoning about the fact under discussion. But, in order to do so reliably, we must recognise that in cases of inexact knowledge there is what Williamson calls a 'margin for error': 'The belief that a general condition obtains in a particular case has a margin for error if the condition also obtains in all similar cases. The degree and kind of the required similarity depend on the circumstances.'⁸⁹ Thus a margin for error principle is formed to govern these cases: "A' is true in all cases similar to cases in which 'It is known that A' is true'.⁹⁰

Imagine, for example, that there is a magazine on the table in front of me which has ninety-two pages (and I do not know this). By looking at it, without picking it up to examine it further, I form the belief, and claim to know, that it has more than thirty

⁸⁸ Ibid, pp. 216-217.

⁸⁹ Ibid, p. 226.

⁹⁰ Ibid, p. 227.

pages, but fewer than two hundred. This belief will be reliable if it is true that in all similar cases the magazine in front of me has between thirty and two hundred pages. What is a similar case in this instance? Presumably one where there is a magazine in front of me with a different number of pages, but one which looks roughly indiscriminable to me in terms of thickness, such that I would be likely to form the same judgement regarding its number of pages. Presuming that I would be able to tell the difference between a magazine with fifty pages more or fewer than the actual magazine, my belief in this case is indeed reliable.

Suppose, though, that I claim to know that the magazine has exactly ninety-two pages. In a relevantly similar case, the magazine could have had ninety-five pages, and I still would have judged that it had ninety-two. That is, the margin for error principle would be violated; 'It is known that the magazine has ninety-two pages' is claimed to be true, but in a case similar to the current case, 'the magazine has ninety-two pages' is false. Thus the initial claim to knowledge is defective. The general idea here is the well-accepted one that we should not let knowledge be lucky. I may in fact be right about the number of pages in the magazine, but I could very easily have been wrong; thus we do not want to ascribe me knowledge.

The case of vagueness is slightly different. In the case of the magazine, or Williamson's example of judging how many people are in a crowd, the reason we have inexact knowledge is that we are unable to gather all the relevant information; if we were able to pick up the magazine and examine it, or study a high-resolution photograph of the crowd, we would have all we need to move from inexact knowledge to 'exact' knowledge. In cases of vagueness this appears not to be so. I could have exact knowledge of Bob's height, as well as the height of every other man in existence, and still not be in a position to make an exact judgement regarding who counts as tall and who doesn't. To put it in terms of the margin for error principle, there are no similar cases in which I am wrong about the fact of Bob's height, and so my ignorance of whether Bob is tall does not stem from that.

The margin for error which is active here is to do with the meaning of the term 'tall'. For Williamson, meaning supervenes on use; 'Words mean what they do because we use

them as we do'.⁹¹ This is as true for vague terms as it is for non-vague terms, although there is a slight difference in how this operates:

For any difference in meaning, there is a difference in use. The converse does not always hold. The meaning of a word may be stabilized by natural divisions, so that a small difference in use would make no difference in meaning. A slightly increased propensity to mistake fool's gold for gold would not change the meaning or extension of the word 'gold'. But the meaning of a vague word is not stabilized by natural divisions in this way. A slight shift along one axis of measurement in all our dispositions to use 'thin' would slightly shift the meaning and extension of 'thin'. On the epistemic view, the boundary of 'thin' is sharp but unstable.⁹²

Thus it is true that, were the community to have used 'tall' in a slightly different manner, it would have had a slightly different meaning. These are the 'similar' cases that we need to examine to find out if the margin for error principle has been broken. In this similar case, the boundary for tallness would have been five feet and eleven and a half inches, instead of five feet and eleven inches. Bob would not have counted as tall in this similar case, but I would still have believed him to have done, because I am not aware of the minute differences in use (and therefore meaning) of the term 'tall'. Thus I violate the margin for error principle, and I do not have knowledge.

Someone who asserts 'Everyone with physical measurements m is thin' is asserting a necessary truth, but he is still lucky to be speaking the truth. He does not know the truth of what he says. Although he could not have asserted the proposition he actually asserted without speaking truly, he could very easily have asserted a different and necessarily false proposition with the same words. The extension of the word 'thin' could very easily have been slightly different, so that it would have excluded everyone with physical measurements m . What distinguishes vagueness as a source of inexactness is that the margin for error principles to which it gives rise advert to small differences in meaning, not to small differences in the objects under discussion.⁹³

So we have the explanation we were looking for. Epistemicism has a straightforward account of what the precise facts are in cases of vagueness; they are determined by our uses of the terms involved. And once this is given, we can see why it is that we are ignorant in such cases; even a small change in use can

⁹¹ Ibid, p. 205.

⁹² Ibid, p. 231.

⁹³ Ibid., pp. 230-231.

lead to a change in these meanings, changing the boundaries of vague concepts. We are simply not able to reliably track these changes, and so fall foul of the relevant margin for error principles.

Williamson has, I believe, done a good job in warding off the incredulous stare provoked by a bald and basic statement of the epistemicist position. He has situated the ignorance characteristic of vagueness within a larger, plausible category. No-one will deny the existence of inexact knowledge, as Williamson defines that term. And the use of margin for error principles is wide and accepted as a plausible mechanism for explaining how we can deny 'lucky' guesses from counting as knowledge (see the cottage industry surrounding 'Gettier cases'). Williamson's epistemicism, then, cannot be summarily dismissed. But that does not mean there are not some objections remaining.

2.1 How are the sharp boundaries determined?

The first question I wish to look at in detail is the following: what, for the epistemicist, determines the boundaries of vague terms? More attention has generally been paid to explanations of the ignorance aspect of epistemicism; after all, it is indeed puzzling to be told that a sharp boundary exists, but we cannot know where it is. But the location of these boundaries deserves scrutiny. According to the epistemicist, there is a certain point at which a man becomes tall; why is that point 184.5cm, say, instead of 184.6cm? As Williamson said, 'the meaning of a vague word is not stabilized by natural divisions'; the boundaries are not determined by the world, so the epistemicist cannot answer that one boundary is metaphysically privileged over others.

An understandable answer, and the one given by Williamson, is that it is our use of the terms which determine the boundary. The way a community uses 'tall' is determinative of that term's meaning and extension. For instance, the basketball community's use of 'tall' will determine a boundary of, say, 200cm, while the Danny DeVito Impersonator Club will have a very different boundary. Williamson says that to 'say that use

determines meaning is just to say that meaning *supervenes* on use'.⁹⁴ The meaning of a term cannot change unless there is a change in the community's use of that term. Thus the facts about how we use vague terms determine the facts about the boundaries drawn by those terms.

Keefe notes that this explanation seems insufficient in one major way; Williamson has said that use determines meaning, but has given no indication as to how it does this. Williamson's supervenience thesis 'is compatible with almost *any* meaning being paired with a given pattern of use'.⁹⁵ Firstly, there is nothing in the bare fact that meaning supervenes on use which guarantees that the meanings all draw sharp boundaries; as McGee and McLaughlin ask, how is it that 'the thoughts and practices of English speakers are able to determine the meaning of vague terms with such unexpected precision'?⁹⁶ Supervenience equally holds in a case where 'if a sentence is neither true nor false, it will be neither true nor false unless there is a difference in the subvening base'.⁹⁷ Williamson himself recognises the possible objection that 'our use leaves not a line but a smear'.⁹⁸ Secondly, even if we accept that sharp boundaries are drawn by our use, Williamson has given no explanation of why a given use determines the specific boundary it does, as opposed to any other; 'What is it that speakers of English do, say, and think that determines, down to the minutest angle of wing contour, exact boundaries of application of the phrase "looks sort of like a June bug"?'⁹⁹

Williamson's reply to those concerned about this is to point out that no-one is able to give a thorough and convincing account of how use determines meaning, and so the burden of proof is not on the epistemicist to do so. 'The inability of the epistemic view of vagueness to provide a successful recipe [for extracting meaning from use] is an inability it shares with all its rivals'; thus it is not a problem if epistemicism means that

⁹⁴ Ibid., p. 206.

⁹⁵ Keefe, *Theories of Vagueness*, p. 80.

⁹⁶ Vann McGee and Brian P. McLaughlin, 'Review of Vagueness by Timothy Williamson', *Linguistics and Philosophy*, 21.2 (Apr., 1998), p. 222.

⁹⁷ Ibid.

⁹⁸ Williamson, *Vagueness*, p. 206.

⁹⁹ McGee and McLaughlin, 'Review of Vagueness by Timothy Williamson', p. 222.

meaning ‘may supervene on use in an unsurveyably chaotic way’, as he admits it might.¹⁰⁰

This is somewhat underwhelming. Of course, Williamson is right that there are no alternative complete and comprehensive accounts of use determining meaning on offer. However, this does not mean that any proposed interaction between use and meaning is equally plausible, and this illustrates the fact that I think Williamson elides the two slightly different forms of the objection. Williamson’s argument that no-one can give an explanation of exactly how meaning determines use given the incredibly complex interplay of various usages and contexts is convincing, I think, against the charge that there is no motivated reason why use determines the sharp boundary of some vague predicate to lie at one particular place over another, minutely different, place.

However, it does nothing to answer the question of why the way that people use vague terms determines that they have a sharp boundary in the first place. Given the somewhat jumbled and unstable pattern of use exhibited by vague terms, and especially the fact that often language users are hesitant to apply them at all in borderline cases, surely the obvious response is to say that the usage determines a fuzzy region of applicability, rather than an extremely sensitive, ever-changing precise extension. As Keefe puts it, Williamson’s supervenience thesis guarantees that if there are sharp boundaries for vague predicates they will be determined by use (and how they are determined may be impossible to explain), but this ‘does not touch the implausibility of the claim that such lines are drawn’.¹⁰¹ While the lack of a ‘recipe’ for how meaning determines use is not a problem unique to the epistemicist, the lack of any positive motivation for the *kind of meaning* being posited is not necessarily shared by all rival theories.

In addition, consider cases such as Everest or the foetus developing into a person, where there is (most would agree) an actually existing object in the world which is the subject of vague talk. Presumably Williamson would not countenance the existence of

¹⁰⁰ Williamson, *Vagueness*, pp. 208-209.

¹⁰¹ Keefe, *Theories of Vagueness*, p. 80.

vague boundaries in the world, and so would say that these questions do have a precise answer. The precise boundaries in these cases, though, are not forever shifting.

Williamson can continue to argue that our use of such terms as ‘mountain’, ‘Everest’, ‘person’, etc. follow the pattern he has given for tall, determining precise boundaries which render users subject to the margin for error principle. But this has the corollary that these terms are almost always wrong about how they describe the world, given that the boundaries they determine are constantly shifting, while the worldly facts are not. This is not necessarily an issue (language can be wrong about the world), but it is worth noting as a consequence of the epistemicist view; it entails that the boundaries drawn by such terms will only every so often, and arbitrarily, match the boundaries which are actually determined by worldly facts.

Williamson will argue that the explanatory power and simplicity afforded by epistemicism in general is enough to outweigh the issues surrounding meaning; if it was just this issue he may be right, but, as we will see, there are other problems facing epistemicism.

2.2 How can language users know the meanings of vague terms?

For instance, Tye makes the point that Williamson’s account makes it hard to see how we can have knowledge of the meanings of vague terms. As he puts it, how can I be sure of the meaning of a vague term ‘if there is a relevantly similar situation in which the word has a slightly different meaning even though I do not recognize that any shift in meaning has occurred’?¹⁰² Why does that possibility of a slight change in meaning disqualify knowledge of the boundaries of vague terms, but not disqualify knowledge of the meaning of those same vague terms? In both cases, my putative knowledge claim is unchanged, while the facts do not comply, violating the margin for error principle.

Williamson does anticipate this in *Vagueness*, and offers the following reply. We take ourselves to know the identity of our friends and family by looking at their faces. While it may be true that there could be almost exact lookalikes of these people, such that we

¹⁰² Michael Tye, ‘On the Epistemic Theory of Vagueness’, *Philosophical Issues*, 8 (1997), p. 252.

could easily mistake one for the other, this does not disqualify us from recognising our family. Thus, Williamson, reasons, ‘why should our ability to recognise the meaning of utterances in our language be undermined by the mere possibility of indiscriminably different meanings?’¹⁰³ This, of course, is precisely what Williamson is asking us to accept in the case of vague term boundaries; it is the mere possibility of different meanings which denies us knowledge. So why is knowledge regarding meanings different? Williamson offers no further explanation. Sainsbury thinks that it may be possible to thread this needle with ‘some further work on very delicate matters concerning evaluating reliability’, but does not sound optimistic.¹⁰⁴

I think there is, though, a less ambitious answer open to the epistemicist, which Williamson hints at. He posits that one might respond to this issue by accepting that ‘speakers only roughly know what their utterances mean’.¹⁰⁵ In this case knowledge of the meanings of vague terms is a paradigm instance of inexact knowledge. Just as I was justified in knowing that the magazine has between ten and a hundred pages, I can have knowledge that ‘bald’ is true of someone with fewer than five thousand hairs, and false of someone with more than thirty thousand hairs. This knowledge, though inexact, is enough for our purposes. Williamson seems to take this option to be an unnecessary concession, but it in fact places the epistemicist in no worse a position than any other theorist. For someone who takes there to be no sharp boundaries, our knowledge of the meaning of vague terms must take this same form; we know that it applies in determinate cases, and doesn’t apply in determinately false cases, but no more. More generally, this is often part of language knowledge, even disregarding considerations of vagueness. Most of us cannot provide a complete and accurate grammar and semantics, but we still know how to use our language, and so have knowledge. However, one complication for epistemicists in particular is that this seems to reinforce an earlier objection: if speakers only roughly know what their utterances mean when they discuss vague instances, then how plausible is the claim that the jumble of uses they exhibit determines a precise boundary for the predicates under discussion?

¹⁰³ Ibid, p. 236.

¹⁰⁴ R. M. Sainsbury, ‘Vagueness, Ignorance, and Margin for Error’, *The British Journal for the Philosophy of Science*, 46.4 (Dec., 1995), p. 598.

¹⁰⁵ Williamson, *Vagueness*, p. 236.

2.3 The implausibility of extreme sensitivity to meaning

However, even if this is not a large problem for the epistemicist, another is lurking in the vicinity. In Williamson's reply, he mentioned the possibility of 'indiscriminably different meanings'; this should give us pause for thought. As Keefe notes, the combination of this extreme sensitivity of meaning to use, and the use in question being that of the community as a whole, means that 'the pattern of other people's assertions could stop me being able to have a belief with a given content'.¹⁰⁶ Now, one could respond that while this is technically true, the actual content of my belief would be indiscriminably similar to the given content, which is not necessarily an awful consequence. As Schiffer points out though, there is the further consequence that 'ascriptions of the form 'A said (asserted, stated, etc.) that S' are virtually never true'.¹⁰⁷ If any change in use results in a change in meaning, then meaning will be in almost constant flux, and the chances of the meaning at the time of the original statement and the meaning at the time of the report being precisely the same are very small. The fact that Williamson's account renders almost all statements of disquotational belief attribution false is a concrete cost, which should be considered when evaluating the theory as a whole.

Williamson argues for this extreme sensitivity by appealing to the implausibility of what he sees as the alternative. This is that the 'reference of a vague term would be insensitive to most small shifts in use, but would change in large jerks as the pattern of use shifted across critical boundaries'.¹⁰⁸ Williamson argues that this privileges some boundaries of application as natural, which is implausible.

More importantly, this 'jerky' account of meaning undermines Williamson's account of our ignorance in borderline cases. Recall that we violate the margin for error principles because we would have made the same judgement, even though use, and therefore boundary, has slightly changed. If only large changes in use lead to discriminable changes in meaning, this explanation is no longer available. In a similar case where the boundary of the term has changed, I would know this because I would appreciate the

¹⁰⁶ Keefe, *Theories of Vagueness*, p. 70.

¹⁰⁷ Stephen Schiffer, 'The Epistemic Theory of Vagueness', *Philosophical Perspectives*, 13 (1999), p. 494.

¹⁰⁸ Timothy Williamson, 'Reply to Commentators', *Philosophy and Phenomenological Research*, 57.4 (Dec., 1997), p. 948.

difference in use (and meaning), and therefore would not make the same judgement. Therefore Williamson's entire epistemic account relies on this extreme sensitivity of meaning to use. While it is not incumbent on Williamson to insist that every single change in use leads to a change in meaning, it is necessary to hold that small changes in use lead to indiscriminable changes in meaning. It is the indiscriminability of the changes which is essential to the explanation of ignorance, and the very same indiscriminability which seems problematic. It is a cost we have to accept in order to reap the benefits of the overall theory. Williamson himself acknowledges this: 'the only available explanation of our ignorance assumes the smooth hypothesis'.¹⁰⁹ But he argues that the explanatory benefit of his overall account means that this is a price worth paying.

The reader will have noticed that Williamson states that his explanation of the ignorance in vague cases is the only available explanation. Horwich, for one, would disagree. He believes he has an account of ignorance which is compatible with there being a pattern of 'jerky' shifts in meaning. His main contention is that it is a fundamental fact about vague predicates that we experience a paralysis of judgement about whether or not they apply in borderline cases. Their 'basic conceptual-role', to use Horwich's phrase, is such that it is left unspecified what to say in borderline cases.¹¹⁰ This fact determines a pattern of use for vague terms whereby we are highly likely (if not certain) to apply them to determinate cases, then simply unwilling to apply them in borderline cases, before we are highly likely (if not certain) to deny them of cases where the predicate determinately does not apply. As Horwich puts it, 'no one who has a full grasp of "F"'s meaning will confidently apply it to things that are identified as being in the middle range'.¹¹¹ Horwich then says that, given this fact, we can never be justified in coming to a belief about whether a vague predicate applies to a borderline case, because the facts of usage will not support such a belief. Our ignorance, therefore, is explained by our lack of judgement on borderline cases.

¹⁰⁹ Ibid, p. 949.

¹¹⁰ Paul Horwich, 'The Sharpness of Vague Terms', *Philosophical Topics*, 28.1 (Spring 2000), p. 86.

¹¹¹ Ibid, p. 87.

There is an immediate objection to this view. Horwich takes the meaning of a term to be given by its having a 'basic acceptance property', with an acceptance property determining when sentences containing the term will be accepted, and thus determining our use of the term.¹¹² This is consistent with there being 'jerky' shifts in meaning, which occur when the basic acceptance properties of a term change. It is therefore the case that whether or not a vague term applies to a borderline case is determined by its 'acceptance property'; and for Horwich it is in fact determined in every case, since he wants to protect classical logic. Why, then, does a full grasp of a vague term's meaning not include this fact, and thus justify knowledge in borderline cases?

To put this another way, for Horwich the facts about whether or not a vague term applies to a borderline case supervene on facts about that term and its acceptance properties. In order to explain our ignorance in borderline cases Horwich must explain why we cannot have knowledge about this supervenience relation. He does not do this as far as I can tell, and merely presupposes that such knowledge is impossible, in order to state that 'full knowledge' of a vague term is knowledge of its usage patterns. As such, as Williamson states 'he has not produced an alternative non-trivial explanation of our ignorance in borderline cases'.¹¹³

Perhaps the lesson we can draw here is that it is impossible to maintain both an acceptance of 'jerky', discriminable shifts in meaning, and a convincing explanation of our ignorance in borderline cases. Horwich saves the former, while Williamson places more importance on the latter. Of course, this is all on the assumption that the only choice is between indiscriminable changes in meaning, and jerky shifts in meaning, which it may be if you are committed to keeping classical logic. But, as Keefe points out, 'outside the classical framework, we may maintain that there is sometimes indeterminacy about whether a given range of changes in use is sufficient to alter meaning'.¹¹⁴ The dilemma appears unique to epistemicism.

¹¹² Paul Horwich, *Meaning* (Oxford: Oxford University Press, 1998), p. 44.

¹¹³ Williamson, 'Reply to Commentators', p. 950.

¹¹⁴ Keefe, *Theories of Vagueness*, p. 76.

2.4 Why do we believe there are no sharp boundaries?

Setting aside this issue, though, still leaves another potential problem for the epistemicist. Suppose that we have been given a thorough and convincing explanation of why any belief regarding the precise location of a cut-off point in a sorites series cannot attain the status of knowledge. The question remains as to why, in general, we refrain from forming these beliefs in the first place. That is, why do we seem drawn to the idea that there just is no cut-off at all (the fact that earns epistemicists their incredulous stares)?

This can be framed in terms of Schiffer's notion of a happy-face solution to a paradox. A happy-face solution is one which achieves two things: first, it picks out 'the member of the [paradoxical] set that's not true; and second, it shows us why this spurious proposition deceived us'.¹¹⁵ The second part of a happy-face solution is what is sometimes termed an error theory: it is a theory explaining why a widespread error was made in the first place, in order to strengthen the case for supposing that it is indeed an error. In the case of epistemicism, the first objective is achieved by picking out the inductive step of the sorites argument as being false. It is false that 'if a man with n hairs on his head is bald, then a man with $n+1$ hairs on his head is also bald', because at some (unknown) point, having n hairs constitutes baldness, while having $n+1$ does not. However, epistemicism struggles to conjure a solution to the second objective.

It cannot simply be the case that the lack of knowledge concerning the location of the cut-off automatically leads us to assume that no such cut-off exists, as Williamson's initial attempt runs: 'one might naturally suppose that if they [sharp boundaries] exist then we should be able to find them, and so regard our inability as evidence of their non-existence'.¹¹⁶ I will quite happily admit that for every lottery ticket, I cannot know whether it is a winner or not; this does not lead me to accept the generalisation that no ticket is a winner. A further explanation is needed.

Williamson would presumably want to move to a more sophisticated position, whereby something about the way in which we lack knowledge in vague cases causes us to

¹¹⁵ Stephen Schiffer, 'Skepticism and the Vagaries of Justified Belief', *Philosophical Studies*, 119.1-2 (2004), p. 179.

¹¹⁶ Williamson, *Vagueness*, p. 234.

come to the conclusion that we lack the knowledge because there is no fact of the matter for us to know. As the key issue, for Williamson, in vague cases is the violation of the margin for error principle, the explanation as to why we refuse to form beliefs in these cases should also arise from this principle. You could be tempted to argue that it is the fact that people understand the epistemicist theory itself which stops them from forming the relevant belief, but this seems implausible, and even Williamson does not assume that a lay non-philosopher will reason in accordance with epistemicist arguments.¹¹⁷

Whatever margin for error-related justification is proposed, though, we can go back to one of Williamson's own examples which we discussed earlier for a counterexample. He explicitly says that even though there is the possibility that almost-exact lookalikes of our family members may cause us to misidentify them, we do not take this to mean that we cannot recognise our own family. That is, in the family recognition case, the fact that there is a possible margin for error violation does not cause us to refrain from forming the relevant belief (that we can recognise our family members), whereas in the vagueness case it does. What is the principled reason for this?

Maybe it has to do with the fact that in the family case, the relevant cases which would trigger the margin for error violation are not actually extant. We are not, as a matter of fact, constantly running into almost-exact lookalikes of our relatives; if we were, perhaps we would also refrain from believing that we had the ability to recognise them. This seems plausible, but it also seems plausible that this would occur through a process of trial and error; we would sometimes correctly recognise a relative, and oftentimes be mistaken, and eventually we would learn that we have no way of telling one from the other and so learn to not rely on purely visual recognition. This is not how the process goes in cases of vagueness. Pre-theoretically, people do not begin by assuming that they are correctly able to identify the precise location of a sharp cut-off in a sorites series, realise that they are often mistaken, and revise their belief formation system to refrain from making the same mistake in the future. As soon as someone

¹¹⁷ Ibid.

grasps the meaning of vague terms, they realise that it is simply not the sort of situation where it is appropriate to form a belief about the location of a sharp cut-off.

Keefe presses a similar objection. She notes that there are two ways of being ignorant regarding p ; 'either by having a belief that p ... which fails to count as knowledge... or by *not* believing that p '.¹¹⁸ In the case of vague terms, we are more often (almost always?) ignorant in the second sense; we simply don't form a belief that p in cases where p is a borderline case. However, Williamson's account of ignorance is predicated on the first type; it asserts that if we form a belief regarding a borderline case, it will fail to count as knowledge. As Keefe observes, the epistemicist has given an explanation of 'why we *cannot* know rather than why we happen not to know'.¹¹⁹ Any proposal for a happy-face solution which is based upon the mechanism of forming beliefs about vague cases which turn out to be false will necessarily be a failure, as this is not the mechanism which is actually in play in these instances.

Sorensen has provided a separate and interesting proposal for why we may form the related belief that there are no sharp boundaries between something being p and being not- p (related because plausibly one reason why we refrain from believing that p is we believe that there is no sharp boundary where things switch from being p to being not- p). Sorensen recalls John Stuart Mill's statement that the 'The most deeply-rooted [fallacy] ... is that the conditions of a phenomenon must, or at least probably will, resemble the phenomenon itself'.¹²⁰ While Mill was referring to cases such as the respiratory powers of foxes meaning that they might be a source of a cure for asthma, Sorensen argues that the same sort of thinking might underly the intuition that a small change cannot have a substantial effect; the fallacy is thinking that a large effect must have had a similarly large cause. However, Sorensen points out that science can provide us with counter-examples: 'for instance, an arbitrarily small change in the velocity of an object can make the crucial difference between whether it achieves escape velocity and travels far out into space, or fails, and crashes to earth'.¹²¹

¹¹⁸ Keefe, *Theories of Vagueness*, p. 71.

¹¹⁹ *Ibid.*

¹²⁰ John Stuart Mill, *A System of Logic* (Toronto: University of Toronto Press, 1974), p. 765.

¹²¹ Roy Sorensen, 'Vagueness, Measurement, and Blurriness', *Synthese*, 75.1 (Ar., 1988), p. 68.

According to Sorensen, then, our hesitance to ascribe sharp boundaries or adjudicate on borderline cases may be deep-seated but defeasible.

While Sorensen's point is taken, I think there is an important disanalogy between the cases. While it is true that in the escape velocity case, just as in the vagueness case, a tiny change in speed has massive consequences, it is also the case that the reason for this is readily understood. We know that there is a certain speed which needs to be reached to overcome gravitational forces, and that there is a fine line between achieving it and not; moreover, this fine line is backed up by an objective reality. In the vagueness case, though, there does not seem to be any analogous underlying principle which can help understand why in this particular case a small change is the particular small change needed to effect a large change; it appears completely arbitrary (and, indeed, on the epistemicist view it may be determined by effectively arbitrary changes in language use). We still have not achieved our happy face.

In any case, the epistemicist may fall back on the argument that they do not have a happy-face solution to the paradoxes of vagueness, but neither does anyone else. This is in fact what Schiffer himself proposes; it is his view that 'some classical philosophical paradoxes – for example, the problem of free will and the sorites – don't have happy-face solutions'.¹²² He says the following of free will, which I believe could equally stand for the sorites: 'If the problem of free will had a happy-face solution, I think we would have heard about it by now'.¹²³ This is of course a somewhat informal argument, but it is illustrative of a strategy which is open to the epistemicist. A happy-face solution is a nice-to-have for any philosophical problem, but one may not be available if the paradox under discussion is well and truly paradoxical. We may therefore be reduced to searching for the best 'weak unhappy-face solution', to use Schiffer's terminology. This is a solution which, while not satisfactorily explaining away all the intuitions and errors which lead to the paradox, provides a 'glitch-free version of the concept' or concepts under examination which can replace the original and do the work we expect of it.¹²⁴ In the case of epistemicism, the glitch-free version is saying that vague terms do actually

¹²² Schiffer, 'Skepticism and the Vagaries of Justified Belief', p. 179.

¹²³ Ibid.

¹²⁴ Ibid., p. 181.

have precise but unknowable sharp boundaries (as opposed to the original thought that they do not have sharp boundaries), which can allow us to block the paradox from forming and retain classical reasoning. The literature around vagueness, on this way of thinking, can be seen as a quest to produce the most agreeable weak unhappy-face solution.

2.5 Conclusion

We have now seen that the two most strongly supported accounts of the phenomena of vagueness, supervaluationism and epistemicism, each face various problems. None of these problems are probably fatal by themselves; as mentioned, the paradoxes surrounding vagueness are thorny enough to preclude a clearly convincing solution which can be widely agreed on, and every theory will face particular areas of difficulty. But together they undermine the oft-presupposed idea that we already have a working theory or two on what is going on in cases of vagueness, so we can pick our favourite and there is simply no need to go to the lengths of upending the worldview that the world is precise and there are determinate facts about it. In actual fact, neither offer a strong enough general theory of vagueness to overturn the intuition that the world itself is vague; there is therefore at least a *prima facie* motivation for exploring the possibility of a metaphysical account. And, hopefully, I will show that once we begin this exploration, we realise that there are advantages offered by a metaphysical account which may be enough to persuade some to admit this into their view on vagueness.

Part B

We are now able to move on to the second part of the project, having shown that, due to the deficiencies of the leading competing accounts, there is at least *prima facie* reason to accept the possibility of a positive answer to the research question (is there genuine metaphysical vagueness?). That is, accounts which locate vagueness entirely within the semantic or epistemic realm are not able to offer such complete and convincing explanations of the phenomena that we would be licensed to accept that the intuitions in favour of genuine metaphysical vagueness are mistaken.

In Part B I will first, in Chapter 3, briefly examine some general arguments which are given against the very idea of metaphysical vagueness; that it is incoherent, implausible, or necessarily leads to unpalatable conclusions. I then consider what form genuine metaphysical vagueness should take; that is, what are we accepting when we say that genuine metaphysical vagueness does exist? Chapter 4 will consider previously offered accounts of genuine metaphysical vagueness and conclude that, while each has its merits, none are fully convincing. In Chapter 5 I will present my own account of what genuine metaphysical vagueness consists in.

3. A general defence of metaphysical vagueness

Before specific accounts of metaphysical vagueness are examined, and my own account put forward, though, it is necessary to mount a general defence of metaphysical vagueness. Many dismiss the possibility of genuine metaphysical vagueness offhand, arguing that it is implausible, inexplicable, or even incoherent. These general charges are dismissed. In addition, the burden of proof is switched. Many take it to be incumbent on the proponent of metaphysical vagueness to offer arguments for its existence. This is based on assumptions regarding specificity and idealisations of language and concepts which are unjustified, and goes against general intuitions regarding indeterminacy in the world. It is argued that the background assumption should be that there may be vagueness in the world, unless arguments are supplied to show otherwise. More specific arguments are then discussed, such as the famous Evans/Salmon argument against the existence of vague objects, the Sider argument against the vagueness of the number of objects in the world, and Williamson and Sorensen's insistence on classical logic. It is concluded that none of these arguments force the conclusion that metaphysical vagueness is untenable.

3.1 General arguments against metaphysical vagueness

Metaphysical vagueness is often dismissed as an incoherent notion, something strange that does not so much need to be argued against as pointed to as a cautionary tale of allowing one's thoughts to wander too far from the path of reason. Russell gave perhaps the most famous articulation of this view of metaphysical vagueness:

[I]t is often suggested that as our language becomes more precise, it becomes less adapted to represent the primitive chaos out of which man is supposed to have evolved the cosmos. This seems to me precisely a case of the fallacy of verbalism – the fallacy that consists in mistaking the properties of words for the properties of things. Vagueness and precision alike are characteristics which can only belong to a representation, of which language is an example. They have to do with the relation between a representation and that which it represents. Apart from representation, whether cognitive or

mechanical, there can be no such thing as vagueness or precision; things are what they are, and there is an end of it.¹²⁵

Dummett later took up the cudgel, noting that ‘the notion that things might actually be vague, as well as being vaguely described, is not properly intelligible’.¹²⁶ The nice thing about comments such as these is that they are not arguments, so they do not require counter-arguments. I think the sheer amount of literature on vagueness in the last fifty years, while maybe not all being completely intelligible, is enough to prove Dummett mistaken; and while things may be what they are, that is not to say that they might not be vague.

In fact, I might like to pull a similar strategy to Dummett and Russell in favour of metaphysical vagueness, in order to balance the scales somewhat. I think a great many non-philosophers (and perhaps philosophers too) have the intuition that the world itself is vague, or fuzzy, or indeterminate, or not having a fact of the matter one way or the other, or however it is to be phrased. As Tye puts it, it is ‘part and parcel of our common-sense view that these items [mountains, deserts, and clouds] are not perfectly precise, that they have fuzzy boundaries’.¹²⁷

Historically, the burden of proof has often been seen to be on the proponent of metaphysical vagueness; after all, anyone making an outlandish claim should surely need to go out of their way to provide a convincing argument for accepting it. I would like to remove this burden, and potentially reverse it. Following Tye, I think there is a background intuition in favour of metaphysical vagueness. Of course, that in no way means that metaphysical vagueness definitely exists; but it does mean that any account of the world which rules out metaphysical vagueness would be improved by being able to give an explanation of why that mistaken intuition has arisen, or why the world might seem (pre-theoretically) to be fuzzy.

However, it must be admitted that not all of the antipathy towards metaphysical vagueness is based purely on feelings and intuitions. There have been some notable arguments which have sought to show that metaphysical vagueness gives rise to

¹²⁵ Bertrand Russell, ‘Vagueness’, *Australasian Journal of Psychology and Philosophy*, 1.2 (1923), pp. 84-85.

¹²⁶ Michael Dummett, ‘Wang’s Paradox’, *Synthese*, 30 (1975), p. 314.

¹²⁷ Michael Tye, ‘Fuzzy Realism and the Problem of the Many’, *Philosophical Studies*, 81 (1996), p. 215.

contradictions or unpalatable consequences, and therefore is not an advisable path to go down. I will now look at the two most prominent of these.

3.2 The Evans/Salmon argument from identity

Perhaps the best-known argument against metaphysical vagueness is the Evans/Salmon argument from vague identity. In brief, the argument is this. If there truly is vagueness in the world, then one of the consequences of this is that there might arise vague identity statements (for instance, Everest might be vaguely identical with the collection of matter including all its determinate parts and some of its indeterminate parts). Imagine, then, that we do have two objects, *a* and *b*, which are indeterminately identical due to the vagueness of *a*, or *b*, or both. It could be said of *a*, therefore, that it has the property of being indeterminately identical with *b*. Considering *b*, however, we should say that it is determinately identical with itself; after all, what use is an identity relation if we cannot say that everything is identical with itself? In other words, *b* does not have the property of being indeterminately identical with *b*. By application of Leibniz' law of the indiscernibility of identicals, *a* and *b* cannot be identical because *a* has a property which *b* does not have. But we had thought that *a* and *b* were indeterminately identical, rather than determinately non-identical. Our initial assumption that there may be vagueness in the world has led us to contradiction.¹²⁸

Thomason quickly pointed out a hidden premise in the argument which could weaken its thrust, which is that the move from 'It is indeterminate whether $a = b$ ' to '*a* is such that it is indeterminately identical with *b*' is only valid if *a* is referentially determinate. That is, if the term '*a*' indeterminately picks out one of a number of candidate objects in the vicinity, then this move is not licensed; there may be some candidates for *a* which are determinately identical with *b*, and some which are not. Thomason draws a modal analogy; in the case of non-rigid designators, such as 'the teacher of Alexander', we are not licensed to move from 'it is contingent whether the teacher of Alexander is identical with Aristotle' to 'the teacher of Alexander is such that it is contingent whether they are

¹²⁸ Gareth Evans, 'Can There Be Vague Objects?', *Analysis*, 38.4 (Oct., 1978), p. 208; Nathan Salmon, *Reference and Essence* (Oxford: Blackwell, 1982), p. 308. I will mainly focus on Evans' formulation, and the ensuing literature.

identical with Aristotle'. Assuming that all referring terms are rigid designators would allow us to (fallaciously) conclude that Aristotle is necessarily identical with the teacher of Alexander; Evans' assumption that all terms are referentially determinate has similarly allowed him to conclude that *a* and *b* are necessarily non-identical.¹²⁹

Lewis has no time for this counter-argument, arguing that it completely misunderstands Evans' intentions; he is not offering 'a fallacious proof' of the 'absurd conclusion' that there can be no vague identity statements, but instead is 'making a good argument in favour of a very different conclusion'.¹³⁰ The different conclusion, as indicated by the title of Evans' article, is that there can be no vague objects, rather than no vague identity statements. This is because Evans apprehended the possibility that identity statements could be made vague due to referential indeterminacy, but was also aware that the proponent of vague objects does not have recourse to this evasion of the argument. For them, a name such as 'Everest' 'rigidly denotes a certain vague object', and so Evans' conclusion cannot be avoided in the manner Thomason suggests.¹³¹ According to Lewis, then, what Evans' argument causes real trouble for is 'the vague-objects view combined with the view that vague identity yields identity statements with indeterminate truth value'.¹³²

At this point, one option open to the proponent of vague objects is to in fact accept Evans' argument; that is, to accept that there can be no vague identity but there still can be vague objects, because each vague object is simply distinct from every other vague and precise object, even the precise objects which are candidates for its own composition. Noonan even goes so far as to state that this option is widely accepted, ascribing its explanation to Edgington and Tye:

Everyone knows that Evans's argument against vague identity in-the-world doesn't show that there aren't vague objects (Edgington 2000; Tye 2000). Even if the argument succeeds all it proves is that every vague object is determinately distinct from every precise object and every other vague

¹²⁹ Richmond H. Thomason, 'Identity and Vagueness', *Philosophical Studies*, 42.3 (Nov., 1982), p. 331.

¹³⁰ David Lewis, 'Vague Identity: Evans Misunderstood', *Analysis*, 48.3 (Jun., 1988), p. 128.

¹³¹ *Ibid.*, p. 129.

¹³² *Ibid.*

object. So it is consistent to hold both that there are vague objects and that the identity relation is precise.¹³³

Williams, however, notes that there are powerful reasons for the believer in vague objects, or metaphysical vagueness more generally, to accept that they do imply vagueness in identity; however, he shows that Evans' argument can still be resisted. The first point can be appreciated by considering an object, *a*, about which it is vague whether *a* exists.¹³⁴ Now consider the mereological fusion of all things that exist (call it *b*) and the mereological fusion of all things that exist which are not identical with *a* (call this fusion *c*). It can be seen that *b* and *c* will only be distinct if *a* exists; but ex hypothesi *a* exists vaguely, and thus *b* and *c* are distinct (or identical) only vaguely.¹³⁵ Vague existence (and the existence of vague objects) can indeed therefore lead to vague identity statements, placing us back in Evans' crosshairs.

Williams' next move is to argue that we can again appeal to referential indeterminacy to defuse Evans' argument, but not this time out of misunderstanding (so as not to incur Lewis' wrath), but out of an appreciation that metaphysical vagueness can give rise to referential indeterminacy. The thought is this: to be sure, there are cases where referential indeterminacy arises due to language users not doing what we could to secure a determinate referent. However, there can also be cases where language users have in fact done everything in their power to secure a determinate referent, but the process has still failed due to worldly indeterminacy; the world simply didn't play ball and provide us with a determinate referent to the term 'Everest', for example, because of its inherent vagueness. The above-sketched refutation of Evans' argument is now no longer a sideshow, avoiding the real issue at hand, but an affirmation of metaphysical vagueness; the argument cannot go through *because* there is vagueness in the world (which has infected our language with referential indeterminacy).¹³⁶

¹³³ Harold W. Noonan, 'Are There Vague Objects?', *Analysis*, 64.2 (Apr., 2004), p. 131; Dorothy Edgington, 'Indeterminacy de re', *Philosophical Topics*, 28.1 (2000), pp. 27-44; Michael Tye, 'Vagueness and Reality', *Philosophical Topics*, 28.1 (2000), pp. 195-209.

¹³⁴ To see how this scenario may follow from the existence of vague objects, consider the mereological sum of all the determinate parts of Everest and some of its indeterminate parts; does this sum compose an object? That is, does the sum exist as an object in the world, over and above being an abstract set? Plausibly, for the believer in vague objects, it could be said to vaguely exist.

¹³⁵ J. Robert G. Williams, 'Multiple Actualities and Ontically Vague Identity', *The Philosophical Quarterly*, 58.230 (Jan., 2008), p. 144.

¹³⁶ *Ibid.*, pp. 147-153.

Tye offers another resolution of the argument, available to those who do not wish to be forced to accept referential indeterminacy.¹³⁷ This is to say that the argument may show that identity statements cannot be indefinite in truth value (they are straightforwardly true or false), but the statements remain vague:

To say that an identity statement is vague, on my view, is to say that it has a vague meaning. This will be the case, I maintain, if either of the singular terms flanking the identity sign is vague. But the vagueness of 'a' or 'b' in ' $a = b$ ' does not require that ' $a = b$ ' might be indefinite in truth-value.¹³⁸

This may be seen as more biting the bullet than resolving the issue, however, and it could be argued that it gives too much ground to Evans/Salmon. A final option open to the stronger-stomached proponent of genuine metaphysical vagueness is to accept something that Evans dismisses the possibility of in his formulation; that objects may be indeterminately identical with themselves. This blocks the application of Leibniz' law, derived from the fact that while a may be (*ex hypothesi*) indeterminately identical with b , b must be determinately identical with itself. Perhaps, when we take the idea of a genuinely metaphysically vague object seriously, it can be the case that an object is only indeterminately identical with itself? This does not necessarily imply that the identity relation itself is vague. Consider the property of 'being precisely 150cm tall'; many vague objects may indeterminately satisfy this property, creating indeterminate statements ascribing them the property, without implying that the property itself is anything less than precise. Plausibly, the identity relation could be put in a similar position by vague objects.

The particular intricacies and consequences of the above rebuttals deserve further discussion, but are not necessary to go into here. I have merely aimed to show that the Evans/Salmon argument from identity against vague objects does not deal a knockout blow against metaphysical vagueness.

¹³⁷ I previously argued that the inability to preserve a notion of correspondence truth was a disbenefit for supervaluational theories, and do take it to be a positive for accounts of metaphysical vagueness that they are able to countenance a straightforward link between the term 'Everest' and a single (vague) object in the world.

¹³⁸ Michael Tye, 'Vague Objects', *Mind*, 99.396 (Oct., 1990), p. 556.

3.3 Sider's argument for unrestricted composition

Let us now move on to another supposed argument against the very possibility of vague objects; Sider's argument for unrestricted composition. Sider, in considering the Special Composition Question of whether there could be a coherent, principled answer to the question of when does composition occur and not occur, argues (following Lewis) that any form of restricted composition (i.e. that composition sometimes occurs and sometimes does not) will have to be vague in order to satisfy our everyday intuitions about objects: 'the only plausible restrictions on composition would be vague ones'.¹³⁹ Sider then states that any vague restriction on composition necessarily leads to vagueness as to how many concrete objects exist (which is plausible, see footnote 55, above). If this were the case, we could formulate what Sider calls a 'numerical sentence' asserting the number of concrete object in the world of the form 'there exist exactly n concrete objects in the world'. Given that by hypothesis composition is vague, there will be vagueness for some collection of things whether they compose an object; it will therefore be vague whether this object should count towards the total of n in the above numerical sentence. The numerical sentence will thus come out as vague; but Sider takes this to be an impossibility, because numerical sentences are unambiguous and contain only precise terms.¹⁴⁰

One option of countering Sider's argument is to point out that while for each of the candidate collections of things under consideration it is indeterminate as to whether they compose an object, it may be determinate that there is one object composed in that vicinity. For instance, while individual candidate clusters of matter may only indeterminately compose Everest, it may be determinate that there is only one composite mountain object in the vicinity of these clusters; Everest itself. Thus, when we consider the question of how many concrete objects are in the world in the total, the local indeterminacy regarding the individual clusters disappears, and we simply add one (Everest) to the total, without having to worry about exactly what the boundaries of Everest are.

¹³⁹ Theodore Sider, *Four-Dimensionalism: An Ontology of Persistence and Time* (Oxford: Clarendon, 2001), p. 121.

¹⁴⁰ *Ibid.*, pp. 127-128.

If this does not convince, there is the perhaps more straightforward option of simply denying Sider's assumption that 'numerical sentences' cannot be vague because they only contain logical terms. After all, one of the terms in such sentences will be 'exists' (or some variant of it); for someone who believes that vague existence is a property, there is no issue in holding that this may provide the vagueness of a numerical sentence. Sider's argument looks like the sort of thing which is very convincing to someone pre-disposed against metaphysical vagueness, but has little traction when metaphysical vagueness is genuinely taken seriously; one person's modus ponens is another's modus tollens.

3.4 Williamson and Sorensen's argument for classical logic

The final general argument against metaphysical vagueness is one promoted by the epistemicists Williamson and Sorensen (we shall examine the epistemic theory of vagueness shortly). In short, it is that the existence of metaphysical vagueness, in whatever form, would necessarily lead to the abandonment of classical logic, for classical theorems would fail in the presence of genuine, irreducible metaphysical vagueness. In the most basic case, there will be statements about the world ('Everest has a surface area of exactly 20 square kilometres') which are neither true nor false.

Williamson lays out the position clearly:

If one abandons bivalence for vague utterances, one pays a high price. One can no longer apply classical truth-conditional semantics to them, and probably not even classical logic. Yet classical semantics and logic are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains. It would not be wholly unreasonable to insist on these grounds alone that bivalence must somehow apply to vague utterances, attributing any contrary appearances to our lack of insight. Not every anomaly falsifies a theory. That attitude might eventually cease to be tenable, if some non-classical treatment of vagueness were genuinely illuminating. No such treatment has been found.¹⁴¹

Sorensen offers a similar thought when considering whether it is advisable to move down the path of revising classical logic: 'Retreat is often wise. I have done my share!

¹⁴¹ Williamson, *Vagueness*, p. 186.

But we should not retreat from standard logic to rescue speculative hypotheses about how language operates'.¹⁴² Sorensen names language because it is true that even non-metaphysical accounts of vagueness can result in the need to move away from classical logic (at least in some instances); the most prominent theory in this vein being supervaluationism, which we examined in Chapter 1, and which both Williamson and Sorensen see as misguided.

The key point of contention is how to resolve the conflict between our intuitions that vagueness occurs in many scenarios, and that this vagueness contradicts theorems of classical logic. Going back to our characterisation of vagueness from the previous chapter, there is a deep intuition that in cases such as sorites series, there is the phenomenon of tolerance; wherever it is located, something is happening which means that small enough changes can be tolerated without resulting in a change of status. At the same time, large changes do have the ability to result in a change of status. As small changes accrete to large changes, we are able to build contradictions through the application of seemingly benign logical constructions.

For the believer in metaphysical vagueness (as well as the supervaluationist), the intuition that tolerance is a genuinely existing phenomenon is enough of a motivation to decide that classical logic does not apply in these situations; the existence of vagueness demands a revision (on a larger or smaller scale). Williamson and Sorensen refuse. For them, classical logic (and semantics) is 'vastly superior', has been proven to work over hundreds (thousands?) of years, and underlies vast swathes of other theoretical work. The apparent existence of some troubling cases, then, is not nearly enough to warrant the overturning of this system. Far better, they say, to offer an explanation of Soritical situations which retains classical logic, at the price of offending against the driving intuitions for vagueness (and a few incredulous stares). The success of their explanation will be examined in the following section.

This argument is not one which seeks to prove the incoherence of taking vagueness seriously, but is instead an appeal to theoretical efficiency, parsimony, and power. As is implied by the Williamson quote above, it is (hypothetically) possible that were a non-

¹⁴² Roy Sorensen, *Vagueness and Contradiction* (New York: Oxford University Press, 2001), p. 8.

classical system developed which could approach classical logic in its ‘simplicity, power, past success, and integration with theories in other domains’, then he would be willing to accept it. This seems vanishingly unlikely. The more promising avenue for overcoming Williamson and Sorensen’s resistance is for a non-classical account of vagueness to be ‘genuinely illuminating’. It is hard to imagine that any such account would clear Williamson’s bar for illumination, but I believe it may be possible for less partisan members of the conversation, or for non-philosophers generally.

I think the efforts of supervaluationists and proponents of metaphysical vagueness have been well spent in both justifying any revisions to classical logic by appeal to the pervasive nature of vagueness, and in making clear where and when classical logic must be amended, and where it can be retained; no-one is advocating for revisionist systems for the pure joy of it (apart from possibly Priest), but sincerely trying to build the best system for making sense of the phenomena of vagueness as we experience them. The rest of this thesis can perhaps be taken as an attempt to show that ‘genuinely illuminating’ non-classical theories of vagueness are available which justify moving away from Williamson and Sorensen’s beloved classical world.

In examining all of the above arguments, I have not done justice to the full discussion on these interesting issues; I hope, however, that I have shown that none of them are powerful enough to rule out metaphysical vagueness as a possibility worth exploring. In doing so, I aim to have kept up my project to remove the burden of proof from the believer in metaphysical vagueness; there is nothing inherently incoherent or inconsistent about it, and it should be considered on a par with linguistic or epistemic accounts of vagueness. We can now move on to examining these competing frameworks in more detail.

4. Previous accounts of metaphysical vagueness

Now the stage has been set, I will examine some previous attempts to offer accounts of metaphysical vagueness. Before doing so, I discuss what is meant by metaphysical vagueness, and therefore make clear what it is that the theories are trying to account for.

In discussing the theories, I organise them into two categories (though this is primarily an administrative division, and does not necessarily reflect a difference in kind between the account). The first are what I call ‘specific’ theories, which centre their accounts around a particular type of metaphysical vagueness. These include the vague mereology of Nicholas Smith, Michael Morreau’s account of vague objects as having ‘questionable parts’, and Neil McKinnon’s account of vague simples. Each of these theories is found to have challenges, and to be insufficiently general (in their current form, without expansion) to account for the myriad manifestations of metaphysical vagueness in the world.

The second category of theories are what I call ‘general’ theories, which do meet this bar, and attempt to offer an account of metaphysical vagueness more widely. The first two general theories of metaphysical vagueness I discuss are the supervenient theories of Akiba, and Barnes with Williams. Supervenientism is the dominant account of semantic vagueness, and so understandably many have tried to apply it to the metaphysical phenomena too. Akiba does this by creating a heavily metaphysical supervenient structure, in which there are actually existing ‘precisificational worlds’. I identify several problems with this, and overall I do not find it to be plausible. Barnes and Williams utilise an ersatz view of precisificational worlds, and so do not run into these metaphysical problems. As Akiba says, though, their account turns out to be insufficiently metaphysical, and is really more of a model than a theory. I argue that any supervenient account will be stuck on the horns of this dilemma; a robustly metaphysical supervenientism is implausible, while a plausible supervenientism is insufficiently metaphysical. Finally we come to Wilson’s ‘general’ theory of indeterminate states of affairs, under which entities may have a determinable property,

but have no unique determinant under that determinable. This ingenious account is more successful than the supervenientists in locating metaphysical vagueness in the world itself (rather than in a meta-structure), but, as I will demonstrate, it ultimately has its own drawbacks.

4.1 What is metaphysical vagueness?

Before we begin our analysis of some existing theories of metaphysical vagueness, we should discuss what exactly it is we mean when we use the term ‘metaphysical vagueness’. It is something of a general term; at base, it refers to any phenomena of vagueness which occurs in the metaphysical realm, as opposed to vagueness that arises out of semantics or epistemology. That is, metaphysical vagueness is vagueness which occurs *in the world*, and would occur even if humans had never existed or developed consciousness, language, beliefs, and knowledge. In Sattig’s admirably clear wording, metaphysical vagueness is metaphysical ‘in the sense of being independent of conceptual, linguistic or epistemic representation’.¹⁴³ Metaphysical vagueness is posited to manifest in various ways, and in Chapter 1 we briefly introduced some of these instances. I will now go through these in a bit more detail, with the proviso that not everyone will agree with the existence of all of them, or agree on exactly how they interact with each other.

The first manifestation was mereological vagueness, or vagueness as to whether something is a part of another thing. These are cases such as the droplet on the edge of a cloud; the droplet is not so close to the centre of the cloud as to be definitely part of it, nor so far away as to be definitely not part of the cloud. Plausibly, such a droplet is indeterminately part of the cloud; it is in the penumbral region between parts and non-parts. Moreover, if one considers clouds to be genuine metaphysical objects, then the status of this vague parthood relation is metaphysical in nature (of course, if the reader does not consider clouds to be the sort of things which would appear in a carving-

¹⁴³ Thomas Sattig, ‘Mereological Indeterminacy: Metaphysical but Not Fundamental’, in *Vague Objects and Vague Identity: New Essays on Ontic Vagueness* eds. Ken Akiba and Ali Abasnezhad (Dordrecht: Springer, 2014), p. 26.

nature-at-the-joints classification of objects, they should feel free to substitute in the sort of thing which would).

A related manifestation is compositional vagueness; vagueness as to when a group of things compose another thing. This could be a case such as a fairly diffuse collection of water droplets hanging in the air, with the question being whether they together compose a cloud. Plausibly, they vaguely compose a macro-object, not definitely doing so, but not definitely not doing so either. Another example which may exhibit compositional vagueness is the above droplet on the edge of the (definitely existing) cloud; does the set of all droplets which are definitely parts plus the droplet in question compose an object? As you can imagine, depending on your theory of metaphysical vagueness it may well be the case that any instance of mereological vagueness automatically implies a related instance of compositional vagueness.

Another potential manifestation of metaphysical vagueness is existential vagueness, that is vagueness as to whether an object exists or not. Again, it may be the case that instances of compositional vagueness, as sketched above with the diffuse collection of water droplets, are also cases of existential vagueness; if it is vague as to whether the droplets compose a cloud, and it is the case that if they do not compose a cloud then there will be no object existing in that vicinity, then it follows that the potential cloud composed by them will exhibit existential vagueness. However, existential vagueness does not necessarily have to go hand-in-hand with compositional vagueness, as the notions could plausibly come apart in scenarios such as that of the statue and the clay. Imagining for the sake of argument that the lump of clay has purely determinate parts, and none are lost or gained during the modelling of the clay by the artist, there may still be existential vagueness regarding whether the lump of clay, the statue, or both, exist (at various times).

The final important instance of metaphysical vagueness is potentially the most controversial: vagueness in identity. Potentially, the statue and the clay example above could be an instance of vagueness in identity, with it being a vague matter as to whether the lump of clay is identical with the statue (if both are taken to exist). Yet another well-known potential example of vagueness in identity is the case of Theseus' ship, whereby a ship slowly has pieces removed and replaced, whilst a new ship is created out of the

removed pieces. At some point in this sequence, it is plausibly vague as to whether the original ship is identical with the ship containing replaced parts, or the new ship made out of the removed parts.

Sainsbury calls this sort of vagueness individutive vagueness (' x is individutively vague if for some y it is indeterminate as to whether $y = x$ '), and supposes that its potential existence is why many are leery of accepting metaphysical vagueness, and the vagueness of objects; this is because they assume that individutive vagueness necessarily leads to a vague identity relation, which is often taken to be contradictory (see section 2.1.1).¹⁴⁴

At the risk of slightly digressing, I think it is worth summarising Sainsbury's argument for why this is not the case, and why we can have, according to his 'slogan', 'vague objects without vague identity'.¹⁴⁵ He starts by pointing out that, if you accept the possibility of vague objects, then an atomic sentence ascribing a property to an object may be made vague by vagueness inherent in the object, rather than the property. For instance, one could imagine that the sentence 'Snowdon has a surface area of 150 acres' might have an indeterminate truth value, if Snowdon is mereologically vague (Sainsbury calls this type of vagueness compositional vagueness, somewhat confusingly, but I will stick with my above terminology). The property of 'having a surface area of (exactly) 150 acres' is intuitively sharp, but it does not definitely apply to Snowdon, nor definitely not apply to it, due to the vagueness of whether some things are part of Snowdon or not.

Sainsbury then argues that the same considerations apply in cases of individutive vagueness; the vagueness as to whether ' x is identical with y ' can reside in the vagueness of x , or y , or both, rather than the vagueness of the identity relation. For instance, the sentence 'Snowdon is identical with the collection of matter m ', where m is some collection of all the determinate parts of Snowdon and a few indeterminate parts, is presumably going to be vague. This vagueness, though, is not due to the identity relation, but due to the vagueness of Snowdon and its parthood relations; the

¹⁴⁴ R. M. Sainsbury, 'What Is a Vague Object?', *Analysis*, 49.3 (Jun., 1989), p. 101.

¹⁴⁵ *Ibid.*, p. 103.

conditions for identity to hold are clear, but we cannot say whether they are fulfilled or not due to the vagueness of Snowdon.

Now that we have an idea of what I mean by metaphysical vagueness, let us move to examine some previously offered theories of how it works.

4.2 Specific theories

4.2.1 Smith

As mentioned earlier, a common motivating factor for the positing of the existence of metaphysical vagueness, and compositional vagueness in particular, is the consideration of composition and when it occurs. The paradigmatic instance of vagueness, the paradox of the heap, can be thought of as such a case: two grains of sand definitely do not compose a heap, and twenty thousand grains definitely do, but adding just one grain is not enough, at any particular stage, to move from non-composition to composition. To take a slightly less formless example, consider the process of building a ship. To begin, the assortment of planks in the workshop do not compose an object, but eventually they do compose a ship; at what point does this occur? van Inwagen called this the ‘Special Composition Question’; ‘In what circumstances do planks compose (add up to, form) something?’.¹⁴⁶ This is specifically a metaphysical question and, as Nicholas Smith notes, it seems intuitively to be a locus of vagueness: ‘nonphilosophers tend to think that *mereological composition is a vague matter*: sometimes it occurs, sometimes it does not, and sometimes it *sort of* occurs’.¹⁴⁷

As van Inwagen outlines, there are three main avenues of reply to the special composition question. Two are what he calls extreme answers: nihilism, which argues that composition never occurs, and universalism, which says that composition always occurs. The moderate answer is known as restricted composition; sometimes composition occurs, and sometimes it does not, for various proposed reasons. These

¹⁴⁶ Peter van Inwagen, *Material Beings* (Cornell University Press, 1990), p. 21. It is the ‘Special’ composition question as opposed to the ‘General Composition Question’, which concerns the nature of composition in itself.

¹⁴⁷ Nicholas J. J. Smith, ‘A Plea for Things That Are Not Quite All There: Or, Is There a Problem about Vague Composition and Vague Existence?’, *The Journal of Philosophy*, CII.8 (Aug., 2005), p. 381.

reasons are generally taken to be precise criteria, which are either fulfilled or not. This is because of the prevalence of what Smith calls the ‘Orthodox Argument’:

O1: Whenever there is a case of vague composition, there is a case of vague existence.

O2: There are no cases of vague existence (because vague existence is at least impossible, if not nonsensical).

OC: Therefore there are no cases of vague composition.¹⁴⁸

Since there is no vague composition, a restricted composition view entails that at some point in the building of the boat, adding one plank (or even one nail) transforms the collection from composing nothing to composing a boat. Such sharp restricted composition views have few adherents, though not on account of their sharpness; it is difficult to find a principled criterion for composition which accords with intuition in all proposed cases of composition. Nihilism has a few notable proponents, including Peter Unger, Gideon Rosen and Cian Dorr.¹⁴⁹ Unrestricted composition is the most widely accepted answer, with David Lewis and Theodore Sider among its followers.¹⁵⁰

We have three answers to the Special Composition Question, then, which Smith notes, and I agree, are ‘each horribly implausible’; in fact, ‘the reader should not need much reminding that each of these views is intuitively repulsive’.¹⁵¹ This is of course not an argument against the views, but an incitement to possibly reconsider the basis on which they have come to be accepted. If the ‘Orthodox argument’ leads us to accept such ‘prima facie absurd’ theories, ought it to be accepted? Smith argues not.

Premise O1 is not really in doubt; whatever one takes vague composition to consist in, it seems inevitably to lead to vague existence. After all, if it is vague whether these five planks of wood compose a bench, the question can be asked of whether the bench in question exists. The only basis we have for deciding whether the bench exists is whether

¹⁴⁸ Ibid, p. 381.

¹⁴⁹ See Gideon Rosen and Cian Dorr, ‘Composition as a Fiction’, in *The Blackwell Guide to Metaphysics* ed. R. Gale (Oxford: Blackwell, 2002), pp. 151–174; Peter Unger, ‘The Problem of the Many’, *Midwest Studies in Philosophy*, 5.1 (Sept., 1980), pp. 411–67.

¹⁵⁰ See Theodore Sider, *Four-Dimensionalism: An Ontology of Persistence and Time* (Oxford: Clarendon, 2001); David Lewis, *On the Plurality of Worlds* (Oxford: Blackwell, 1986).

¹⁵¹ Smith, ‘A Plea for Things That Are Not Quite All There’, p. 382.

there are some parts which compose it (since we can assume that such a bench is not a mereological simple). Since this question has a vague answer, so must the question of whether the bench exists. Theoretically it is possible to resist this fairly quick argument, but it does not seem promising as a view. Especially since, as Smith argues, we are able to resist O2 more easily. O2 states that there can be no cases of vague existence; Smith takes it as his 'aim... to *make good sense* of vague composition, vague existence, and the relations between them'.¹⁵²

He does this by building on two observations: one is the aforementioned intuition that during the building of a boat, the boat only gradually comes into existence as it gradually comes to be composed by the various planks and nails; the other is that such *recherché* mereological fusions as Westminster Bridge plus my left hand, while not actually being an object, are possible objects. The combination of these observations provides the basis for drawing the distinction, as Smith does, between *notional* and *concrete* mereology. Notional mereology is the mereology which governs possible objects; as any purported mereological fusion is at least a possible object, notional mereology 'is classic extensional mereology with unrestricted composition'.¹⁵³ Concrete mereology, conversely, governs the objects which actually exist (at this world or another); as such objects can go in and out of existence gradually, concrete mereology is vague. Smith takes it to have a lot to do with 'contact and adhesion'; as a plank is placed in position during the building of the boat, glued on, and the glue dries, 'it gradually becomes a concrete part of my boat to a greater and greater degree, until finally it is a definite concrete part'.¹⁵⁴

To go into slightly more detail, the picture is somewhat modal in nature. There is an overall domain, D , which includes all objects from all possible worlds; this is the domain of notional mereology. Then we have a subset of D , D_1 , for each world W_1 which contains the objects existing at that world; D_1 , however, is a fuzzy subset of D , since when it is vague whether a given set of objects compose another, it is vague whether their fusion exists at that world. More specifically, 'a notionally composite object (an

¹⁵² Ibid, p. 383.

¹⁵³ Ibid, p. 385.

¹⁵⁴ Ibid, pp. 385-386.

object in D with more than one notional part) fades into/out of concrete existence (at a given world) as its notional parts become concrete parts to a greater/lesser degree (at that world).¹⁵⁵

It is worth noting here that Smith explicates his theory using a system of fuzzy logic/set theory, as he believes this to be ‘the correct framework for the treatment of vagueness’; however, it is not essential to the theory, and it is possible to reformulate the account using a different background analysis of vagueness, while still maintaining its overall structure.¹⁵⁶

What we have then, is a two-tiered system of mereology. Notional mereology is precise, maintains unrestricted composition, and is the universal background; everything in the set D ‘exists *simpliciter*’, and it is D over which the existential quantifier ranges.¹⁵⁷ In more detail, Smith defines a primitive binary, non-fuzzy, relation, proper parthood, which is irreflexive, transitive, and asymmetric. He introduces a supplementation principle, by which if there exists a proper part, y , of an object x , then y must be supplemented by at least one other proper part, z , disjoint from y . Notional mereology is also extensional; objects having exactly the same proper parts are identical. Finally, unrestricted composition is given by the following principle: For every set S , there exists exactly one x such that x is a fusion of S .

Concrete mereology restricts composition, and is really the way in which certain properties are determined; an object can have ‘the property of existing at world w ’ to ‘any intermediate degree’.¹⁵⁸ It is concrete mereology in which we are presently most interested, as that is where the vagueness resides. Smith explores the details by considering, in turn, each of the principles which hold for notional mereology, and discussing whether they hold for the fuzzy notion of concrete parthood. Firstly, it is important to note that in order for some x to be a concrete part of y (to any degree), x must be a notional part of y : ‘notional parthood is one of the determinants of concrete parthood (but not the only one)’.¹⁵⁹ Concrete parthood is also irreflexive, asymmetric,

¹⁵⁵ Ibid, p. 386.

¹⁵⁶ Ibid, p. 384.

¹⁵⁷ Ibid, p. 390.

¹⁵⁸ Ibid.

¹⁵⁹ Ibid, p. 396.

and transitive, although the transitivity is slightly modified. As concrete parthood is a relation of degree, its transitivity is stated thus: 'at any world and time, the degree to which x is a concrete part of z is greater than or equal to the minimum of the degrees to which x is a concrete part of y and y is a concrete part of z '.¹⁶⁰

Next, Smith considers supplementation: do we want it to be the case that 'no object can have one proper concrete part without also having another proper concrete part which is disjoint from the first'?¹⁶¹ He answers in the negative, influenced by examples such as part-way completed sculpture. Supposing that a part, b , of a sculpture is the only part to be so far mounted on a plinth. Smith takes it to be the case that b is a concrete part of the sculpture, even though there are as yet no other concrete parts of it.

Now we turn to extensionality, and again Smith denies that we should embrace such a principle for concrete mereology. He argues against a world-bound extensionality principle by imagining a case where you are building a boat, and have cut and numbered every part of it. Accidentally, you have cut and numbered two foredecks, a and b . When everything is constructed apart from the foredeck, we have the following situation; every part, z , of the semi-constructed boat is a concrete part of the fusion of the semi-constructed boat plus a , and a concrete part of the fusion of the semi-constructed boat plus b , to the same degree. Yet we do not want to say that these two fusions are identical.

But what about a universal concrete extensionality principle, not tied to one world? What about cases where z is a concrete part of x to the same degree as it is a concrete part of y not just in this world, but at every world and time? That is, there is never a difference in concrete parthood at any world; perhaps in that case we should say that $x=y$. Again, Smith answers that we should not accept such a principle, but for a different reason. In order to invalidate such a principle, x , say, would have to have a notional part which y lacks, since their concrete parts never differ. But, since their concrete parts never differ, this notional part is never realised as a concrete part of x ; that is, 'there is a notional fusion which is never fully realized, at any time in any world'.¹⁶² A notional

¹⁶⁰ Ibid.

¹⁶¹ Ibid, p. 397.

¹⁶² Ibid, p. 398.

fusion was meant to capture the idea of a possible object, so having a possible object which does not exist at any possible world may offend against this. However, whether or not you agree with this principle, it ‘has nothing to do with concrete mereology as such’.¹⁶³ Concrete mereology does not care whether there is only one possible world or a million, and whether there are infinitely many notional fusions which are never fully realized or none; thus there is no need to take a stand one way or another on this principle.

Finally, we come to the question of the existence of fusions; the question of ‘under what conditions (specified in terms of facts about concrete parthood) this notional fusion *concretely* or *actually* exists at a time in a world’.¹⁶⁴ Smith’s answer is as follows: for any notional fusion x , its degree of concrete existence is the sum of the degree of concrete parthood of its notional parts multiplied by their degree of importance to x .¹⁶⁵ Degrees of importance are introduced to deal with the fact that my head, say, is more important to my existence than one of my fingernails. Were my head to become concretely part of me to degree 0, that would be highly damaging to my degree of concrete existence; my fingernail, less so.

Having outlined Smith’s account of concrete mereology, we can return to the previously-mentioned argument that the vagueness of mereology leads to vagueness regarding ‘*how many things there are in the world*’; this is seen as ‘impossible/obscure/nonsensical’.¹⁶⁶ Smith points out that such problems arise for ‘anyone who countenances fuzzy sets’; for instance, someone who takes ‘bald’ to be an indeterminate predicate for linguistic reasons faces the question of how many men standing in a room fall under the predicate, when five of the men are determinately bald and there are a few borderline cases.¹⁶⁷ For Smith, ‘concretely exists’ is simply another predicate, picking out a fuzzy subset of the overall domain D , and so there is no special problem raised here. This is fortunate, for Smith notes that ‘someone who thinks that the unrestricted existential quantifier \exists is vague... really would have a big problem

¹⁶³ Ibid.

¹⁶⁴ Ibid, p. 399.

¹⁶⁵ Ibid, p. 401.

¹⁶⁶ Ibid, p. 412.

¹⁶⁷ Ibid.

here'.¹⁶⁸ That is, Smith's strategy is to agree that vagueness about how many things genuinely exist *simpliciter* is impossible, but his theory does not entail such vagueness.

Now we have been through the main points of Smith's system, we can appreciate its unconvincing nature. It is again an account which has precise objects as its fundamental suite; that is, the vagueness is not fundamental, as we desire. Vague concrete objects, for Smith, are instances of vagueness affecting precise notional objects. When setting out his theory of concrete parthood, he states that 'nothing can be a positive-degree concrete part of an object, at any world or time, unless it is a *notional* part that object';¹⁶⁹ Smith only ever delineates objects based on their notional, precise, state. What it is to be a vague object, for Smith, is for some of the object's precisely fixed parts to be fastened to less than degree 1 in a particular world, or to have the property of concrete existence to less than degree 1. A corollary of this is that in order for an object to exist to degree 1 at a world, all its notional parts must exist and be concrete parts of it to degree 1; that is, it must be fully precise! Only precise objects can fully concretely exist for Smith. This counts against his claim to be respecting our intuitions regarding vague parthood and existence.

Another problem which becomes apparent when we focus on the link between notional and concrete objects is the 'problem of the many'. Unger introduced the problem of the many with the example of a cloud in an otherwise clear sky, and the attempt to identify which set of water droplets compose the cloud. The problem arises when, after picking out a suitable candidate for that set, we consider the 'very many similar complexes each of which "overlaps" it just slightly, sharing constituents with it, except for a peripheral droplet or two, here or there'.¹⁷⁰ Whichever set of droplets is chosen, it 'has nothing objectively in its favor to make it a better candidate for cloudhood than so many of its overlappers'.¹⁷¹ We are left with a multitude of candidate clouds, and no principled way of privileging one over any other. Each have as good a case as any other to be

¹⁶⁸ Ibid, p. 412 fn. 60.

¹⁶⁹ Ibid, p. 395.

¹⁷⁰ Unger, 'The Problem of the Many', p. 415.

¹⁷¹ Ibid.

considered a cloud; either they all pass, and we have thousands of clouds, or they all fail and we have none.

Smith's notional composition is unrestricted, so, just as we had a multitude of candidate clouds, there exists the fusion of all my determinate parts plus 10,000 hairs, the fusion of all my determinate parts plus 10,001 hairs, and so on. Each of these notional fusions will concretely exist to a fairly high degree, since all my determinate parts are fully concrete parts of me, and various hairs are concretely part of me to various degrees. None of them appears to have a privileged claim to being 'me', and again we either have thousands of me or none.

This is, suffice to say, an unappealing consequence of the view. Smith admits that all these concretely existing objects are non-identical with each other, since identity is classical and based on notional parts. However, while identity is classical, there is a somewhat related notion which is vague: concrete overlap. Notional fusions which are non-identical can concretely overlap to a higher or lesser degree; in 'problem of the many' cases, there will be a high number of non-identical fusions which concretely overlap to a very high degree. Lewis argued that when there are many things which are 'almost identical' to each other, in everyday cases it can be 'a blameless approximation' to say that there is only one thing present, even though in stricter circumstances this would be false.¹⁷² Smith's account leaves open the same reply; it is acceptable to sloppily talk of there being only one concretely existing fusion in the vicinity, since the many strictly non-identical notional fusions concretely overlap to a high degree. This is fine as far as it goes, but is another blow to Smith's claim to be respecting intuitions; we believe there to be one vague object in the vicinity, not a multitude of precise fusions which overlap in their degree of concrete existence. Moreover, this intuition does not seem to dissipate when we closely attend to the problem, as it might be supposed to when we move away from loose ways of talking.

This problem, of course, is not unique to Smith, and is not even a consequence of the vague aspects to his theory of mereology. In fact, it is an area where one might expect a vague theory to be free of worry, since the problem of the many is motivated by

¹⁷² Lewis, 'Many, But Almost One', p. 178.

considering objects which are finely and precisely individuated. The reason it is faced by Smith is his insistence on a two-level account, where the vagueness occurs in concrete mereology, and there is a higher, classical, level of notional mereology.

This insistence is additionally underwhelming, since there does not seem to be a compelling reason to introduce the notional level at all; it brings in the question of levels of existence, which is notoriously tricky. Smith claims that having existence *simpliciter* as notional existence, and concrete existence being a predicate, is unproblematic, as it simply tracks the ‘*intuitively* clear distinction...between the ones which exist *around here*...and the rest’.¹⁷³ This is debatable; it is not clear that the notional/concrete, precise/vague distinction offered by Smith is the same as the actual/non-actual distinction. Lewis, for one, would argue that the difference between things which exist ‘around here’ and those which don’t is merely one of location.¹⁷⁴ Even on Smith’s own account, objects which exist in other worlds, i.e. not around here, have concrete existence too. The distinction is not between things around here and the rest, but between things as they exist *at any world* and their nominal, non-world-bound status. As the name of ‘concrete mereology’ implies, the real distinction seems to be operating along the lines of concrete versus abstract, although with the abstract fusions having the ability to become concretised (to some degree or other).

Moreover, Smith sees it as a virtue of the system that the broadest domain is governed classically, and so can accommodate the various classical intuitions. However, he has to introduce a non-classical system to deal with the vagueness of concrete mereology and accommodate the various intuitions regarding vague phenomena. Smith believes that this double system gives him the ‘double advantage’ of being able to ‘*both* make perfectly clear sense, *and* satisfy widespread intuitions about vague composition and existence’.¹⁷⁵ I would argue, though, that Smith’s proposal ends up as a halfway house, convincing on neither count. It does not do full justice to our intuitions about vague objects, and will not be amenable to someone classically-minded and predisposed to dismiss the possibility of vague composition. That is, while it may seem tempting to pick

¹⁷³ Smith, ‘A Plea for Things That Are Not Quite All There’, p. 389.

¹⁷⁴ Lewis, *On the Plurality of World*; in particular section 1.9 (‘Actuality’).

¹⁷⁵ Smith, ‘A Plea for Things Which Are Not Quite All There’, p. 405.

up the advantages of both approaches, one must be mindful of the possibility of inheriting both sets of objections.

4.2.2 Morreau

Michael Morreau offers an alternative account of the vagueness of objects, which is also mereologically based; as he puts it, ‘Some ordinary material objects are vague because they have questionable parts’.¹⁷⁶ It is easy to get ‘an intuitive hold on questionable parthood by thinking of parts that have come loose and will be lost, such as the cat’s whisker’ which moves from definite parthood to definite non-parthood by a gradual process.¹⁷⁷ Before this is explained in more detail, though, a contrast is drawn with objects which cannot have questionable parts: ‘*quantities of matter*’, such as ‘the soap in China’.¹⁷⁸

Such objects do not lose parts when they become physically disjoint, whether this is sudden or gradual; if I cut a Chinese soap bar in half, ‘the soap in China’ has not become smaller, but merely changed (some of its) location. This is because quantities of matter ‘have their parts essentially’.¹⁷⁹ We do not think that ordinary material objects act in this way; if my finger is cut off, it is no longer part of me. Thus there is ‘more to part-whole relations among material objects than just inclusion among quantities of matter’.¹⁸⁰ Moreover, unlike in the case of quantities of matter, not every subset of the matter of a material object is a part of that object. According to Morreau, while a whisker or a foot should be counted as part of a cat, we should not count an arbitrary fusion of various cells from around its body. Morreau argues that the language of ‘English makes a handy distinction’ between these two types of cases: the ‘arbitrary portion of cat is “some of” or “part of” the cat... But it is not “a part of” the cat, in the same full sense in which a firmly attached whisker is.’¹⁸¹ According to Morreau, this is

¹⁷⁶ Michael Morreau, ‘What Vague Objects Are Like’, *The Journal of Philosophy*, 99.7 (Jul., 2002), p. 339.

¹⁷⁷ *Ibid.*

¹⁷⁸ *Ibid.*

¹⁷⁹ *Ibid.*

¹⁸⁰ *Ibid.*, p. 340.

¹⁸¹ *Ibid.*

because the whisker is a functional part of the cat, whereas the arbitrary selection is not.

In order for something to count as a functional part of a cat (or any organism), it must play a proper role in the functioning of the cat, and be ‘furthermore suitably integrated into the rest of the cat for fulfilling this function’.¹⁸² Of course, there is the further question of what exactly is the functioning of the cat (to survive? To catch mice? To ignore its owner apart from when hungry?). Morreau states that he is only interested in ‘*proper functions of the sort that interest functional anatomists and reverse engineers*’; for example, it is the proper function of the cat’s whisker to detect movement and objects in the vicinity of the cat’s head.¹⁸³

The vagueness arises due to the fact that organisms are always in the process of losing (and gaining) functional parts:

Crucially, for my present purpose, organisms are always gaining and losing functional parts. Once the loose whisker was firmly attached to the rest of the cat. Then it was suitably integrated for detecting objects and it was a functional part of the cat. Soon, when it has dropped off, it will no longer be suitably integrated and it will no longer be a part of the cat. This is why organisms are vague. They gain and lose their functional parts gradually. There is no precise moment at which the whisker becomes so loose that it can no longer fulfill its function. It does not go in an instant from being a part of the cat to being a nonpart but is, for a time, a questionable part of the cat. Indeed organisms, maintaining themselves in continuous flux, are more vague than meets the eye. As a flake of the cat’s skin loosens, it does not instantly become so loose that it can no longer do its bit to keep the inside environment of the cat separate from its external environment. It too takes a while to quit its small part in the life of the cat and is, for a while, a questionable part of the cat.¹⁸⁴

This model, Morreau says, can also be extended to non-organisms, such as a statue, by introducing the relation of composition. Composition holds between an object and a quantity of matter, and is supposed to be distinct from identity; a quantity of clay may compose a statue, without being identical to it, or a quantity of animal matter may compose a cat, without being identical to it. Morreau argues that material objects such as statues ‘can be vague if they are indefinitely constituted by quantities of matter

¹⁸² Ibid.

¹⁸³ Ibid.

¹⁸⁴ Ibid., p. 341.

without being indefinitely identical to them'.¹⁸⁵ The notion of a functional part seems to have been jettisoned somewhat, but the idea seems to be this; the statue is not identical with any candidate quantities of clay, but it may be indeterminately composed by them at one time or another. As we observed earlier, quantities of matter cannot lose parts, so all of the candidate quantities of clay are permanent and unchanging in their membership; the statue could not have questionable parts (and therefore be vague) if it were identical to any of these quantities of matter. However, it is plausible that we should not want the statue to be identical with a quantity of matter; a statue can be altered, worn down, retouched, while retaining its identity. Therefore it makes sense, according to Morreau, to say of the statue that over its lifetime it is indeterminately composed by various quantities of clay.

Objects, for Morreau, can also exhibit higher-order vagueness:

Natural change is continuous. Just as there is no precise moment at which the loose whisker becomes a nonpart of the cat, there is none at which it becomes a questionable part, either. As well as a first-order penumbra of questionable parts, in limbo between parthood and nonparthood, a vague object can have second-order penumbrae of objects that questionably re questionable parts, in limbo either between parthood and questionable parthood, or between questionable parthood and nonparthood. There can be third-order penumbrae of objects in limbo between penumbrae of lower orders, and so on. Ranking vague objects according to their highest-order penumbrae, we can distinguish vague objects of various orders of vagueness.¹⁸⁶

Morreau also makes an effort to argue that his account of vague objects can actually be consistent with the principles of classical mereology, if one thinks those are important and worth preserving. I will not go into the details here, primarily because Morreau himself says that whether or not to stick with classical mereology is independent of accepting his theory:

I do not mean to suggest that mereology has to be classical, though. My point is just that it can be classical, even if there are vague objects. There might be good reasons to reject classical mereology. But the need to accommodate ontological vagueness, by itself, is not one of them.¹⁸⁷

¹⁸⁵ Ibid., p. 342.

¹⁸⁶ Ibid., p. 344.

¹⁸⁷ Ibid., p. 356.

The next move which Morreau makes which is interesting for our purposes is the introduction of what he calls ‘sharpenings’. A sharpening of a vague object is a way of getting rid of one level of vagueness by imagining that the part(s) in question are suitably distributed. At the limit case where there is only one level of vagueness, a sharpening is thus a way of imagining the object in question to be fully precise. Morreau says we can imagine the sharpenings of an object to be ‘just it but in another possible world’.¹⁸⁸ As can be imagined, Morreau utilises a supervaluational approach to evaluate sentences over all possible sharpenings; a sentence such as ‘this cat has four legs’ will be supertrue if it is true on all possible sharpenings of the cat, superfalse if it is false on all possible sharpenings, and neither if it is true on some and false on others.

Morreau’s account then, is an interesting one, utilising supervaluational evaluations over the possible sharpenings of vague objects, while the vagueness of the objects themselves is located in the indeterminacy of whether their potential parts are integrated well enough into the object’s function. However, it seems to me to face a number of drawbacks. Firstly, the account is quite specific, and only applicable to a small class of things. It is really only an account of the vagueness of objects, in particular their mereological vagueness, to use the taxonomy I introduced at the start of the chapter. Even then, it is unclear whether the account can apply to all objects, or only those which have some function or intention behind them. Morreau begins by discussing in some depth the way that whiskers contribute to the proper function of being a cat, and then moves on to apply his theory to a statue which may be vaguely composed by various lumps of clay; he is not clear as to whether the relation of composition is based on some considerations of form, or intended form, or whether any macro object can stand in the composition relation to a quantity of matter. Do collections of rocky matter compose Kilimanjaro, on Morreau’s view, or is Kilimanjaro simply identical with one of them? As I have intimated, it is preferable for a theory of metaphysical vagueness to apply as generally as possible, but Morreau does not seem to aspire to this.

¹⁸⁸ Ibid., p. 347.

The second objection to Morreau's view is in some senses the opposite, that it is too general, or perhaps too non-committal; it does not actually say very much about what it means for something to be a questionable part of an object. It is therefore unsatisfying in being relatively uninformative, about a fairly restricted class of metaphysical vagueness.

4.2.3 McKinnon

The final 'specific' theory of metaphysical vagueness we will examine is perhaps the most specific of them all; McKinnon's brief argument for the possibility of vague mereological simples. McKinnon begins by considering the Evans/Salmon argument against vague identity, as well as Sider's argument against vague composition.

McKinnon contends that even if both of these arguments are wholly conceded (which, of course, many do not agree with), then it could still be the case that objects have vague boundaries.

He proceeds by imagining mereological simples, objects which have no proper parts, but which do have spatial extension. That is, mereological simples which are not mere point-particles, but which extend through space; they are located at more than one (though presumably contiguous) point or region. McKinnon then posits that while these simples may be determinately located at many of the points or regions, there is nothing to stop it being the case that 'there are also points and regions at which the simple is indeterminately located'.¹⁸⁹ That is, it is possible for these spatially extended simples to have vague boundaries. That is the theory, as it goes.

McKinnon considers the objection that such mereological simples just do not have the requisite structure to ground such boundary uncertainty, but claims that, once spatially extended simples are accepted, there is no principled reason to deny that they might have vague boundaries. The argument is by analogy to a similar objection which might be levelled against extended simples having a certain shape:

If an object has spatial extension then it has a shape. An object's having spatial extension is consistent with its having all sorts of shapes. We can

¹⁸⁹ Neil McKinnon, 'Vague Simples', *Pacific Philosophical Quarterly*, 84 (2003), p. 395.

explain why different objects have different shapes if the objects have proper parts; the differing shapes are due to the different configurations of each object's proper parts. But differing shapes among simples cannot be accounted for in this way. Whatever shape a simple has, it has as a matter of brute fact.¹⁹⁰

McKinnon admits that it is indeed a brute fact, but argues that there is no principled reason for denying that such brute facts could exist, and that it is not blatantly incoherent to say so. Extended simples are indeed strange objects, different in kind to mereological complexes, and they have their shapes in a different manner too.

A similar strategy can be used to defend the possibility of vague boundaries for the vague simples. While it may be the case that for mereological complexes, the existence of a vague boundary must be explained in terms of vague parthood or vague composition, there is no principled reason to suppose that this shows that extended simples cannot have vague boundaries as a matter of brute fact. In fact, it seems plausible to me that the sort of brute facts about an extended simple which will show its shape will also show if it has a vague boundary; these being the facts about which points and regions the object is located in, whether it be determinately or indeterminately so.

McKinnon finishes by mentioning, in passing, that in addition to the possibility of mereological simples with vague boundaries, there may also be 'vague complexes composed of those simples'.¹⁹¹ The details of this are not fleshed out, but it is at least theoretically possible to see that complex objects could be composed of simples, some of which could have vague boundaries. Presumably the complex object could also inherit some of the boundary vagueness from its vague parts.

In any case, we know enough about McKinnon's proposal to see its limitations as a theory of metaphysical vagueness. While such extended simples with vague boundaries are an interesting proposal, and do indeed provide a rebuke to those who would think that Evans/Salmon or Sider-type arguments are knockdowns against the very possibility of metaphysical vagueness, they are not (at least on first inspection) plausible candidates for offering an explanation of what is going on in the instances of

¹⁹⁰ Ibid., p. 396.

¹⁹¹ Ibid.

metaphysical vagueness mentioned at the start of this chapter. It just does not seem right to say that the reason Kilimanjaro is metaphysically vague is that it is composed of simples which have vague boundaries. To be fair to McKinnon, I do not think he was attempting to offer such an explanation. For that sort of answer, we shall need to move to more ‘general’ theories of metaphysical vagueness, which we shall now do.

4.3 General theories

Supervaluational strategies are common (even dominant) in explanations of semantic vagueness. It might be initially plausible, then, that something similar will work in the metaphysical case. In this section I will examine two attempts to do just that. I will first look at the account of Akiba, which he has developed over a number of years. Akiba opts for a hardcore metaphysical interpretation of supervaluationism, on which there are concrete ‘precisificational worlds’ that everyday objects straddle. It will be shown that Akiba’s account faces many problems while invoking heavy metaphysical commitments, and as such is unconvincing.

I will then look at Barnes’ and Williams’ more recent supervaluational theory. They opt for an ersatz view of precisificational worlds on which to base their supervaluational strategy. Metaphysical vagueness, for them, consists in it being the case that more than one precisificational world maximally represents reality. I will agree with a criticism Akiba that makes of Barnes and Williams, that their account is insufficiently metaphysical; in fact, it should be accepted by a linguistic supervaluationist. However, this does not mean that Akiba is better off. In fact, I will conclude that any supervaluational account of metaphysical vagueness will be stuck with the dilemma of being plausible, but metaphysically trivial (à la Barnes and Williams) or metaphysically heavyweight but implausible (à la Akiba). Thus metaphysical supervaluationism should be abandoned.

4.3.1 Akiba

Akiba characterises his account of what he terms *worldly* vagueness by two analogies. The first is, as previously mentioned, to the structure of linguistic supervaluationism; in fact, he says, ‘it may even be considered to be the worldly version of supervaluationism’.¹⁹² The second, and more interesting, is to metaphysical and temporal modalities. Akiba proposes that we should not rest satisfied with the five dimensions currently (generally) accepted, these being the three spatial dimensions, the temporal dimension, and the modal dimension. In order to explain the phenomena of vagueness (and thus more fully explain the phenomena of reality) we must add a sixth, ‘precisificational’, dimension. The precisificational dimension is, as the name implies, ‘made up of *precisified* worlds, in which everything is precise’.¹⁹³

Turning to the notion of coincidence, Akiba can then give his definition of a vague object. Akiba defines coincidence as indiscernibility within a world. He then states that what it is for an object to be vague is that the object coincides with one precise object in one precisified world, and another precise object in another precisified world. To take the example of Tibbles the cat, in one precisified world there is a precise cat, with all Tibbles’ essential properties, and with a precise number of hairs (say, 50,000); in another precisified world, there is a precise Tibbles-cat with 50,001 hairs. Tibbles the everyday, vague object, coincides with both those objects in their own worlds. That is, vague objects extend over precisified worlds. Precise objects extend over precisified worlds, too, but they always coincide with the same precise object in every world.

If the idea is not clear to the reader, it will hopefully be more fully illuminated by noting some consequences of the view. Akiba admits that on this account, a vague object ‘has perfectly sharp boundaries...It is just that those boundaries are what may be called “indeterminate”’; that is, their locations are slightly different in different precisified worlds’.¹⁹⁴ There is also no vague identity on this view; vague objects are simply not identical to any of the precise objects with which they may coincide.

¹⁹² Ken Akiba, ‘Vagueness in the World’, *Noûs*, 38.3 (2004), p. 408.

¹⁹³ *Ibid.*

¹⁹⁴ *Ibid.*, p. 409.

Akiba then seeks to head off an initial objection: namely, that if in every precisified world, Tibbles (and every other vague object) is actually precise, then how (and why) is it that we perceive Tibbles to be a vague object? He gives the following answer: perceiving itself is a vague relation, holding between a vague object and a vague subject. ‘When I perceive Tibbles, one precisification of me perceives₁ one precisification of Tibbles in one precisified world’, another perceives₂ another in another, and so on, ‘where *perceiving_n* is a precisification of the vague relation *perceiving*’.¹⁹⁵ Thus Tibbles (the vague object) will appear to me (the vague subject) as the ‘superimposition’ of his various precisifications.

This reply seems rather implausible. It is acceptable to imagine a perception of a whole arising as an amalgamation of various discrete objects (or parts); after all, when we watch a film, we are actually perceiving thousands upon thousands of still images, which we amalgamate into smooth video. But what does it mean for there to be various precisifications of me *doing* the perceiving, and somehow amalgamating together? No further explanation is given, and it is hard to see what form it could take.

The supervenient element of the account has not been discussed yet explicitly. According to semantic supervenientism, there is indeterminacy as to whether ‘Tibbles’ refers to a cat with 50,000 hairs or a cat with 50,001 hairs; a statement such as ‘Tibbles has 50,000 hairs’, therefore, is neither true nor false, but indeterminate. However, some statements, such as ‘Tibbles has a heart’ are true whichever precise object ‘Tibbles’ refers to; thus we can say that they are super-true (and vice-versa for super-false statements). Akiba’s supervenientism works on the level of coincidence, rather than reference. ‘Tibbles’ always refers to the vague object Tibbles (or ‘Everest’ to Everest, as in Akiba’s below explanation); however, that vague object coincides with different precise objects at different precisificational worlds.

Mt. Everest is itself a vague area, that is, a transworld object that coincides with different precise areas in different precisified worlds; and the term ‘Mt. Everest’ wholly (not partially) refers to the vague object Mt. Everest. If Jones, another transworld object, is on Mt. Everest in some precisified worlds but not on it in other precisified worlds, then it is indeterminate whether Jones is on Mt. Everest. Correspondingly, ‘Jones is on Mt. Everest’ is true in (or with

¹⁹⁵ Ibid, p. 410.

respect to) the former worlds and is false in the latter; thus, it is indeterminate in truth value. 'Jones is either on Mt. Everest or not on Mt. Everest' and 'If Jones is on Mt. Everest and Mt. Everest is in Nepal, then Jones is in Nepal' are determinately true because in each precisified world, Jones is either on Mt. Everest or not on Mt. Everest, and if he is, and if Mt. Everest is in Nepal, then Jones is in Nepal.¹⁹⁶

The advantages which have made supervaluationism so popular as a solution for semantic vagueness carry over to the metaphysical case; Akiba, for instance, is keen to keep classical logic intact. Some disadvantages may appear to carry over as well, though. The phenomenon of higher-order vagueness is often thought to pose a problem for supervaluationism. After all, supervaluationism is essentially positing a third category, between determinate truth and determinate falsity, because a sharp divide was implausible. But in doing so, it introduces two new sharp divides, between determinate truth and indeterminateness, and indeterminateness and determinate falsity; these divides are no more plausible than the original one. More distinctions must thus be made, until there are an infinite amount of divisions between determinate truth and determinate falsity.

On Akiba's metaphysical account, the problem surfaces like so. A new (sharp) divide is introduced between the precise objects which Tibbles never coincides with and the precise objects which Tibbles coincides with at some precisified worlds. This seems as implausible as a sharp divide between Tibbles having 50,000 hairs or 50,001 hairs. Akiba combats this by allowing that the very notion of a precisified world is vague. That is, it could be indeterminate 'whether there is a precisified world that meets a certain specification'.¹⁹⁷ As such, it may be the case that there isn't in fact a sharp divide between the precise objects which Tibbles sometimes coincides with and the precise objects Tibbles never coincides with, as it may be indeterminate whether there is a precisified world containing an intermediate object.

However, this reply seems to me to raise two problems. Firstly, higher-order vagueness, if it is taken to exist, is seen as a pervasive phenomenon; it arises for almost every case of genuine vagueness. Thus, if Akiba's response is to be effective, it must in fact be the

¹⁹⁶ Ibid., p. 411.

¹⁹⁷ Ibid, p. 422.

case that it is indeterminate whether precisificational worlds exist that meet a possibly infinite number of specifications. This seems an unattractive position to be forced into. If it is indeterminate as to whether all sorts of precisificational worlds exist, how much clearer does our account of vagueness become by acknowledging them? Secondly, this strategy only works for the second level of vagueness (taking the first level to be the initial vagueness). While semantic supervenience can accept indeterminacy regarding whether it is indeterminate whether a precisification is acceptable, it is hard to see what mechanism Akiba can invoke. First level vagueness is coincidence with various objects; second level vagueness is indeterminacy of existence for precisificational worlds; what metaphysical phenomena do higher levels of vagueness consist in?

This brings to attention a more general problem: Akiba is keen to emphasise the metaphysical nature of the precisificational dimension. It is meant to be seen on a par with the temporal and modal dimensions. But with this heavy-duty metaphysical commitment comes problems, like those above. Perhaps a less metaphysically robust supervenient strategy is required; this is exactly what Barnes and Williams attempt to do.

4.3.2 Barnes and Williams

Barnes and Williams explicate a framework for metaphysical vagueness which is not analogous to modality, as Akiba does, but in fact uses modality itself. Their basic idea is that there can be multiple possible worlds which are maximally accurate in their representation of reality; this captures what it is for reality to be vague (in itself).

They, in contrast to Akiba, work with a theory of ersatz possible worlds, possible worlds being 'abstract objects which represent (classically complete) ways the world might be'.¹⁹⁸ They do not take a stand on precisely what these abstract objects are; anything representational will do the job. As on any ersatz possible worlds theory, it is taken to be

¹⁹⁸ Elizabeth Barnes and J. Robert G. Williams, 'A Theory Of Metaphysical Indeterminacy', in *Oxford Studies in Metaphysics: Volume 6* eds. Karen Bennett and Dean W. Zimmerman (Oxford: Oxford University Press, 2011), p. 114.

the case that the actual world is the possible world which represents reality as it is (to use Barnes and Williams' terminology, the maximally accurate world is *actualized*).

Barnes and Williams then introduce the notion of a 'precisificationally possible' world; this is a world 'that does not determinately misrepresent reality'.¹⁹⁹ If reality is fully determinate – if there is no metaphysical vagueness – then there will be just one precisificationally possible world, the world which represents reality precisely as it is.

However, if it is the case that there is metaphysical vagueness in reality, there will be more than one precisificationally possible world; more than one world which fails to determinately misrepresent reality, and is 'a candidate to be actualized' (I prefer Williams' earlier formulation, that there are 'multiple actualities').²⁰⁰ To illustrate with one of Barnes and Williams' examples: consider Fred, 'a foetus at an intermediate stage of development'.²⁰¹ There will be one ersatz possible world which represents Fred as instantiating the property of *being a person*; there will be another that represents Fred as not instantiating that property. Neither world is determinately correct, but neither determinately misrepresents reality, either; they are both precisificationally possible candidates for actualization. As Barnes and Williams put it:

If it is fundamentally unsettled whether p , there are two candidate representations for actualization—the abstract world which represents that p , and the abstract world that represents that $\neg p$. Neither of these are determinately correct, but neither is determinately incorrect, because in reality it's simply unsettled whether p or rather $\neg p$ obtains. If there is fundamental unsettledness in the actual world, then there will be no determinately correct way of representing how things are in reality.²⁰²

At this stage, we can bring in the familiar supervenient strategies for analysing metaphysically vague states of affairs. It is indeterminate whether Fred is a person, because there are precisificationally possible worlds in which he is, and precisificationally possible worlds in which he isn't. It is determinate that Fred is an organic being, because he is in every precisificationally possible world. To the question

¹⁹⁹ Ibid., p. 115.

²⁰⁰ Barnes and Williams, 'A Theory Of Metaphysical Indeterminacy', p. 115; J. Robert G. Williams, 'Multiple Actualities and Ontically Vague Identity', *The Philosophical Quarterly*, 58.230 (Jan., 2008), p. 135.

²⁰¹ Barnes and Williams, 'A Theory Of Metaphysical Indeterminacy', p. 114.

²⁰² Ibid., p. 115.

‘what exactly is metaphysical vagueness?’, Barnes and Williams reply, ‘indeterminacy over which maximally precisified possible world is actualized’. The indeterminacy settles in the lack of a one-to-one correspondence between the actual world and one of the maximally precisified possible worlds.

Barnes and Williams are careful, as was Akiba, to preserve classical logic. Each of the ‘worlds in the space of precisifications are themselves maximal and classical’; thus $p \vee \neg p$ will always be determinate, since it is the case in every precisificationally possible world.²⁰³ Bivalence is also secured by taking truth to be what is the case at the actualized world: in cases of metaphysical vagueness, there are multiple candidates for actualization, but on each of them, $p \vee \neg p$ is the case. In a sense, ‘*there is a precise way that things are* — so long as by ‘precise’ one means that for every p , either p or $\neg p$, and either p is True or $\neg p$ is’; for Barnes and Williams, this is compatible with it being ‘primitively indeterminate which precise way things are’.²⁰⁴

More formally, Barnes and Williams define truth simpliciter as truth at the intended model, where the intended model is the combination of the set of ersatz possible worlds, the set of all objects whatsoever, the interpretation of the language, and the designated actual world element. As stated, in cases of vagueness there will not actually be a member of the set of ersatz possible worlds which accurately represents the concrete world and is therefore suitable to determinately be the designated actual world. Nevertheless, Barnes and Williams say we are licensed to talk about ‘the intended model’:

For according to our guiding conception, there's exactly one world that depicts reality correctly, and adding this world to our space, domain, and interpretation will give a unique, intended model. But as it is indeterminate which world that is, it will be indeterminate which model is the intended one.²⁰⁵

Bivalence is thus delivered in the usual supervaluationist manner, with the disjunction ‘ p or not- p ’ being true simpliciter for every p , even though in cases of vagueness, both p and not- p will, on their own, be assigned indeterminate truth values; for even though we

²⁰³ Ibid., p. 116.

²⁰⁴ Ibid.; emphasis in original.

²⁰⁵ Ibid., p. 126.

are allowed to speak of ‘the intended model’, due to the indeterminacy over which is the actual world element of the intended model, both p and not- p will neither be determinately true nor false on the intended model, and therefore neither true nor false simpliciter.

Barnes and Williams say that their model can be usefully applied to instances of vagueness such as parthood, survival, identity, and the open future. It might seem that this lightweight, ersatzist supervaluational model is just what we have been looking for. Akiba, though, disagrees. He argues that the account which Barnes and Williams have presented is not, in fact, metaphysical in nature; moreover, any semantic supervaluationist ought to accept it. He begins by focusing on the fact that although Barnes and Williams take precisifications to be possible worlds, these possible worlds are abstract representations, rather than concrete. He then gives various definitions derived from the notions of actualization and indeterminacy, which will be accepted by all. Finally, he shows that, given a *de dicto* interpretation of ‘it is not determinate that p in reality’ (‘It is not determinate that p is true in reality’) which the semantic supervaluationist accepts, one must accept Barnes and Williams’ main claim, that if it is metaphysically indeterminate whether p , ‘then it is indeterminate whether a ‘ p ’ world is the actualized world or a ‘not p ’ world is the actualized world’.²⁰⁶

The technical details are not important; the key point from Akiba is that a linguistic ersatz possible world, ‘i.e., a maximal consistent set of sentences, is virtually no different from a complete set of interpretations that assigns truth values T and F to those sentences’.²⁰⁷ Once this is seen, it is understood that Barnes and Williams are essentially saying that if there is metaphysical indeterminacy, then there will be indeterminacy of interpretation. The semanticist would be perfectly willing to accept this, while still holding that there is in fact no metaphysical vagueness.

Akiba takes this to be a major criticism of Barnes and Williams’ work; he believes it shows that the account ‘is not an intelligible ontic theory; it is either a semantic theory

²⁰⁶ Ken Akiba, ‘How Barnes and Williams have failed to present an intelligible ontic theory of vagueness’, *Analysis*, 75.4 (Oct., 2015), pp. 569-570.

²⁰⁷ *Ibid.*, p. 572. Barnes and Williams, of course, did not commit to linguistic ersatz worlds; even if we take possible worlds to be non-linguistic, though, they can be translated to a set of sentences due to their essentially representational nature, and the point stands.

in disguise or at best an ontic theory not adequately characterized'.²⁰⁸ I believe that while Akiba's disjunction is accurate, Barnes and Williams would (reluctantly) agree with it. This is because Barnes and Williams do not intend for their ersatzist supervenient strategy to be an analysis of metaphysical vagueness. Williams, in an earlier paper, notes that 'my intention is not to reduce ontic vagueness, but only to develop a framework for theorizing about it'.²⁰⁹ Barnes and Williams together state it explicitly more than once: 'we will not be attempting to offer any kind of *reduction* or *analysis* of metaphysical indeterminacy';²¹⁰ 'obviously we're not attempting to (conceptually or metaphysically) "analyse" the notion of metaphysical indeterminacy... we conceive of the above as offering a way to show someone how to *work with* the notion'.²¹¹ When discussing the semantics and logic that they think best fits their account, they admit that 'even one completely sceptical of the coherence of the notion of m-indeterminacy [metaphysical vagueness] will be able to follow much of the discussion', as their 'notion of logical consequence, and of truth-on-a-model, never appeals to indeterminacy as such'.²¹²

Taking this account, then, Akiba's criticism loses much, if not all, of its force. Barnes and Williams are offering a translation, you might say, from metaphysical vagueness talk, which many profess to find unintelligible or incoherent, to ersatz modal talk, which is widely accepted. It is no problem for them if semantic supervenientists could (and should) accept their translation as an actual theory, for they were not attempting to offer anything heavily metaphysical as an alternative.

However, there is still a problem lying in wait for Barnes and Williams. Have they not attempted to analyse metaphysical indeterminacy because it was not necessary for their purposes, or because they believe it cannot be done? It is clear that they choose the latter: 'when p is metaphysically indeterminate, there are two possible (exhaustive,

²⁰⁸ Ibid.

²⁰⁹ Williams, 'Multiple Actualities and Ontically Vague Identity', p. 149.

²¹⁰ Barnes and Williams, 'A Theory Of Metaphysical Indeterminacy', p. 105.

²¹¹ Ibid., p. 118.

²¹² Ibid., p. 123.

exclusive) states of affairs... and it is simply unsettled which in fact obtains. No further explication is possible or needed.²¹³ This may be unsatisfactory for many.

At this point it is worth noting another feature of Barnes and Williams' account which may be taken as a positive by many, but to a true believer in genuinely metaphysical vagueness could be a significant drawback. It is linked to their ability to retain classical logic (at least in some sense), due to each possible world being maximally precisified and classical. In Barnes' own words, on this account 'indeterminacy does *not* entail imprecision'; in fact it is 'perfectly determinate that everything is precise, but... it's indeterminate which precise way things are'.²¹⁴ This seems to be a slightly odd thing for a proponent of metaphysical vagueness to say, as one of the driving intuitions behind such vagueness is exactly that everything is *not* precise. As Michael Tye puts it, 'common sense has it that the world contains countries, mountains, deserts, and islands, for example, and these items certainly do not seem to be perfectly precise'.²¹⁵ While it is easy to appreciate that Barnes and Williams' approach is fairly streamlined, and allows them to proceed along well-trodden supervenient lines without major revisions of classical logic, it is hard to lose the suspicion that their account, at heart, explains *away* metaphysical vagueness, rather than explains it.

In fact, they themselves admit to this in so many words:

When *p* is metaphysically indeterminate, there are two possible (exhaustive, exclusive) states of affairs—the state of affairs that *p* and the state of affairs that not-*p*—and it is simply unsettled which in fact obtains. No further explication is possible or needed. A primitivist about indeterminacy might take a different line: claiming that there is a tripartite (exhaustive, exclusive) division amongst states of affairs: the state of affairs that *p*, the state of affairs that $\neg p$, and, incompatible with either, the state of affairs of *p* being indeterminate. This is not the conception to be pursued here, and leads to quite different pictures of primitive indeterminacy.²¹⁶

They explicitly do not allow that it is possible for there to be a state of affairs where *p* is indeterminate; their taxonomy includes determinate states of affairs, and it is

²¹³ Ibid, pp. 113-114. Williams, in an earlier paper, thinks that 'one attractive "null" option for a constitutive account of vagueness would claim that this sort of operator is a metaphysical primitive': Williams, 'Multiple Actualities and Ontically Vague Identity', p. 149.

²¹⁴ Elizabeth Barnes, 'Ontic Vagueness: A Guide for the Perplexed', *Noûs*, 44.4 (2010), p. 622.

²¹⁵ Tye, 'Vague Objects', p. 535.

²¹⁶ Barnes and Williams, 'A Theory Of Metaphysical Indeterminacy', pp. 113-114.

only unsettled as to which is actualized. For Barnes and Williams, the world at base is precise, but can appear vague (or indeterminate, or unsettled) from one point of view.

It seems clear that any attempted supervenient treatment of metaphysical vagueness will be stuck on the following dilemma: be truly metaphysically robust, à la Akiba, and incur heavy metaphysical commitments and objections, or be metaphysically lightweight, à la Barnes and Williams, to the point where the account is potentially trivial, and acceptable even to those who deny the existence of metaphysical vagueness.

This is a more general issue with supervenient type theories, whether they be metaphysically heavyweight or lightweight; they result in general statements being made true even though there is no specific instance which can be pointed at as the one which makes the general statement true. This was mentioned earlier in section 1.1 as an issue for linguistic supervenientism, labelled as ‘truth-value shifts’. For instance, the statement ‘Mount Everest has a sharp boundary’ will be true in every precisified world, due to the very fact that they are precisified. Therefore, on a supervenient system, the statement is simply true. However, there is no one sharp boundary which makes the statement true in every world; the truthmaker shifts from cut-off to cut-off in each world. While this is seen by many to be problematic and unappealing (how can it really be true that there is a sharp boundary if no specific sharp boundary can be pointed to?), the supervenientist can console themselves that the truth of the general statement is something they wish to hold, generally for the purposes of maintaining some classical elements in their theory.

But to bring out a specific issue for metaphysical supervenientists, consider the following statement: ‘There is no indeterminacy in the world’. Presumably Barnes and Williams want this statement to be false, as they are arguing for the existence (or at least the possibility) of metaphysical indeterminacy. However, when we analyse its truth value in the possible worlds, we quickly see that it is true at every world, due to the fact that every world is maximally precisified. Thus, when we consider the truth of the statement at the actual world and calculate whether it is true in the ‘intended model’, we come to the conclusion that even if it is indeterminate as to which exactly is the

intended model, the statement will be true in any case, and thus is true simpliciter. The model, which supposedly allows us to talk about metaphysical vagueness, makes true a statement denying that any such vagueness exists.

As such, it may fairly be concluded that supervaluationism is not a promising strategy for anyone wishing to theorise about metaphysical vagueness, although it may be a useful paraphrase, as Barnes and Williams intend it to be. Let's now look at another 'general' account of metaphysical vagueness which utilises a different strategy.

4.3.3 Wilson

Jessica Wilson rejects what she terms the 'meta-level' approach to metaphysical vagueness; those theories on which metaphysical vagueness is a matter of it being '*indeterminate* which *determinate* SOA [state of affairs] obtains'.²¹⁷ This characterisation includes the supervaluationist accounts of Barnes and Williams and Akiba discussed above, as well as that of Barnes and Cameron; they may define what states of affairs are slightly differently, but the key point is that they are determinate, or precise, and the indeterminacy resides in the question of which of them are actually obtaining in the world.²¹⁸ Wilson makes the following observation of such accounts, which I find myself in agreement with:

Finally, metaphysical supervaluationist accounts are notoriously hard to 'grok': it is clear how (semantic) indeterminacy might reflect our having not yet decided how to use our language, but what would it be for the world to be undecided about, for example, what boundary a given macro-object has? There have been valiant attempts to make sense of metaphysical supervaluationism—as reflecting, e.g. that it may be '*indeterminate* which world is actualized', or that there may be multiple '*actual*' worlds. But even if such conceptions are coherent, I think I speak for many in saying that these accounts of MI occupy a metaphysically tenuous region of logical space.²¹⁹

²¹⁷ Jessica Wilson, 'Are There Indeterminate States Of Affairs? Yes', in *Current Controversies in Metaphysics* ed. Elizabeth Barnes (Abingdon: Routledge, 2017), p. 105 (emphasis in original).

²¹⁸ See Elizabeth Barnes and Ross Cameron, 'Are There Indeterminate States of Affairs? No', in *Current Controversies in Metaphysics* ed. Elizabeth Barnes (Abingdon: Routledge, 2017), pp. 120-132.

²¹⁹ Jessica Wilson, 'A Determinable-Based Account of Metaphysical Indeterminacy', *Inquiry*, 56.4 (2013), p. 364.

Wilson instead favours what she calls an ‘object-level’ approach to metaphysical vagueness; this is one on which it is ‘*determinate* (or just plain true) that an *indeterminate* (imprecise) SOA obtains’.²²⁰ The distinction between meta-level and object-level is intended to reflect a difference in opinion between where the metaphysical vagueness truly resides; for Barnes, Cameron, Williams, and Akiba, it is located in the meta-structure, in the relationships between the concretely existing world and the world of states of affairs, whereas Wilson locates the vagueness in the objects themselves.

For Wilson, an indeterminate state of affairs is ‘one whose constitutive object has a determinable property, but no unique determinant of that determinable’.²²¹ She explains that she makes use of determinable properties because they are distinctively unspecific properties, which by their nature are irreducibly imprecise, and only admit of precisification by specific determinate properties underneath them. This is taken to be good basis for providing an understanding of metaphysical vagueness and its fuzzy nature.

An initial response, which Wilson anticipates, is that determinables are generally taken to be ‘either eliminable, reducible to determinates, or (even if existing and irreducible) in any case less fundamental than their associated determinates’.²²² These arguments usually proceed along the lines of saying that determinates can already ‘do all the work’ that determinables would be involved in, whether it be causal work, or the grounding of facts, or some other metaphysical job; therefore, it is profligate to posit determinables as well as determinates. On this view, determinables are seen as something like an abstraction from groups of determinates, rather than a genuinely existing metaphysical category. Wilson rebuts this by arguing that when such powers are exerted, the determinable is exercising one and the same power wielded by its associated determinate in that instance, and therefore there is no duplication or overdetermination. There is a rejoinder available that, if it is indeed one and the same power, then determinables aren’t actually making a difference, in violation of

²²⁰ Ibid., p. 360 (emphasis in original).

²²¹ Ibid.

²²² Ibid., p. 365.

Armstrong's 'Eleatic Principle' (so named by Graham Oddie after the Eleatic stranger in Plato's *Sophist*)²²³ that 'Everything that exists makes a difference to the causal powers of something'.²²⁴ Wilson, however, offers that determinables 'make a difference' in that they have a distinctive set of powers to that of any of their determinates, and therefore determinables extrinsically contribute to a difference in powers.

The more serious objection is that of fundamentality; that 'once the specific/determinate facts are fixed, all the more general/determinable facts are thereby fixed, with the latter being something like ontologically real abstractions from the former'.²²⁵ Wilson argues that this supposition is based on a tacit premise that a fundamental base has only to ground non-modal facts at a world. This premise, though, is reasonably denied; modal facts about entities in a world are part of that world. And, Wilson, says, determinables are a key part of grounding modal facts; they ground facts such as that an instance of a particular shade of red might have been a slightly different shade. This is enough to deny the suggestion that determinables are non-fundamental, and ultimately eliminable.

We can now move on Wilson's account of how determinables figure in her account of metaphysical vagueness. As we have seen, she says that a state of affairs is metaphysically vague if it constitutively involves an entity which has a determinable property, but does not have a unique determinate of that property.²²⁶ The first question to answer is how can it be the case that something has a determinable property but no determinate; traditionally, as in the discussion above, it is assumed that when a determinable is present, so is (one and only one) relevant determinate, which instantiates the determinable. Wilson offers the example of the iridescent throat feathers of a hummingbird as an everyday counterexample; the feathers may appear different colours to observers viewing the bird from different angles at the same time. Plausibly, then, this is a case where the feathers are in possession of the determinable

²²³ Graham Oddie, 'Armstrong on the Eleatic Principle and Abstract Entities', *Philosophical Studies*, 41.2 (Mar., 1982), p. 286.

²²⁴ David Armstrong, *A World of States of Affairs* (Cambridge: Cambridge University Press, 1997), p. 41.

²²⁵ Jessica Wilson, 'Fundamental Determinables', *Philosopher's Imprint*, 12.4 (Feb., 2012), p. 11.

²²⁶ This is actually slightly simplified and removes some of the caveats in Wilson's formulation, as they are not germane to the general discussion here. See Wilson, 'A Determinable-Based Account of Metaphysical Indeterminacy', p. 366.

‘coloured’, while possessing no unique determinate of it; they may be thought of as possessing multiple determinates, each relativized to the individual observers.²²⁷

Moving on to the more familiar case of Everest and its vague boundaries, Wilson suggests that such cases be handled by positing that Everest is in possession of a determinable boundary property, while not having any determinate boundary property under it. What is a determinable boundary property? Wilson does not go into massive amounts of detail, but refers to macro-object boundary properties which seem to consist in simply ‘having a boundary’. What is a determinate boundary property? It is a micro-boundary property, that is a boundary property based on an aggregate of the micro-objects which are candidates for making up Everest; these can be seen as analogous to the multiple possible sets of matter which are candidates which we have encountered before. On this view, then, Everest has a determinable boundary property (the property of having a boundary), while not having a micro-boundary property (a specific and precise micro-boundary determined by any aggregate of micro-objects).

Wilson says this accommodates intuitions about metaphysical vagueness better than meta-level accounts; while they (in different ways) say it is indeterminate which precise boundary Everest has, Wilson’s theory can just state that Everest has no precise boundary (in the form of a determinate boundary property), even though it does have a boundary (a determinable boundary property). It also heads off Evans-type arguments involving indeterminate identity, as an object with no determinate boundary property will simply be non-identical with any object which has such a determinate. In a similar way, problems of the many never arise, as ‘there is just one mountain (table, statue) there: the one with the determinable boundary’.²²⁸

Overall it is an elegant and coherent account of what it means for something to be metaphysically vague. However, I believe there are certain weaknesses. There is one aspect of the theory which Wilson herself acknowledges, and which may not be a weakness, dependent on your viewpoint. It is that on this account, in a sense vagueness is not fundamental; Wilson is ‘at least weakly reductive’, in that she defines

²²⁷ Ibid., p. 367.

²²⁸ Ibid., p. 378.

metaphysical vagueness 'in terms of (a certain pattern of possession of) determinables and determinates, as opposed to taking this to be a primitive phenomenon'.²²⁹ Not everyone will find this a drawback, but as I will argue in the next chapter, I think it makes most sense, and respects intuitions best, to treat metaphysical vagueness as a primitive and fundamental part of metaphysics. It feels somewhat underwhelming to have an account which says that vague objects are just like precise objects, the only difference being in some respects they may lack a determinate property in some domain. I acknowledge this is not a particularly formal argument, though.

A more searching critique is that the account appears to rest on a somewhat idiosyncratic understanding of the nature of determinables and determinates. Wilson's arguments against determinables being reducible to, or less fundamental than, determinates are all well and good. But they only seem to have traction on a framework where either notion is seen to be part of our fundamental metaphysical description of the world, and I am not sure I agree with this framework. Are determinable and determinate not both abstractions, ways of categorising macro-level properties seen from a human-level perspective, which are all reducible to more micro-level properties?

For instance, take Wilson's feather case. She takes it to show that the feathers have the determinable of 'being coloured', while holding more than one determinate colour property. For one, I would not think it part of the fundamental description of the world to say that a feather has a specific shade of colour, any more than I would to say it is 'coloured' in general. Both are ways of describing, or organising, the way our perception engages with the feathers in different environments; and they both arise from certain facts about the physical and chemical makeup of those feathers (the surface molecules being such that they reflect different wavelengths of light differently when hit at different angles, for instance – I'm afraid I'm no physicist or zoologist so the specifics of this may be entirely wrong). That is to say, I am not convinced that either determinate or determinable properties are the kind of things that should be used in an account of metaphysical vagueness.

²²⁹ Ibid., p. 382.

The final objection is that, even admitting that determinables and determinates are admissible in such an account, I think that metaphysical vagueness is very widespread. It is hard to think of a macro-level object which does not appear to be vague in at least some way. In which case, almost every object will be the holder of many determinable properties but very few determinate properties. In fact, it may well be possible to build sorites-type series out of almost any property which is capable of holding a determinable-determinate distinction, to the extent that, if Wilson is right, determinates might be an endangered species, with only more or less specific determinables taking their place.

On the other hand, due to what I believe is the widespread nature of metaphysical vagueness, there may be instances which are not covered by Wilson's theory of determinables. As Barnes and Cameron point out, it is not particularly plausible to say that in a candidate instance of existential vagueness (where it is vague whether an object A exists or not), this could be due to A having a determinable property but no determinate. There is no obvious determinable which has existence and non-existence as determinates, and even if there were 'in saying that A has the determinable, we are presupposing the existence of A, and it's not even settled that there is such a thing [as A]'.²³⁰ Similar problems occur for vagueness in identity. Wilson is free to reject the possibility of existential vagueness or vagueness in identity, but it is at least a *prima facie* drawback that her supposedly general account of metaphysical vagueness commits her to the denial without further argument.

4.4 Conclusion

We have not been able to find a fully convincing theory of metaphysical vagueness from the options already proposed, then. Smith, Morreau, Tye, and McKinnon's specific theories each had some merits, but were lacking in other ways, and fell short of offering an account of widespread, general metaphysical vagueness. Akiba and Barnes and Williams fared better on that score, both proposing wide-ranging accounts which could accommodate the varied instances of metaphysical vagueness seen across the actual

²³⁰ Barnes and Cameron, 'Are There Indeterminate States of Affairs? No', p. 129.

world. However, both were meta-level accounts built on supervenational foundations, and as such failed to respect the fundamentally vague nature of the world. They both, in their different ways, reduced metaphysical vagueness to interactions between fundamentally precise entities, and were therefore unsatisfactory. Wilson offered a genuine object-level general theory of metaphysical vagueness, positing that it is the result of entities having determinable properties but no unique determinants. While the object-level nature of the account was encouraging, the reliance on determinables and determinants left it vulnerable to a number of objections, and made it hard to straightforwardly accept.

The examination of these theories has taught us some important lessons, though. We are looking for a general theory which can handle multiple types of metaphysical vagueness. It should be an object-level theory, that is one that locates the vagueness genuinely in the world itself, rather than in indeterminacy over which precise state of affairs the world instantiates. It should also not rely on controversial interpretations of other notions, and should be able to be accepted without major revisions of other theoretical commitments. I hope to make a start on building such an account in the next chapter.

5. My account of metaphysical vagueness

Let us now take stock of where we are. We have looked at what the phenomena of vagueness are; the various ways in which vagueness presents itself, and which of those should count as genuine cases of vagueness. In doing so, we have clarified what we actually mean, and consequently what we are actually talking about, when we discuss vagueness. We have offered *prima facie* reasons for thinking that vagueness really might be metaphysical (rather than merely semantic, or merely epistemic). We have also explored the leading semantic and epistemic theories of vagueness, which have been the object of much discussion and propagation; and have discussed the problems each of them face. We have also discussed some previously offered metaphysical theories of vagueness, but found none to be totally convincing. It is therefore now time to put one's money where one's mouth is, so to speak, and to offer a positive theory which will hopefully prove to be more successful than the metaphysical and ontic theories already discussed.

My view is an amalgamation and extension of the views of Wilson and Tye, taking what I see to be the most advantageous parts of each in order to amalgamate a general theory of genuine metaphysical vagueness. I agree with Wilson that we should take an object-level approach to metaphysical vagueness, and hold that it is (determinately) true that vague states of affairs obtain. That is, the world is such that some states of affairs are vague, and some of these states of affairs do in fact obtain. Thus my view is a 'Third-Way' view, to use Barnes and Cameron's terminology, in contrast to 'Unsettledness' views such as theirs, where there are only precise states of affairs which indeterminately obtain.²³¹ I do not agree with Wilson, though, that a state of affairs is vague if it contains the existence of a determinable with no unique determinant.

For me, a vague state of affairs will be one which contains a vague constituent; the vague constituent may be a vague object, a vague property, a vague set, or a vague relation. I will follow Tye in his definition of what it means for these constituents to be

²³¹ Barnes and Cameron, 'Are There Indeterminate States of Affairs? No', p. 122.

vague; for instance, a concrete object O is vague if and only if '(a) O has borderline spatio-temporal parts and (b) there is no determinate fact of the matter about whether there are objects that are neither parts, borderline parts, nor non-parts of O'. Similar definitions will be given later for other sorts of constituents.

My view, then, is that genuine metaphysical vagueness exists because vague states of affairs determinately obtain; that is, states of affairs which contain vague constituents determinately obtain, independent of thought and representation. Before I get into the details, I will first start by laying out what I take the desiderata to be for an effective theory of metaphysical vagueness, these having been informed by the previous discussions of other theories. These desiderata are what will then guide me in laying out my positive proposal.

5.1 What are the desiderata?

I will set out three general considerations that any philosophical theory must bear in mind; the list does not aim to be exhaustive, but they are key considerations which will guide my discussion. First, a good theory must be explanatorily useful; it must actually provide value in explaining some phenomena or other which are proving puzzling, or in need of systematic explication. That is, it is not enough for a theory to be merely internally consistent. It must offer a way of explaining the world (for a theory of metaphysics) which allows the reader to think that they have gained a greater understanding of how the world actually is.

A good theory should also take the issues under consideration seriously, which entails one of two things. Either the theory will accept that the apparent phenomena are real, and offer an account of what they consist in. Alternatively, the theory may deny that the phenomena under discussion are truly occurring (in some sense); in which case the theory must still take the issue seriously by offering a reason for denying the phenomena, and an 'error theory' which can explain why they appear to occur. A simple denial is merely a sidestepping of the issue.

The final general desideratum is that a good theory should not create needless complexity, in a few different dimensions. One is internal simplicity; a theory which can be explained concisely, without requiring the use of an excessive number of moving parts will, all else being equal, be more attractive than a convoluted one which implies ad hoc modifications. Another is coherence with other, well-established systems and frameworks; the complexity introduced by needing to modify various other theories will be a disbenefit to an account. Finally, if offered the choice between two theories which offer to explain the same phenomena to the same level, we will choose the one which fits with our wider commitments and beliefs, rather than the one which entails an overturning of them (for instance, by requiring us to admit that there are no composite objects).

Now, we move on to the desiderata which are more specific to a theory of metaphysical vagueness. Firstly, and most importantly, is that a theory of metaphysical vagueness must explain why the various phenomena which we have observed earlier exhibit the epistemic tolerance which we have agreed is representative of vagueness. That is, they must answer the basic question of why vagueness happens. It is key to note that this need not necessarily be a unified answer; plausibly, vagueness arises for different reasons in different situations, and therefore there are different underlying causes for the surface instances of tolerance. A non-unified account, though, must offer some explanation as to why these separate causes happen to produce a similar epistemic effect.

Secondly, a theory of metaphysical vagueness must explain why some of the instances of metaphysical vagueness interact with each other in the manner that they do. For instance, an account of compositional vagueness could also include an explanation of its relation to existential vagueness (or an explanation of why existence cannot in fact be vague); when does one lead to the other, and why is that the case. Of course, the theory might state that not all the apparent connections are actual, in which case an explanation would be needed of why we have the intuition that they are.

Thirdly, and relatedly, a theory must explain why metaphysical vagueness can lead to, or at least be connected to, instances of semantic and epistemic vagueness (that is, ignorance whether something is the case or not). This explanation could take many

forms, from showing how metaphysical vagueness leads to semantic or epistemic vagueness, to an account of how certain types of situations give rise to more than one form of vagueness.

Fourthly, as we argued in the previous chapter, it should take the idea of genuine metaphysical vagueness seriously. To put it another way, it should not quickly reduce metaphysical vagueness out of existence, and offer a view of the world which is inherently precise (albeit with a smattering of vague description on top).

5.2 Vague states of affairs

The first aspect of my account to be clear on is the statement above that genuine metaphysical vagueness consists in the definite obtaining of vague state of affairs (for instance, the state of affairs that contains the cat Tibbles is such that it is vague whether a certain hair is part of the cat ‘Tibbles’). We have already encountered a view with a similar background framework; that of Jessica Wilson, discussed in section 4.3.3. Wilson labelled her view an ‘object-level’ approach to metaphysical vagueness, on which it is ‘*determinate* – or just plain true – that an *indeterminate* SOA [state of affairs] obtains’.²³² This is opposed to the ‘meta-level’ views of Barnes and Williams or Akiba, on which it is ‘*indeterminate* which *determinate* SOA obtains’.²³³ The object/meta-level distinction here captures the fact that for Wilson (and myself) vagueness is located within individual states of affairs themselves; for Barnes and Williams or Akiba, vagueness only arises when we ascend a level and evaluate across multiple (determinate) states of affairs.

Barnes and Cameron themselves discuss this difference, although they use the terms ‘Unsettledness Views’ and ‘Third-Way Views’. On the Unsettledness View, ‘indeterminacy entering the world at the level of what states of affairs obtain or fail to obtain’:

The believer in the Unsettledness View and the epistemicist can agree that every possible state of affairs is a perfectly normal state of some ordinary things instantiating some familiar properties and standing in familiar

²³² Wilson, ‘Are There Indeterminate States Of Affairs? Yes’, p. 105 (emphasis in original).

²³³ *Ibid.*, emphasis in original.

relations. Their disagreement is not about what states of affairs could exist, or about the nature of such things; it is about whether the world can fail to settle which states obtain. Thus on the Unsettledness View there are no indeterminate states of affairs.²³⁴

On Third-Way views, however, it is the case that ‘indeterminacy arises because of the obtaining of a certain special kind of state of affairs – an indeterminate state’:

On the Third-Way View, as well as propositions making demands on the world for their truth, and for their falsity, propositions also make specific demands on the world for their indeterminacy. The defender of the Third-Way View is increasing the stock of possible states of affairs she countenances: she disagrees with the defender of the Unsettledness View and the epistemicist that the only states of affairs that might obtain are those of ordinary objects instantiating familiar properties and standing in familiar relations. Those ordinary states of affairs might be what are relevant to the demands of truth and falsity, but indeterminacy demands the obtaining of a special new kind of state of affairs: perhaps the state of an object indeterminately instantiating a familiar property, or perhaps the state of an object instantiating the non-familiar property of being indeterminately F. Either way, the indeterminacy is part of the very nature of the state of affairs. Thus on the Third-Way View there are indeterminate states of affairs.²³⁵

As can be seen, Barnes and William’s view, being as it is an Unsettledness View, cannot countenance the existence of vague states of affairs. I hope it will be clear from the preceding sections why I do not believe a meta-level view is convincing as an account of genuine metaphysical vagueness. It is my belief that in order to fully locate vagueness in the metaphysical nature of the world, we must accept that indeterminate (or vague) states of affairs actually obtain; I definitely do not wish to agree with the epistemicist that every possible state of affairs is ‘perfectly normal’ and precise.

This is another instance of my belief that an account of metaphysical vagueness should take the intuitions behind metaphysical vagueness seriously. In Tye’s words, ‘common sense has it that the world contains countries, mountains, deserts, and islands, for example, and these items certainly do not seem to be perfectly precise’.²³⁶ Wilson comments that this intuition seems to ‘naturally read as characterizing such objects as *determinately failing* to have precise boundaries, not as being such that it is

²³⁴ Barnes and Cameron, ‘Are There Indeterminate States of Affairs? No’, p. 122.

²³⁵ Ibid., p. 123.

²³⁶ Tye, ‘Vague Objects’, p. 535.

indeterminate *which* precise boundary they have'.²³⁷ In cases of vague objects, the intuition is that precise boundaries are simply not the sort of things that they can have; an account of them which relies on perfectly precise boundaries (occurring in precise states of affairs), then, seems to have missed the point.

At this point it is worth clarifying that Wilson denies that her view is in fact a Third-Way view. She claims that while she does countenance the existence of indeterminate states of affairs, this does not entail the indeterminacy of propositions. Propositions about the world will either conform or not to its indeterminate nature, in which case they will simply be true or false: Wilson's account 'does not posit a third "indeterminate" category of truth value, or an indeterminacy operator on propositions, or any other piece of semantic machinery that would suggest that propositions are ever anything other than true or false'.²³⁸ Straightforwardly, any statement attributing a precise boundary to Everest will be false: 'it is not indeterminate that (or whether) Mount Everest has precise boundary #326 – again, for any (unrelativized) precise boundary, it is false, not indeterminate, that (or whether) Mount Everest has that precise boundary'.²³⁹ More controversially, Wilson implies that pairs of statements such as 'Rock R is part of Everest' and 'Rock R is not part of Everest' (where R is located in Everest's borderline region) will both be false, due to the fact that neither conform to the indeterminate nature of the world: Wilson advocates a system of 'gappy' metaphysical indeterminacy.²⁴⁰

It seems to me that in doing so, Wilson has offended against the very intuition that she herself cited as counting against meta-level views. When common sense says that mountains appear to not have precise boundaries, this does not appear to line up with it being false that rocks in their borderline region are part of them. Being able to reply that the situation isn't as bad as it seems because it is also false that those same rocks are not part of the mountain is scant consolation. Why not take the intuition seriously, and agree that the statement 'Rock R is part of Everest' is vague? This is what I propose to do. I am, then, proposing an account of genuine metaphysical vagueness which is both

²³⁷ Wilson, 'Are There Indeterminate States Of Affairs? Yes', p. 115.

²³⁸ *Ibid.*, p. 110.

²³⁹ *Ibid.*

²⁴⁰ *Ibid.*, p. 111.

object-level and Third-Way. What, then, do the vague states of affairs which I am positing consist in?

5.2.1 What is a vague state of affairs?

I offer the following definition of a vague state of affairs:

VSOA: What it is for a state of affairs *S* to be metaphysically vague at time *t* is for *S* to constitutively involve a vague entity.

In order to explain what counts as a vague entity, we must first explore the account of vague objects given by Michael Tye. Tye begins from the standard intuition that there is no sharp line delimiting Mt. Everest, for example, and that Everest's boundaries are therefore fuzzy. After observing that there are clearly some molecules which are inside Everest (and therefore part of it), and clearly there are some molecules which are outside Everest (and therefore not part of it), he considers the question of whether there are any remaining molecules; that is, whether there are molecules which are neither inside nor outside Everest.

To suppose that it is true that this is the case is to postulate more categories of molecules than are demanded by our ordinary, everyday conception of Everest and hence to involve ourselves in gratuitous metaphysical complications. It is also to create the need to face a potentially endless series of such questions one after the other as new categories of molecules are admitted. On the other hand, to suppose that it is false that there are any remaining molecules is to admit that every molecule fits cleanly into one of the three categories so that there are sharp partitions between the molecules inside Everest, the molecules on the border, so to speak, and the molecules outside. And intuitively, pretheoretically it is not true that there are any sharp partitions here. What, I think, we should say, then, is that it is objectively indeterminate as to whether there are any remaining molecules.²⁴¹

From this intuition, Tye states that an ordinary concrete object *O* is vague if and only if '(a) *O* has borderline spatio-temporal parts and (b) there is no determinate fact of the

²⁴¹ Tye, 'Vague Objects', p. 535.

matter about whether there are objects that are neither parts, borderline parts, nor non-parts of O'.²⁴²

Part (a) of the above definition is fairly standard for an account of vague objects; if an object were to have no borderline parts, then it would be a precise matter as to whether any thing was a part of it, and the object would have precise, determinate boundaries. Part (b), however, is more interesting. This specifies that not only does Everest have a borderline region, but that this borderline region does not have sharp boundaries; essentially it is a requirement that in order for an object to count as vague for Tye, it must be vaguely vague. He admits that some philosophers may wish to classify an object as vague based only on condition (a), but says that for him, 'a material object which has borderline spatiotemporal parts but which is such that the lines dividing those parts from the matter clearly inside the object and the matter clearly outside are sharp... is not indeterminate enough to count as a vague object'.²⁴³ From this we can assume that Tye takes the 'no sharp boundaries' intuition seriously.

However, this is not an invitation to ascend to higher and higher levels of vagueness, as might be supposed. After all, if there is no sharp boundary between the rocks which are parts of Everest and those which are borderline parts of Everest, it is reasonable to conclude that there must be rocks which are located in a borderline region between definite parthood and borderline parthood; these could be said to be 'borderline²' parts of Everest. Of course, we will wish to deny a sharp boundary between parts and borderline² parts (or indeed between borderline² parts and borderline parts), in which case we will identify borderline³ parts, and so on. Tye denies such a climbing of levels, when answering the question of whether there are any molecules which do not fit into the category of part, borderline part, or non-part:

To suppose that it is true that this is the case is to postulate more categories of molecules than are demanded by our ordinary, everyday conception of Everest and hence to involve ourselves in gratuitous metaphysical complications. It is also to create the need to face a potentially endless series of such questions one after the other as new categories of molecules are admitted. On the other hand, to suppose that it is false that there are any remaining molecules is to admit that every molecule fits cleanly into one of

²⁴² Ibid, p. 536.

²⁴³ Ibid., p. 539.

the three categories so that there are sharp partitions between the molecules inside Everest, the molecules on the border, so to speak, and the molecules outside. And intuitively, pretheoretically it is not true that there are any sharp partitions here. What, I think, we should say, then, is that it is objectively indeterminate as to whether there are any remaining molecules.²⁴⁴

This is an acceptance of the vague nature of vagueness itself, and a way of avoiding the further paradoxes of higher-order vagueness which appear when attempting to examine borderline regions in precise detail. Some may find it unintuitive or unsatisfying, but I think it is a good example of being willing to take the phenomena of vagueness seriously and holistically.

Tye goes on to say that some abstract objects, such as sets and properties, can also be vague, and gives analogous conditions for them:

I classify a set S as vague, if, and only if, (a) it has borderline members and (b) there is no determinate fact of the matter about whether there are objects that are neither members, borderline members, nor non-members.²⁴⁵

An example of this would be the set of tall men. While there will be individuals who are determinately members of this set, and individuals who are determinately not members, there will also be individuals for whom it is borderline whether they are a member of the set. Again, Tye agrees that it is 'a mistake to assert that it is true that there are further men who are neither members, borderline members, nor non-members', while it is 'no less mistaken to assert that it is false that there are such men: intuitively it is not true that the dividing lines between the categories are sharp'.²⁴⁶

The conditions for properties are slightly modified:

I take a property P to be vague only if (a) it *could* have borderline instances and (b) there is no determinate fact of the matter about whether there *could* be objects that are neither instances, borderline instances, nor non-instances.²⁴⁷

The condition for properties is necessary but not sufficient due to the possibility of counter-examples to sufficiency such as the following. Imagine the (intuitively precise)

²⁴⁴ Ibid., p. 535.

²⁴⁵ Ibid., p. 536

²⁴⁶ Ibid.

²⁴⁷ Ibid. Emphasis in original.

property of being 2000 feet in height. Now imagine a vague object such that it is indeterminate whether it is 2000 feet in height. In this case, as Tye observes, both condition (a) and (b) will be satisfied, but we do not want to conclude that the property of being 2000 feet in height is vague. Tye notes that we can construct necessary and sufficient conditions if we restrict ourselves to properties of concrete objects: ‘P is vague if, and only if, (a) P could have as a borderline instance a concrete object that does not have borderline spatiotemporal parts, and (b) there is no determinate fact of the matter about whether there could be an object of this sort which is neither an instance, a borderline instance, nor a non-instance of P’.²⁴⁸ This formulation rules out the possibility of the vagueness arising from the object; that is, a property will only satisfy the conditions if it itself is vague.

The conditions for a property to be vague are also modalised (whether in the general form, or restricted to concrete objects) to take into account the fact that a vague property may not have any actualised instances, but still be intensionally vague.

Tye continues by applying the strong Kleene three-valued system of logic to vague statements. This system provides for three truth values; truth, falsity, and neither (Tye also refers to ‘neither’ as ‘indefinite’, and takes ‘indefinite’ to be technically not a truth value in itself, but merely the lack of either of the other two values). When only truth and falsity are present, the logical connectives behave classically, but formulae can also be indefinite if some of their constituent parts are. I will go further into the details of such a system in section 5.5.2, while arguing that, as implied by my account being a Third-Way view, we should take the third value to indeed be a separate value; that of vagueness.

5.2.2 Genuinely vague states of affairs

I wish to generalise Tye’s definition of vague objects in order to provide an explanation of what it means for a state affairs to be vague. As stated earlier, on my account a vague state of affairs is one which constitutively involves a vague entity. A vague entity is one

²⁴⁸ Ibid., p. 537.

which (a) exhibits instances of vague relations, and (b) where there is no determinate fact of the matter as to whether there are things which neither fulfil the relation in question, vaguely fulfil it, nor do not fulfil it.

Tye's definition of a vague object *O*, given above, shows that such objects are vague entities: there are things which are vague instances of the relation of 'being a part of' with *O*, and there is no determinate fact of the matter as to whether there are things which neither fulfil the 'being a part of' relation with *O*, vaguely fulfil it with *O*, nor do not fulfil it with *O*. It can easily be seen how the conditions for vague sets and properties also entail that these are vague entities. I should note, though, that my definition of a vague entity is not modalised; therefore it will not be enough for a state of affairs to qualify as vague for it to contain a vague property without any actual instances. In order for a state of affairs to be vague via a vague property, the property must have an actual vague instance.

My notion of a vague entity can encompass types beyond those mentioned by Tye, however. Consider another type of metaphysical vagueness mentioned earlier, that of compositional vagueness. In this case, we can construct a similar definition. A collection of objects *C* vaguely compose a larger object *O* if and only if (a) *C* and *O* are a vague instance of the composition relation, and (b) there is no determinate fact of the matter about whether there are objects which are neither composed by *C*, vaguely composed by *C*, nor not composed by *C*. As an example, consider a group of water vapour molecules scattered across an otherwise empty sky, not so close together as to obviously compose a cloud, but not so spread out as to obviously not compose a cloud either. In this case, the collection of molecules together are an example of vague composition of a cloud. Furthermore, part (b) is satisfied because there is no sharp boundary available whereby we could say that if the molecules were arranged slightly differently they would go from determinately composing a cloud to vaguely composing one, or equally sharply go from vaguely composing a cloud to determinately not composing one.

Or consider the genuinely vague relation that a vague object might stand in to a location. Assuming that an object is located at all of the locations where it has extension, we could have the situation where, in the borderline region of the object, it is

vaguely located there. On my schema, an object *O* is vaguely located at a place *L* if and only if (a) *O* is a vague instance of being located at *L* and (b) there is no determinate fact of the matter as to whether there are locations which *O* is neither located at, vaguely located at, or not located at. An example of this could be Everest, when we look at the rocks in the foothills which are (genuinely) vague parts of it. Is Everest located at the regions where those vague parts are located? It is vaguely located there. And there is no sharp boundary between the locations where Everest is determinately located, those where it is vaguely located, and those where it is determinately not located. As before, Everest is (genuinely) vaguely located at those regions.²⁴⁹

In each of these cases, if a state of affairs contains such a vague entity, it will be the case that the state of affairs is vague. If such a vague state of affairs (determinately) obtains, then the world will be genuinely metaphysically vague.

To return to the characterisation of vagueness arrived at in the introduction, that instances of vagueness exhibit epistemic tolerance with respect to a property, it can easily be seen how this aligns with the existence of genuinely vague states of affairs. The principle of epistemic tolerance was as follows:

F is *epistemically* tolerant with respect to Φ if and only if any small positive degree of change with respect to Φ is insufficient to make a known difference to the correctness of applying *F* to a particular case.²⁵⁰

Vague states of affairs will satisfy this condition, due to the existence of genuine vague cases, and the lack of a sharp boundary between the vague cases and determinate cases. Given that, a small degree of change with respect to the vague property will necessarily fail to make a known difference to the correctness of application, because a small degree of change cannot make a difference to the actual obtaining of the vague relation. For instance, in the case of Everest with its vague parts, taking *F* to be 'is part of Everest' and Φ to be the particular rock (or perhaps the location of the particular rock under examination), then *F* will be epistemically tolerant; a small degree of change as to which rock is under examination will be insufficient to make a known difference as to whether the rock is part of Everest, due to the existence of vague cases, and the lack of

²⁴⁹ This is an obvious example of the connections that can hold between different types of genuine metaphysical vagueness.

²⁵⁰ Greenough, 'Vagueness: A Minimal Theory', p. 258.

a sharp boundary between determinate cases and vague cases. That is, there is no small degree of change which marks a genuine change from it being correct to apply 'is a part of Everest' to one rock and then to its neighbour in the world itself; there is nothing metaphysically to ground a known difference in correctness of application.

It will be noted that my definition of a vague entity contains the notion of a 'vague instance', which is not further defined or explicated. This is because I take the idea of something's being vague to be irreducible to some other set of facts. To follow Barnes's distinction, I want to have vagueness itself be fundamental to our ideology:

Indeterminacy is itself fundamental just in case our fundamental ideology must involve indeterminacy: that is, if our fundamental ideology must include a primitive indeterminacy operator, or if our fundamental logic be degree theoretic, etc.²⁵¹

'Vague' is a primitive operator in my ideology; something just is a vague instance, or it's not, and this cannot be given an analysis or reduction. It is one of the bedrock notions which must be used to accurately describe the world (which is itself vague).

This is in contrast to Wilson (for whom vagueness can be reduced to the existence of a determinable without the instantiation of a unique determinant) and Barnes and Williams or Akiba (for whom vagueness can be reduced to what's true at a precisificational space, or the existence of multiple precise possible worlds which are all maximally accurate to the actual world). Barnes herself is sympathetic to this point of view; 'I'm sometimes tempted by the thought that a case for metaphysical indeterminacy... requires commitment to fundamentality of indeterminacy itself'.²⁵² She states of Akiba's account that 'this isn't a theory according to which the world really is (bang the table) indeterminate in the interesting, provocative sense that a theory of metaphysical indeterminacy should capture'.²⁵³

This 'interesting, provocative sense' is, I think, a reference to the idea I have argued for a number of times in this thesis; namely that, if possible, we should try to engage

²⁵¹ Elizabeth Barnes, 'Fundamental Indeterminacy', *Analytic Philosophy*, 55.4 (2014), p. 346.

²⁵² Ibid.

²⁵³ Ibid.

seriously with the intuitions that the world is vague, and resists precision, unless we find it impossible to maintain a theory which does so.

5.3 Truth-functional background

We must now explore how truth works in relation to vague states of affairs. I mentioned earlier that Tye favoured a Kleene three-valued system; I will too, but before getting into the details, we should discuss the background motivation for such a system.

Kripke developed an account according to which paradoxical sentences such as the Liar are neither true nor false; they are defective ‘in much the way that... sentences that contain category mistakes have been thought to be defective’.²⁵⁴ I will be following this approach, although importantly, on my view such sentences are not actually ‘defective’; they are accurately representing a vague state of affairs.²⁵⁵

It is important at this point to consider what we seek from a theory of logic. Traditionally, logic is taken to play a normative role; according to Frege, logic’s job is to ‘prescribe universally how one ought to think if one is to think at all’.²⁵⁶ A system of logic ‘should teach us how to reason well by showing us patterns of inference which are reliable’.²⁵⁷ It is intuitive, and surely correct, to see logic as normative for reasoning; as MacFarlane points out, ‘try saying how logic differs from geometry without mentioning thought or reasoning, and try saying how logic differs from psychology without mentioning norms’.²⁵⁸

But the question now arises, if we have more than one system of logic available to us, they may well give us conflicting norms to follow. How, then, are we to choose between them? Plausibly, by comparing how they instruct a thinker to proceed in situations which we are sure of pre-theoretically. We take patterns of inference which we already know to be reliable, and make sure that the putative system endorses them (or,

²⁵⁴ Vann McGee, *Truth, Vagueness, and Paradox* (Indianapolis: Hackett, 1990), p. 87.

²⁵⁵ *Ibid.*

²⁵⁶ Gottlob Frege, *Grundgesetze der Arithmetik* (Paderborn, Germany: Mentis, 2009), p. XV.

²⁵⁷ McGee, *Truth, Vagueness, and Paradox*, p. 95.

²⁵⁸ John MacFarlane, ‘In What Sense (If Any) Is Logic Normative for Thought’ (presented at the 2004 Central Division APA symposium on the normativity of logic), accessed at <<https://johnmacfarlane.net/normative.pdf>>, p. 1.

perhaps, does not disavow them). If a putative logical system did not license the inference from 'P and Q' to 'P', we would be unlikely to follow its other injunctions.

All this seems right enough. But which patterns of inference are so widely accepted pre-theoretically that they are able to stand as these standards against which systems of logic can be measured? As the saying goes, one person's modus ponens is another person's modus tollens. Probably there are no generally agreed upon patterns, and every proposed set will be incomplete for some and overfull for others. What we should do, then, is simply lay out how a given system treats all patterns of inference, and it will be up to each individual to decide whether they agree with its conclusions. This is what I shall attempt to do with the following system, after first giving it a full explanation.

5.3.1 What is a partial model?

In order to build the system which makes space for these truth-value-less sentences, we need to introduce the idea of partial interpretations. Informally, partial interpretations do as their name implies; they provide an interpretation of a language which is not complete. An example of Carnap's may illustrate the concept.

In the context of logical positivism, Carnap discussed the question of the observational content of dispositional terms such as 'solubility' or 'malleability'. What observation can provide truth-values to sentences containing these 'hypothetical' properties? We can observe items dissolving (or not), and being shaped (or not), but we cannot directly observe them 'being such that they would dissolve were they to be put in water'.

Therefore, Carnap settled on the partial definition of the term 'soluble' given by the following 'bilateral reduction sentence' (where the single arrow is used to represent the material conditional, and the double arrow is used to represent the material biconditional):

$$(\forall x)(x \text{ is placed into water} \rightarrow (x \text{ is soluble} \leftrightarrow x \text{ dissolves}))^{259}$$

²⁵⁹ Rudolf Carnap, 'Testability and Meaning', *Philosophy of Science*, 3.4 (Oct., 1936), pp. 440-443.

Everything that is placed into water and dissolves is soluble, and everything that is placed into water and does not dissolve is not soluble. What of the (small number of) things that are not placed into water at some point in their existence? Carnap leaves it undefined as to whether they are soluble or not; this falls some way short of the everyday definition of solubility. McGee remedies it somewhat by adding the condition that ‘If two bodies of the same substance are each placed in water, then either both will dissolve or neither will’. This allows us to conclude that a particular wooden cricket bat, while it may have thankfully managed to remain un-submerged in its lifetime, is insoluble, because other objects made of the same material have been placed into water without dissolving.

We can now place the above example into a slightly more formal context. Following McGee, let’s take a partial interpretation to be ‘a pair (U, Γ) , where U is an ordinary first-order structure and Γ is a first-order theory in a language extending the language of U . The language of U is fully interpreted, but the rest of the language is only partially interpreted.’²⁶⁰

The language U will include the terms ‘is placed into water’, ‘dissolves’, and ‘is the same substance as’. The theory Γ comprises the bilateral reduction sentence ‘ $(\forall x)(x \text{ is placed into water} \rightarrow (x \text{ is soluble} \leftrightarrow x \text{ dissolves}))$ ’, and McGee’s additional ‘substantial similarity’ sentence ‘ $(\forall x)(\forall y)(x \text{ is the same substance as } y \rightarrow (x \text{ is soluble} \leftrightarrow y \text{ is soluble}))$ ’. The sentence ‘our dry cricket bat is soluble’ comes out as definitely untrue under Γ , as wooden items have indeed proved their insolubility. However, a sentence regarding the solubility of a never-submerged item, made of a material which has never been submerged, will be neither definitely true nor definitely untrue under Γ ; the conditions for being assigned a truth value have simply not been met.

Above, I used the term ‘definitely true’. The formal definition of definite truth in a partial interpretation is as follows:

A sentence is *definitely true* under the partial interpretation (U, Γ) in the model theoretic sense $[(U, \Gamma) \models \phi]$ iff ϕ is true in every expansion of U to a

²⁶⁰ McGee, *Truth, Vagueness, and Paradox*, p. 149.

model of Γ . Φ is definitely untrue in the model theoretic sense iff $(U, \Gamma) \models \neg\Phi$.²⁶¹

An expansion of a model is constructed by adding new symbols to the language, with the reference of the symbols being specified, and all previously existing symbols retaining their previous reference. As McGee puts it, 'we might think of the expansions of U as representing the hypothetical future interpretation that we will get when the meanings of all the terms have been completely specified. Γ represents features of that hypothetical future model to which our present-day usage commits us.'²⁶² As you can imagine, not all sentences will retain their truth values in every expansion; something of the form 'all x s are P ', which may be true given the domain of x s in U , might become false when an expansion adds an x which is not- P . On the partial interpretation (U, Γ) , then, this sentence 'all x s are P ' will be neither definitely true nor definitely untrue.

5.3.2 A partially interpreted truth predicate

Let us now put the above comments on partial interpretations on a more formal footing.

McGee defines a partially-interpreted truth predicate as follows:

Given a countable first order language L and an acceptable structure U for L , we form the language $L+$ by adjoining the single new unary predicate 'Tr' to L . We expand U to a classical model (U, E) of $L+$ by picking a subset E of $[U]$, which is to be the extension of 'Tr'. We get a partial model $(U, (E, A))$ by picking two disjoint subsets E and A of $[U]$. The extension E is to consist of those things to which the predicate 'Tr' definitely applies, while the anti-extension A is to consist of those things to which the predicate 'Tr' definitely does not apply. There may be members of $[U]$ which are in neither E nor A ; for such things it remains unsettled whether 'Tr' applies.²⁶³

This is in line with Kripke's introduction of such a predicate in 'Outline of a Theory of Truth' (1975):

Suppose we extend L to a language D : by adding a monadic predicate $T(x)$ whose interpretation need only be partially defined. An interpretation of $T(x)$ is given by a "partial set" (S_1, S_2) , where S_1 , as we said above, is the

²⁶¹ Ibid., p. 150.

²⁶² Ibid.

²⁶³ Ibid, pp. 87-88.

extension of $T(x)$, $S2$ is the antiextension of $T(x)$, and $T(x)$ is undefined for entities outside $S1 \cup S2$.²⁶⁴

Kripke was making use of Kleene's strong three-valued logic schema, and while doing so pushed back against the claim, sometimes raised, that such a system introduces a third truth value, separate to truth and falsity. For Kripke, "'Undefined" is not an extra truth value',²⁶⁵ but simply the lack of application of either truth value. Kleene himself emphasised the same point: the undefined " u is thus not on a par with the other two t and f in our theory".²⁶⁶ Given that on my view vagueness genuinely exists and is part of both our ontology and ideology, I disagree and take the third value to positively represent instances of vagueness; this does not affect the technical framework but is relevant in the metaphysical interpretation of it, under which vague states of affairs determinately obtain, and provide the 'vague-makers' for propositions (as opposed to truth-makers, as previously discussed in section 5.2).

McGee argues that an important aspect of Kripke's 'Outline of a Theory of Truth' is the conclusion that what we require of a successful theory of truth is not actually that it satisfies Tarski's T-schema:

T-Schema: $[\phi]$ is true (in English) if and only if ϕ .

What we really need is that our theory of truth yields the related but different R1-4:

(R1) From $[\phi]$ one may infer $\{[\phi] \text{ is true}\}$

(R2) From $\{[\phi] \text{ is true}\}$ one may infer $[\phi]$

(R3) From $[\neg\phi]$ one may infer $\{[\phi] \text{ is not true}\}$

(R4) From $\{[\phi] \text{ is not true}\}$ one may infer $[\neg\phi]$ ²⁶⁷

This is because Kripke places a great emphasis on people's everyday use of the notion of truth. As Wittgenstein said, 'philosophical problems arise when language goes on holiday',²⁶⁸ so Kripke, in his attempt to build an adequate theory of truth, examines truth

²⁶⁴ Kripke, 'Outline of a Theory of Truth', *The Journal of Philosophy*, 72.19 (Nov. 1975), p. 702.

²⁶⁵ Ibid, p. 700 fn. 18.

²⁶⁶ Raymond Kleene, *Introduction to Metamathematics* (Amsterdam: North Holland Publishing Co., 1952), p. 333.

²⁶⁷ Vann McGee, 'Applying Kripke's Theory of Truth', *The Journal of Philosophy*, 86.10 (Oct. 1989), p. 530.

²⁶⁸ Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1958), p. 19.

in its domestic form of life. One of truth's everyday uses is to convey a statement's truth or falsity 'without actually having to repeat the statement; it is enough to be able to name the statement'.²⁶⁹ As Kripke puts it, 'we are entitled to assert (or deny) of any sentence that it is true precisely under the circumstances when we can assert (or deny) the sentence itself'.²⁷⁰ This is at play in the example of truth-involving inference which McGee gives, paraphrased as follows: A friend who I trust completely tells me that 'Everything Wendy says is true'. I meet Wendy, and Wendy tells me that 'Mogulcorp will go bankrupt'. I sell all my shares in Mogulcorp (and avoid a catastrophic financial hit).²⁷¹ Here I have reasoned that since whatever Wendy says is true, and Wendy says Mogulcorp will go bankrupt, then 'Mogulcorp will go bankrupt' is true. Applying (R2) allows me to infer that Mogulcorp will go bankrupt, and hence make the prudent move to divest of my holdings.

Moving back to the formal theory of truth, Kripke goes on to consider different pairs of extensions and antiextensions of $T(x)$, for obviously we can construct sets (S_n) and (S_{n+1}) with different memberships. In particular, Kripke considers sets (S_1^*) and (S_2^*) which are said to *extend* (S_1) and (S_2) , by only adding members, without losing any members which were in (S_1) and (S_2) : 'any sentence that is true (or false) in (S_1, S_2) retains its truth value in (S_1^*, S_2^*) '.²⁷² We can then imagine extending these sets again and again, with more and more sentences which were previously of undetermined truth-value gaining entry to the extension or antiextension of $T(x)$: as Horsten puts it, 'as we proceed in the sequence of models, more and more sentences... end up in the extension or in the anti-extension of T , and once a sentence is in the extension (anti-extension) of T , it stays in forever after'.²⁷³ At some point we will reach a pair of extension and antiextensions of $T(x)$ to which no more sentences can be added. Kripke labels this model (S_n, S_{n+1}) the *fixed point*.

²⁶⁹ Ibid, p. 534.

²⁷⁰ Saul Kripke, 'Outline of a Theory of Truth', *The Journal of Philosophy* 72.19 (Nov. 1975), p. 701.

²⁷¹ McGee, *Truth, Vagueness, and Paradox*, p. 174.

²⁷² Ibid, p. 703.

²⁷³ Leon Horsten, *The Tarskian Turn: Deflationism and Axiomatic Truth* (Cambridge: MIT Press, 2011), p. 119.

Solomon Feferman converted Kripke's description of the construction of fixed-point models into an axiomatized theory of truth, known as KF, which has proved popular.²⁷⁴

The axioms can be formulated as follows:

$$\text{KF1} \quad \forall x(\text{Trueof}(x) \leftrightarrow x)$$

$$\text{KF2} \quad \forall x(\text{Trueof}(\neg x) \leftrightarrow \neg x)$$

$$\text{KF3} \quad \forall x(\text{Trueof}(\neg\neg x) \leftrightarrow x)$$

$$\text{KF4} \quad \forall x, y(\text{Trueof}(x \wedge y) \leftrightarrow \text{Trueof}(x) \wedge \text{Trueof}(y))$$

$$\text{KF5} \quad \forall x, y(\text{Trueof}(\neg(x \wedge y)) \leftrightarrow \text{Trueof}(\neg x) \vee \text{Trueof}(\neg y))$$

$$\text{KF6} \quad \forall \varphi \forall x(\text{Trueof}(\forall x \varphi(x)) \leftrightarrow \forall y(\text{Trueof}(\varphi(y))))$$

$$\text{KF7} \quad \forall \varphi \forall x(\text{Trueof}(\neg \forall x \varphi(x)) \leftrightarrow \exists y(\text{Trueof}(\neg \varphi(y))))$$

$$\text{KF8} \quad \forall x(\text{Trueof}^{n+1}(\text{Trueof}^n(x)) \leftrightarrow \text{Trueof}^n(x))$$

$$\text{KF9} \quad \forall x(\text{Trueof}^{n+1}(\neg \text{Trueof}^n(x)) \leftrightarrow \text{Trueof}^n(\neg x))$$

$$\text{KF10} \quad \forall x \neg(\text{Trueof}(x) \wedge \text{Trueof}(\neg x))$$

KF1 and KF2 state that 'P is the case' will only be true when P is the case, and 'P is not the case' will only be true when P is not the case. KF3, KF4, and KF5 validate standard logical relations, namely: that not-not-P is true if and only if P is true; that a conjunction is true if and only if both its conjuncts are true; and that a conjunction is not the case if either its conjuncts are not the case. KF6 and KF7 underpin quantification, and state that universally quantified statements are true if and only if every instance is true, and false if and only if there is at least one false instance. KF8 and KF9 govern the application of the truth predicate to itself, and provide that it will be true that P is true if and only P is true, while it will be true that it is not true that P if and only if it is true that not-P. KF10 is a guarantor of non-contradiction, stating that it will never be the case that both P and not-P are true.

²⁷⁴ Solomon Feferman, 'Reflecting on Incompleteness', *The Journal of Symbolic Logic*, 56.1 (Mar. 1991), pp. 1-49.

Tye briefly sketches a similar system:

Turning next to predicates, I suggest that for the purposes of formal semantics the following treatment suffices for any extensionally vague monadic predicate F : given a non-empty domain D , F is assigned an extension S and a counter-extension S' . S is the set of objects of which F is true; S' is the set of objects of which F is false. Since F is vague, S and S' are vague sets. Here the term 'vague set' is to be understood in the way explained in Section 1.

For any individual constant c , let i_c be the object in D assigned to c . Then F_c is true if i_c belongs to S ; F_c is false if i_c belongs to S' and F_c is indefinite if there is no determinate fact of the matter about whether i_c belongs to S (or to S'). The generalization to n -place predicates is straightforward.²⁷⁵

It can be seen that on Tye's system, the elements which are vague members of S and S' play the same role as the elements which are outside the extension and anti-extension in the above partial interpretation system; it is they which make statements neither true nor false.

The key difference between the systems is that Tye allows for the possibility of vague sets, whereas KF is built on classical set theory. This is one area where Tye's insistence on the part (b) of his definitions, the 'vaguely vague' nature of his sets, will play an important role. Recall that for Tye, there is no determinate fact of the matter as to whether there are elements which are neither members of S , vague members of S , or members of S' . This is to guard against the existence of a sharp boundary between determinate members and vague members. In contrast, in any given partial interpretation, every element will either be in the extension, the antiextension, or neither. I agree with Tye's emphasis on the avoidance of sharp boundaries so am happy to utilise his framework of vague extensions and antiextensions, but the more classical KF system is available for those who are squeamish about the idea of vague sets.

Tye proposes the use of a Kleene three-valued logic system to govern vague statements, given in the below tables, in which the truth-values 'truth' and 'falsity' are supplemented by 'indefinite':

²⁷⁵ Tye, 'Vague Objects', p. 545. Tye's conception of vague sets is given above in section 5.5.1.

P	¬P
T	F
I	I
F	T

V	T	I	F
T	T	T	T
I	T	I	I
F	T	I	F

→	T	I	F
T	T	I	F
I	T	I	I
F	T	T	T

∧	T	I	F
T	T	I	F
I	I	I	F
F	F	F	F

It can be easily seen that these tables agree with classical logic when the possibility of an indefinite truth value is removed. Therefore, when we consider precise properties such as ‘is a natural number strictly greater than 3’, it will be the case that, since every natural number will either have this property or not (and therefore definitely be in either the extension or the antiextension), then every sentence ascribing the property to a number will be either true or false.

More generally, as Tye notes, classical tautologies such as the Law of Excluded Middle (LEM) will become what he calls ‘quasi-tautologies’, which have no false substitution instances. While they can take the value of indefinite when P is indefinite, they are in all other cases true.²⁷⁶

But of course the reason for working with such a system is not to deal with precise properties; it is to accommodate cases such as the droplet which is a vague part of a cloud. The sentence ‘Droplet D is part of the cloud C’ will, on this system, be assigned the value I, due to D not being determinately part of the extension of the set ‘parts of C’, and not determinately being part of the antiextension either. That is, due to the genuine metaphysical vagueness of the state of affairs which obtains in this situation, a statement regarding it has a resultant vagueness in its truth value.

That is not to say that all statements involving the cloud will be vague; ‘the cloud C has more than a billion parts’ might just be straightforwardly (determinately) false, while ‘the cloud C has fewer than ten parts’ is straightforwardly (determinately) true. But statements regarding the vague cases, such as ‘Droplet D is part of the cloud C’, will remain vague, due to the obtaining of a vague state of affairs. Importantly, on my

²⁷⁶ Ibid., p. 544. For those who prefer to stick with KF, a similar result can be found: KF1 and KF2 together imply that the following will hold, which guarantee that Bivalence and the Law of Excluded Middle hold for such precise properties:

Precise Law of Excluded Middle (PLEM): $\forall x((x \vee \neg x) \leftrightarrow ((\text{Trueof}(x)) \vee (\text{Trueof}(\neg x))))$

Precise Bivalence (PB): $\forall x((\neg x \rightarrow x) \leftrightarrow (\neg \text{Trueof}(x) \rightarrow \text{Trueof}(x)))$

account there is no need to invoke any further conceptual apparatus to explain or analyse the vague nature of these statements; we do not need to explore the precisificational space, or evaluate its truth-value at various intended models and identify whether any of them disagree. We simply look at the world, see that there is genuine vagueness in the area which the sentence describes, and know that the sentence should be evaluated as vague.

It's important to note that on this account, sentences are not given a truth value of *I* simply because there is a lack of a truthmaker (or falsemaker) in the actual world; *I* indicates a genuine third state, that of there existing an *I*-maker, which is an instance of genuine metaphysical vagueness. As Barnes and Cameron said when they explicated what a Third-Way view is:

...in the presence of indeterminacy, neither the demands for *p*'s truth nor the demands for *p*'s falsity are met. What is met is a third, distinct, demand: the demands for its indeterminacy.²⁷⁷

Barnes and Cameron take it to be a disadvantage of Third-Way views that, unlike the Unsettledness View which 'is entirely non-committal with respect to one's broader metaphysics... The Third-Way View, by contrast, demands a metaphysics that allows for the special states of affairs associated with indeterminacy'.²⁷⁸ I do not see the problem with such a demand. After all, we have seen that there is compelling reason to think that genuine metaphysical vagueness exists, and are trying to provide an account of it; why should that account not demand the existence of a special vague state of affairs?

Interestingly Tye is somewhat less committal on this point, stating that the indefinite status 'is strictly speaking, not a truth-value at all but rather a truth-value gap', before commenting that in his 'view, there are gaps due to failure of reference or presupposition and gaps due to vagueness'.²⁷⁹ He clarifies in a footnote that while in cases of failure of reference or presupposition there is a gap because nothing has actually been said, in the case of gaps due to vagueness 'something is said which is neither true nor false'.²⁸⁰ Tye says no more on the matter, and I think it may well boil

²⁷⁷ Barnes and Cameron, 'Are There Indeterminate States of Affairs? No', p. 123.

²⁷⁸ *Ibid.*, p. 125.

²⁷⁹ Tye, 'Vague Objects', p. 544.

²⁸⁰ *Ibid.*, p. 544, fn. 17.

down to a matter of terminology rather than anything more consequential. Tye distinguishes a special type of gap caused by vagueness; I call that a distinct value.

5.3.3 Introducing a definitely operator

We now move on to the introduction of a definitely operator, to allow us to make higher-order statements of truth. This operator is governed by the following axioms:

$$\text{DT1} \quad \forall x_1, \dots, x_n (D\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n) \rightarrow D\text{Trueof}(\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n)))$$

$$\text{DT2} \quad \forall x_1, \dots, x_n (D\neg\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n) \rightarrow D\neg\text{Trueof}(\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n)))$$

$$\begin{aligned} \text{DT3} \quad & \forall x_1, \dots, x_n ((\neg D\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n)) \wedge \\ & (\neg D\neg\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n))) \rightarrow \\ & (\neg D\text{Trueof}(\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n)) \wedge \\ & \neg D\neg\text{Trueof}(\text{Trueof}(\langle Rv_1, \dots, v_n \rangle, x_1, \dots, x_n))) \end{aligned}$$

Non-technically, DT1 says that if a statement *R* is definitely true, then the statement ‘*R* is true’ is also definitely true. This makes intuitive sense, and avoids unpalatable consequences; we do not want it to be the case, for instance, that ‘Harry is bald’ is definitely true, but we are unable to say for definite whether “‘Harry is bald’ is true’ is true or false. A definitely true statement retains its definite status all the way up the higher orders of truth.

DT2 gives us the reverse of DT1; if a statement *R* is definitely not-true, then ‘*R* is true’ is definitely not true as well. This formalises our intuition that if a statement can be definitely be said not to be true, then a journey up the higher orders cannot unseat this status and lead to the possibility of it being true.

DT3 says that if *R* is not definitely true, and not definitely not-true, then ‘*R* is true’ is neither definitely true nor definitely not-true. This is the axiom which defines the unsettled status of vague instances. For if we take that statement *R* to be, say ‘The droplet *D* is part of the cloud’, then *R* is neither definitely true nor definitely not-true, as

it is a vague matter. DT3 tells us thus that in this case, ‘It is true that the droplet D is part of the cloud’ is also neither definitely true nor definitely not-true. This is precisely what we should want, otherwise we could end up in situations whereby a statement about the world has a vague truth status but higher-order statements regarding the truth of that statement could have a settled, definite, truth value. In effect, DT3 guarantees that vagueness persists all the way up.

5.4 Can we reason in cases of vagueness?

We have set up a three-valued system, whereby sentences concerning instances of genuine metaphysical vagueness are assigned the value of indefinite, due to the obtaining of vague states of affairs which satisfy the vague-making conditions laid down by those sentences.

Importantly, the system does not compel everything to be indefinite: as Tye’s explanation of quasi-tautologies, and the derivability of the Precise Law of Excluded Middle (PLEM) and Precise Bivalence (PB) in precise domains show, if something is determinate in the world, then its truth-value is determinate and classical reasoning can proceed.

More importantly, we can still reason with vague propositions and draw inferences based upon them. What are these permitted patterns of inference? The standard strong Kleene schema, shown above, seems plausible and can validate some intuitive entailments, even when the indefinite value is included. For instance, the following, among others, are valid in K_3 :

$$1. A \wedge B \vDash_{K_3} A, B$$

$$2. A \vDash_{K_3} A \vee B$$

$$3. A, A \rightarrow B \vDash_{K_3} B$$

$$4. A \rightarrow B, B \rightarrow C \vDash_{K_3} A \rightarrow C$$

One can reason in situations where vagueness may occur and still rely on the above patterns of inference. There is, though, nothing inherent to my account about accepting

K_3 as the particular system of logic which best represents how reasoning should work when an unsettled truth value is allowed. Some may prefer systems such as \mathcal{L}_3 , LP, or others; those technical discussions can be had at another time.²⁸¹ For now, it is enough to say that there are plausible systems available which can support the acceptance of a primitive vagueness, and therefore truth values being sometimes indefinite.

5.4.1 Penumbral truths

A potential challenge that has been raised to this kind of three-valued system, most famously by Fine, is that of penumbral truths, and how they fare on a three-valued approach. Penumbral truths are those that arise from considering the logical connections that can hold between indefinite sentences; the intuitive idea is that even in the borderline region, everything should not be lawless, and there may be informative (and, in fact, true) statements which we can make. The example Fine leads with, for instance, is of a blob which is of a hue on the borderline between pink and red. If we call the sentence ‘the blob is pink’ P and the sentence ‘the blob is red’ R, Fine argues that we should want ‘ $P \wedge R$ ’ to be evaluated as false, due to the predicates ‘is pink’ and ‘is red’ being contradictory. However, on the truth table given earlier, as P by itself is indefinite, and R is also indefinite (due to the blob being on the borderline), the conjunction of the two will be evaluated as being indefinite. This goes against our intuition about the sentence’s real truth-value, and according to Fine shows that a three-valued approach cannot provide an accurate portrayal of such situations.²⁸²

Going back to the droplet and the cloud example I have discussed earlier, we can derive sentences which are even more intuitively repellent. Consider the statement ‘droplet D partly overlaps with cloud C’; in a case of genuine metaphysical vagueness, on my view this is neither true nor false. Equally, consider the statement ‘droplet D is disjoint from cloud C’; again, by the same considerations, this is indefinite. Now if we evaluate the conjunction of these statements (‘droplet D partly overlaps with cloud C and droplet D

²⁸¹ See Graham Priest, *An Introduction to Non-Classical Logic* (Cambridge: Cambridge University Press, 2008), pp. 120-185.

²⁸² Kit Fine, ‘Vagueness, Truth and Logic’, *Synthese*, 30.3/4 (Apr. – May, 1975), pp. 269-270.

is disjoint from cloud C') we are forced to conclude that this, too, is indefinite. This appears to be an undesirable consequence of our view.

In the case of the blob which is borderline pink and borderline red, it may be possible to argue that it is not so abhorrent to claim that pink and red are not truly contradictory in order to defuse the problem (recall the example of iridescent feathers from Wilson in the previous chapter). But this avenue does not appear open in the case of a droplet being both overlapping and also disjoint. Any defence of the offending conjunction statement must, I think, begin with the acknowledgement that such sentences will never be assigned the value of truth. (I can already hear Fine objecting that this is rather thin gruel, akin to someone who has just keyed your car protesting that at least they did not break into it and drive away). It could be argued that the relevant intuition here is that such a sentence could never be true, and on a background assumption of bivalence this is equivalent to saying that such a sentence is always false. However, when we move to a three-valued system, the intuition that it could never be true is weakened to saying that it is always either false or indefinite.

The counter-reply here will be, I assume, that the intuition is stronger than that; it is that we would actually want to assert that such a sentence is false. Tappenden has proposed an interesting strategy for defusing this counter-reply. In order to do so, he introduces the notion of *articulation*, as distinct from *assertion*. To motivate the distinction, consider that sentences can be uttered for multiple purposes. One may utter a sentence S to convince the hearer that S is the case. In this case, it is important for S to be true for it to serve its purpose. One may also, though, utter S in order to 'induce a hearer to withdraw an assertion of not-S'.²⁸³ In this case, what is important is that not-S cannot be correctly said to be true; it is this fact that you are trying to convince the hearer of. Importantly, though, it is not strictly necessary for S to be true for the utterance to successfully perform its function. Tappenden calls the former use case *assertion*, and the latter *articulation*, and speculates that the penumbral intuition arises from confusion that when we are articulating penumbral truths (to discourage the

²⁸³ Jamie Tappenden, 'The Liar and Sorites Paradoxes: Toward a Unified Treatment', *The Journal of Philosophy*, 90.11 (Nov., 1993), p. 571. In most instances; there will of course be cases where I am deliberately lying to you and wish you to believe S, even though I believe S to be false (or am indifferent to its truth). Such situations do not detract from the thrust of Tappenden's argument.

assertion of their negation), we think that we are in fact asserting them. He gives the following example:

Say you hear the following (said of two borderline tall men): "Joe is tall and John isn't." You know that John is taller than Joe, so you respond: "If Joe is tall then John is." You may not want to commit yourself to drawing boundaries sharp enough so that 'If Joe is tall then John is tall' would come out true. The point of uttering 'If Joe is tall then John is' is not to state the facts to someone who has not committed himself to an opinion, but rather to induce the withdrawal of the first sentence, the first speaker's assertion of which is a mistake. So in this example, 'If Joe is tall then John is' is articulated. The articulation of the sentence is correct if in fact the first utterance was a mistake. That is, the articulation of 'If Joe is tall then John is' is correct if it is incorrect to regard 'Joe is tall and John isn't' as true. The fact about the truth status of 'If Joe is tall then John is' which is relevant to the articulation being correct is that in this context the sentence is not correctly called false.²⁸⁴

Not everyone will find Tappenden's account convincing, but it is a worthwhile attempt to defuse some of the power of the penumbral intuition.

The final, potentially buck-passing, reply to the objection from penumbral truths is that other views do not get off scot-free (as we saw earlier in section 2.2). For supervenient-type theories (whether they be linguistic or metaphysical), they are saddled with the issue of endorsing what Fine called 'truth-value shifts'. This is because they can endorse the truth (or falsity) of penumbral statements only through appealing to super-truth, and in each precisification the statement is made true (or false) by a different instance (the truth 'shifts' between different disjuncts or conjuncts depending on the precisification). It is uncomfortable to be put in the position of affirming a conjunction while being unable to affirm either conjunct singly.

Epistemic theories are often thought to be immune to penumbral objections, but there is one way of looking at it on which they pose an issue. Take the general inductive statement of a sorites series on colour, where each stepwise pair is indiscernibly different: 'if *a* appears red to me, and *b* is indiscernibly different from *a*, then *b* appears red to me'. This has all the hallmarks of a penumbral truth, because of the nature of indiscernibility; if one thing is indiscernibly different to another, we cannot assign one a colour predicate and not the other. However, this is a statement that the epistemicist,

²⁸⁴ Ibid., p. 572.

with their commitment to sharp boundaries, must make false in one instance; there will be an instance, in such a series, where *a* is red but *b* is not red (even though we are unable to know where this point is).

The point of the above two paragraphs is not just to point out flaws in other theories to distract from my own shortcomings; but it should hopefully cast doubt on the efficacy of relying completely on pre-theoretic intuitions about supposed penumbral truths. In the words of Keefe, 'we already knew that we would have to accept *something* counter-intuitive... so admitting *some* counter-intuitive consequence in connection with the sorites paradox is a disadvantage which will plague *any* theory of vagueness'.²⁸⁵ No-one escapes unscathed.

5.5 What is the sorites solution offered by this?

We may now move back to the motivating question for this whole thesis. How does the above framework, neat and consistent though it may be, actually help us to solve the apparent paradox of sorites series?

Let us review the conditional formation of the paradox, using the example of the question of a rock's being located on Kilimanjaro or not. The below sets out a number of premises, each of which appear to be true, but when followed through, lead to an unpalatable conclusion.

1. A rock located at position *n* is located on Kilimanjaro.
2. If a rock located at position *n* is located on Kilimanjaro, then a rock located at position *n*+1 (where *n*+1 is 1mm further downhill than *n*) is also located on Kilimanjaro.
3. If a rock located at position *n*+1 is located on Kilimanjaro, then a rock located at position *n*+2 (where *n*+2 is 1mm further downhill than *n*+1) is also located on Kilimanjaro.

²⁸⁵ Keefe, *Theories of Vagueness*, p. 183; emphasis in original.

:

:

10,001. If a rock located at position $n+9,999$ is located on Kilimanjaro, then a rock located at position $n+10,000$ (where $n+10,000$ is 1mm further downhill than $n+9,999$) is also located on Kilimanjaro.

Therefore: A rock located at $n+10,000$ is located on Kilimanjaro.

(Let us assume, for the sake of argument, that location $n+10,000$ is clearly not on Kilimanjaro. It makes no material difference to the argument how many intermediate premises it contains, so the number could be increased to however many is needed.)

Tye has outlined a solution to the question of where this argument can be blocked to escape the contradictory conclusion, on his account of vague objects; as my view is a generalisation of Tye's, the solution can be carried over. The initial answer is the customary one from three-valued systems, which is that some of the intermediate premises of the argument will be indefinite in truth value, and therefore the overall argument will not be valid. This is because some of the intermediate premises pertain to vague parts of Kilimanjaro, for which the statement 'a rock located at position $n+x$ is located on Kilimanjaro' will be assigned a value of indefinite. So far so simple.

As Tye admits, though, there is an immediate further query for this solution:

Call these statements [in the Sorites series] $M_0, M_1, \dots, M_{1000000}$. Surely, it may be said, it can be demonstrated that there is some statement, M_k , such that M_k is true and M_{k+1} is not true. For suppose that there is no such statement. Then it follows that for any statement, M_k , if M_k is true then M_{k+1} is true. And from this, given that M_0 is true, by repeated applications of universal instantiation and modus ponens it may be inferred that $M_{1000000}$ is true. But $M_{1000000}$ is false. So, there is a sharp transition from the true statements in the sequence to the indefinite ones. This claim is no more plausible, however, than the already rejected claim that the addition of a single hair changes a bald man into a man who is not bald.²⁸⁶

²⁸⁶ Tye, 'Vague Objects', p. 549.

To put it another way, the motivating intuition behind the sorites paradox is that there is not a sharp boundary in series of this type, and this intuition is just as strong when considering a putative sharp boundary between determinate cases and vague cases.

Tye's answer to the question of the location of a transition in the above series of conditionals is a consequence of, or perhaps the reason for his emphasis on, part (b) of his definition of what it is to be vague. He argues that there is no sharp cut-off not only between the last true statement and the first false one, but there is none between the last statement which is true and the first which is indefinite:

We should hold that it is neither true nor false that there is a first conditional that is not true. Thus, there are true conditionals initially, and indefinite conditionals later, but it is not true that there is a sharp transition from the former to the latter.²⁸⁷

Tye addresses the immediate rejoinder that, as a matter of logic, if we start with true conditionals and end with false conditionals, then at some point there has to have been a transition of some nature; and if the first conditional is true, then there must be a conditional later down the series which is the first to be not true. He blocks this by questioning the unstated assumption that every statement in the series is either true or not true. According to Tye the truth-value predicates themselves are 'vaguely vague': 'there simply is no determinate fact of the matter about whether the properties they express have or could have any borderline instances'.²⁸⁸ That is, there is no determinate fact of the matter as to whether there are statements that are neither true nor false nor indefinite. This allows Tye to reject the assumption that every statement in the above sorites series is either true or not true, since this is equivalent to saying that every statement is either true or false or indefinite.

While it is true that this blocks the argument for a sharp cut-off between truth and non-truth in the series, and is a way of respecting the no-sharp-cut-offs intuition, it can be a struggle to grasp the idea. Tye explicitly argues that true, false, and indefinite are not extensionally vague, and states that it is not true that there are sentences which are

²⁸⁷ Ibid., p. 553.

²⁸⁸ Ibid., p. 551.

neither true nor false nor indefinite; but in the next breath admits that it can't be the case that every sentence is either true or false or indefinite:

For if 'is true' is extensionally vague then it follows that the set of true sentences has borderline members. This requires that there be sentences which are such that it is neither true nor false that they are true. And this, in turn, requires that there be sentences that are neither true nor false nor indefinite. I maintain that it is not true that there are such sentences. So 'is true' cannot be classified as extensionally vague. And the same goes mutatis mutandis for 'is false' and 'is indefinite'. Of course, in taking this view I am not committing myself to the position that these predicates are precise. Indeed, it is crucial to my account that they not be precise. For if they were then every sentence would be either true or false or indefinite, and that would not only generate sorites difficulties of its own (as we shall shortly see) but also run counter to my claim that it is indefinite whether no statement of the form 'Mn is true' is indefinite. Rather my view on the truth-value predicates is that they are vaguely vague: there simply is no determinate fact of the matter about whether the properties they express have or could have any borderline instances. So, it is indefinite whether there are any sentences that are neither true nor false nor indefinite.²⁸⁹

Tye is committed to second-order vagueness regarding the truth-value predicates, whereby they are not vague in the same way that Everest or 'baldness' are, by having vague instances; in contrast, for the truth-value predicates it is vague even as to whether they could be vague (by having vague instances).

I think the account is consistent, although thorny to think through. As we have seen, though, the paradoxes of vagueness, and especially the sorites paradox and considerations of higher-order vagueness, do not admit of clear solutions. Such logical contortions may well be a price worth paying for a generally convincing account of genuine metaphysical vagueness, which I think an extension of Tye's general treatment of vagueness is.

5.6 Conclusion

We have thus arrived at an account of metaphysical vagueness which attempts to address the difficulties discussed earlier for previous theories. One which takes metaphysical vagueness seriously, attempts to respect (at least some) pre-theoretic

²⁸⁹ Ibid., pp. 550 – 551.

intuitions, is not imperialist (which admittedly some might see as a fault, rather than a strength), adds theoretic expressiveness by adding vagueness to our primitives, and makes necessary adjustments to our theory of truth in cases of vagueness, while doing so within a framework that can retain a classical framework for precise properties and objects. In the following section of this thesis I will evaluate it against the desiderata set out earlier, to argue that it is the most promising answer to the research question.

6. Conclusion

We are now in a position to consider an overall response to this thesis' research question:

Is there genuine metaphysical vagueness?

The answer is distinguished into two parts. Initially, I wish to answer affirmatively that genuine metaphysical vagueness should be considered a realistic possibility in the logical space of explanations of the phenomena of vagueness. After exploring how vagueness as a phenomenon should be characterised in the Introduction, I argued indirectly for this first affirmative answer in Part A of the thesis, by critically analysing the two competing alternatives to metaphysical vagueness.

First was the idea that vagueness is wholly located in semantics and our linguistic representations of the world, represented chiefly by supervaluationist theories. I found that, while extensively developed and internally consistent, supervaluationist theories face some serious difficulties. Most importantly, in my opinion, are the issues related to reference and the link between language and the world (which will occur for any view which locates vagueness solely in language). That is, linguistic theories of vagueness imply that the world is fully determinate and precise, but for some reason a disconnect occurs whereby our representation of the world becomes vague; I do not claim that it will be impossible to provide some explanation of why this happens, but find it hard to imagine a plausible systemic answer. In addition, supervaluationist theories in particular struggle to maintain a correspondence theory of truth, due to the lack of a single object which can act as a referent for singular terms.

I also considered the view that vagueness is in fact an epistemic phenomenon; a specific kind of ignorance as to where the boundaries are located in certain cases (even though specific boundaries are determinately in existence). While Williamson's presentation of the theory in particular is elegant and on first appearance offers a concise, if unintuitive, explanation of why the phenomena of vagueness arise, on deeper consideration it also faces numerous pressing issues. Some of these are related, and combine to make the view implausible. The argument that vague predicates determine a sharp boundary of applicability through their use, and are

sensitive to small changes in that use, raises the question of how precisely this determination occurs. Moreover, even supposing that it does, how are language users meant to be able to track this sensitivity to use, and so to gain knowledge of the meaning of such terms? Finally, the intuition that there are no sharp boundaries in cases of vagueness is incredibly widespread, so much so that the entire literature on vagueness depends on it; but epistemicism struggles to provide an explanation for why this should be so, given the existence of sharp boundaries on their view.

If either of these models were able to provide a comprehensive and coherent account which ably explained all instances of vagueness, our intuitions surrounding them, and how they interacted with our systems of reasoning, then there would be a case for concluding that positing metaphysical vagueness is unmotivated, and ontologically and ideologically profligate. However, I have shown that neither account is free from problems. Of course, this is generally unavoidable when you are in the presence of paradox; as Keefe has observed, ‘admitting *some* counter-intuitive consequence in connection with the sorites paradox is a disadvantage which will plague *any* theory of vagueness’.²⁹⁰ Overall, though, in Part A I have shown that an affirmative answer to the research question deserve serious consideration.

In Part B, I argued more directly for the existence of genuine metaphysical vagueness. I first considered the arguments that have historically been advanced to show that the existence of vagueness in a metaphysical sense, that is vagueness that exists independent of our thoughts and representations, is an incoherent or inconsistent notion. Some were found to essentially be invocations of intuitions, which are incapable of blocking off a whole avenue of enquiry. Others, such as the argument from identity, while worthwhile and requiring serious engagement, have coherent replies available to them.

I then sought to examine what particular shape an affirmative answer to the research question should take; that is, if we want to argue that genuine metaphysical vagueness does exist, then what does it consist in? To that end, I discussed various accounts of genuinely metaphysical vagueness which have been proposed in the literature. While I

²⁹⁰ Keefe, *Theories of Vagueness*, p. 183; emphasis in original.

argued that none were wholly convincing, each is worthwhile in opening up the debate on metaphysical vagueness. Furthermore, each made an honest attempt to respect the intuitions, which I believe to be widespread, regarding the pervasiveness of vagueness in the fundamental make-up of the world itself. The supervaluationally-structured systems of Akiba and of Barnes and Williams were the most fleshed-out attempts to offer a wide-ranging account of all metaphysical vagueness. However, both suffer from their use of such a system, which I argued did not adequately locate the vagueness in the world; in each case, the apparent vagueness was reducible to fully precise structures, whether they be states of affairs or possible world. Jessica Wilson offered a general account of metaphysical vagueness that does indeed locate the vagueness at the 'object-level' rather than the 'meta-level', but her doctrine of determinables without unique determinants was argued to be both somewhat obscure, and failed to offer a plausible account of the various types of metaphysical vagueness.²⁹¹

Finally, I outlined my own attempt at such an account. It is one which takes vagueness to be both genuinely metaphysical, and also fundamental to our ideology; vagueness is irreducibly present both in our theory and descriptions of the world, and in the world itself. Indeed, the former is the case because of the latter. I explained Tye's account of genuinely metaphysically vague objects, and argued that it could be profitably expanded. My view, then, is one whereby genuine metaphysical vagueness consists in the determinate obtaining of vague states of affairs; vague states of affairs being those which constitutively involve vague entities. The view requires some alterations to classical systems of truth and logic, which I have provided a sketch of; for instance, I follow Tye in adopting a three-valued logic system, although disagree with him in that I believe we should interpret the third value as a genuine third (vague) value, rather than as simply a truth-value gap. I have argued that each of the alterations are justified and acceptable, in that they maintain some classical characteristics when the domain is restricted to precise objects and instances, while providing the flexibility to handle and reason in vague contexts.

²⁹¹ Wilson, 'A Determinable-Based Account of Metaphysical Indeterminacy', p. 360.

I will now consider my view against the desiderata set out earlier in the thesis, in order to argue that it deserves to be seen as the best available answer to the research question.

6.1 Evaluation of my account against desiderata

Having sketched out my account of metaphysical vagueness, it is hopefully fairly clear to see how it fulfils the desiderata we set out for a theory of metaphysical vagueness at the start of Chapter 5. One was to explain why the existence of metaphysical vagueness leads to the phenomenon of epistemic tolerance which we are taking to be characteristic of vagueness. In instances of genuine metaphysical vagueness, there is a fundamental vagueness as to whether some thing is part of another, or whether some group of things compose another, etc.; a vague state of affairs obtains such that, in the borderline region, a certain rock is a vague part of Everest, or a certain droplet is a vague part of the cloud. There is no fact of the matter as to where the border of the macro-object is; or to put it another way, there is the fact that the border is vague. Thus it is clear why small differences cannot make a known difference to the application of the relevant predicate – there is no cut-off to make the proposed statement true or false, so there is no difference which can be known.

The second desideratum was to explain the apparent links between the various types of metaphysical vagueness. This explanation fairly falls out of the primitive nature of my use of the notion ‘vague’, and a consideration of the sort of metaphysical relations under examination, without requiring (or admitting of) much analysis. If it is genuinely metaphysically vague as to whether some thing is part of another thing, then (taking into account general principles of parthood and composition) it will be genuinely metaphysically vague as to whether the collection of the determinate parts of the larger object, plus that thing which is vaguely a part of it, together compose an object. Similarly, it can be imagined that in a case where there are only two things which might vaguely compose one object, it is really one and the same as saying that it is vague whether each one is part of the potential object composed of them both. The same

considerations can apply in potential cases of vague composition leading to vagueness as to whether an object exists or not, or whether two objects are identical.

In actual fact, given that my account takes vagueness to be primitive, and given that the notions of parthood, composition, etc. may be taken to interact in various ways, I think that my account can actually be consistent with different theories of composition or parthood, as long as the possibility for fundamental vagueness is included; the specifics of the interactions, then, might vary on different interpretations. This may be seen as a positive of the theory.

The third was to explain why metaphysical vagueness can lead to, or at least be connected to, instances of semantic and epistemic vagueness (that is, ignorance whether something is the case or not). Considering that, on my account, the world itself is genuinely vague, and our language has developed in order to accurately represent the world and enable us to communicate with each other about what is going on in the world, does it not stand to reason that language might be vague because (at least in part) the world is vague? That is, the reason that some predicates have vague regions of application is that they apply to areas of the world where there are no determinate facts of the matter available to ground a complete bivalent semantics. More straightforwardly, if you take the meaning of a name to be its referent, then linguistic vagueness can be inherited directly:

In the case of 'Everest', its meaning is, I believe, its reference. And its reference is the vague object, Everest... Linguistic vagueness, then, is as widespread as non-linguistic vagueness. Neither, I maintain, is confined to a narrow domain.²⁹²

This is related to an important point, and I believe a positive for, my account; that it is not meant to apply in each and every apparent case of vagueness. That is, like Barnes and Williams, I am not intending to present an 'imperialistic' theory which purports to explain all vagueness; rather, I am 'quite happy to think that there will also be semantic and perhaps even epistemic indefiniteness also'.²⁹³

²⁹² Tye, 'Vague Objects', p. 540.

²⁹³ Barnes and Williams, 'A Theory Of Metaphysical Indeterminacy', p. 106, fn. 3.

Take the following example. Suppose we say of a man (a) who is 170cm tall that ‘a is short for a man in the UK in the 21st century’. This presumably has a settled truth-value (namely, truth). Now produce a sorites series of men of increasing heights, going up in 1mm increments, until reaching a man who is 190cm tall. This is a case where vagueness occurs, as surely as it does in the case of the rocks on Kilimanjaro. However, it is not obvious to me that this is a case of metaphysical vagueness, as I do not believe there to be metaphysical facts regarding the class of men in the UK in the 21st century, their average heights, etc. These are semantic questions, and the vagueness in this case is therefore best explained as semantic – vagueness as to the extension of ‘short’ in this country at the current time, for instance. Of course, once we probe further there may be issues of higher-order vagueness in this case where epistemic concerns might come into play, introducing another type of vagueness.

However, it is important to note that while the vagueness in the above example may not be fundamental ontology-wise, the semantic (and epistemic) explanation needed may well involve the ideologically primitive vagueness concept argued for above. That is, the semantics of ‘short’ will plausibly involve vagueness, in that its explanation will utilise the fundamental property of vagueness. Tye raised a similar thought in his discussion of the vagueness of ‘bald’:

In the case of 'bald', its intensional meaning is the property or concept it expresses and its extensional meaning (assuming that meaning can have an extensional sense) is the set of bald men. These, too, are vague.²⁹⁴

It thus seems unwise for the proponent of metaphysical vagueness to endorse the view that its existence obviates the need to include semantic vagueness in our worldview.

The fourth desiderata was that a theory should not reduce metaphysical vagueness out of existence; that is, it should not end up with vagueness dropping out of the fundamental metaphysical account of the world. Hopefully it should not need further explanation as to how my account fulfils this.

²⁹⁴ Tye, ‘Vague Objects’, p. 540.

6.2 General considerations

We also mentioned wider desiderata for any philosophical theory: explanatory usefulness, internal consistency and coherence, a lack of needless complexity, genuinely engaging with the problem at hand, etc. Hopefully the issues of explanatory usefulness and internal coherence and consistency have been addressed throughout the exposition of my account, and its situation within the wider debate over metaphysical vagueness.

I wish to draw particular attention to the ideal of genuinely engaging with the intuitions behind metaphysical vagueness. As Tye put it when concluding that none of the arguments which supposedly preclude the existence of vague objects (or metaphysical vagueness more generally) are conclusive, ‘philosophical arguments rarely, if ever, demonstrate that fundamental common-sense beliefs are mistaken or incoherent’.²⁹⁵ There is a widespread intuition that vagueness exists as a genuine phenomenon, and moreover that some of it seems to be located in the world itself. It is right to respect that intuition if possible, and only deny it if the path becomes impossible. I hope to have shown that there is no need to deny; we can build an account of genuine metaphysical vagueness which does in fact seriously engage with the idea that the world itself is vague.

This comes back to Barnes’ thoughts about fundamental vagueness (or to use her term, indeterminacy):

I’m sometimes tempted by the thought that a case for metaphysical indeterminacy—or at least an interesting case for metaphysical indeterminacy—requires commitment to fundamentality of indeterminacy itself (i.e., to a fundamental theory which is indeterminacy-involving). If we want to argue that the world could really be indeterminate, this plausibly requires not simply that there must there be indeterminacy in what’s fundamental... but also that indeterminacy must itself be fundamental.²⁹⁶

I have attempted to do both. My account argues that there is vagueness in what’s fundamental in the world; without taking a stand on the debates about what fundamentality is (or whether indeed there is such a thing), the ‘ontological bedrock’ (in

²⁹⁵ Tye, ‘Vague Objects’, p. 557.

²⁹⁶ Barnes, ‘Fundamental Indeterminacy’, p. 346.

whichever way you flesh it out) will contain vagueness. Moreover, this vagueness cannot be reduced, analysed, or explained by recourse only to precise concepts. Vagueness itself is fundamental on my account.

As Barnes herself argues, ‘derivative’ vagueness, that which arises from the combination of determinate, fundamental components, is too easy to dismiss. Firstly, it is unclear how any combination of purely determinate, precise components can ever produce vagueness out of nothing, as it were. Even if this is allowed, though, once vagueness is identified as derivative, its puzzling and paradoxical consequences become eminently deniable. Take the positing of sharp, apparently arbitrary, cut-offs with regard to composition, in order to escape sorites-type considerations:

If derivative entities exist, but don’t ‘really exist’ or don’t exist ‘in reality’, then a certain amount of arbitrariness in derivative ontology looks wholly unproblematic. But even those with a more inflationary view of derivative ontology might be happier to locate arbitrariness and cut-offs in derivative ontology than in fundamental ontology—and likely far happier to allow such costs in derivative ontology than to embrace indeterminacy.²⁹⁷

In effect, relegating vagueness to the derivative reduces the need to genuinely engage with it. It becomes an inconvenience, albeit one which can be ignored when considerations are restricted to the real, fundamental nature of the world. This is, I believe, a negative answer to the research question by the back door; as such, I believe my commitment to the fundamentality of vagueness in both senses is *prima facie* a benefit in its accordance with common sense.

I think the charge that will be levelled most forcefully against my account is that of needless complexity, both in terms of the apparatus of genuine metaphysically vague relation, and the three-valued logic invoked to support it.

The criticism would be something like the following. We saw earlier that a metaphysical supervaluational semantics, such as that proposed by Barnes and Willaims, can offer an account of metaphysical vagueness which retains a fully classical framework, requires minimal revision to previous theoretical commitments, and is compatible with various wider metaphysical interpretations. Assuming that genuine metaphysical

²⁹⁷ *Ibid.*, p. 344.

vagueness exists and we should provide an account of it, why not stick with one like Barnes and Williams and avoid the complexity of mine?

The answer, as prefigured earlier, is that I do not believe such a supervenient view does enough justice to the truly metaphysical phenomena of vagueness; the account, relying as it does purely on maximally precise possible worlds and the indeterminacy of the actualising relation, does not make enough space for the fundamentality of vagueness. That is, while such a view avoids complexity and can retain classical logic, it does so at the much larger expense of being able to offer a genuine explanation of the phenomena at hand. My view, while introducing genuinely vague states of affairs and requiring a move to a three-valued system, does manage to maintain a fundamentally metaphysical vagueness, while continuing to be able to communicate and reason about it (via the formal structures outlined above). My response, then, is that the complexity introduced by my account is not needless: it is in fact the minimal level of complexity needed in order to fully reckon with the nature of genuine metaphysical vagueness.

The final objection is one common to all novel accounts; what positives does the model give us, and are these positives worth changing our theory for? The answers can hopefully be found above, but to briefly recapitulate: the account is intended to be simple and intuitive, with vagueness taken as a primitive both metaphysically and ideologically; in doing so it respects metaphysical vagueness as a phenomenon, without analysing or reducing it out of existence in deference to a pre-conceived bias towards determinacy; it provides an explanation of where (some) semantic vagueness comes from – some terms represent objects or properties which are metaphysically vague, and therefore are vague themselves; relatedly it leads to a simple account of reference, with one (vague) object per (vague) term, and correspondence truth (see the discussion in section 2.2.5).

6.3 Areas for further research

While I hope to have offered a convincing positive answer to the research question, there are of course many aspects which could have warranted much further discussion;

indeed some could sustain whole theses on their own, and could be explored in further research.

For instance, a full exploration of how my account deals with various soritical instances would be worthwhile, along with a discussion of other ways in which metaphysical vagueness can manifest. Barnes and Williams discuss the example of the vagueness as to when a foetus becomes a person; such a discussion might entail both the metaphysical discussion of how such vagueness should be characterised, if and how we are able to gain knowledge of the relevant facts, as well as the ramifications for ethical theories pertaining to the value of such vague instances of personhood.²⁹⁸

Relatedly, most of the examples considered in this thesis have ignored change across time, but this phenomenon may plausibly be thought to give rise to a large number of cases of vagueness. These could range from the purely mereological ‘Ship of Theseus’-type considerations, to more *recherché* (but perhaps increasingly realistic) science fiction scenarios. Imagine that an individual’s brain states are able to be transferred into a machine, or a different set of organic matter; it may be vague as to where the individual is located, or it may be vague as to what percentage of brain states must be transferred before there is any change in the individual’s location. There may be a vague process by which memories and characteristics are forgotten or lost by the original individual, while they are recreated elsewhere. All of these might plausibly involve metaphysical vagueness, and a rich discussion could be had of the involved consequences for our views of personhood, persistence, responsibility, etc.

The more technical and formal aspects of the account that I have provided of metaphysical vagueness would also merit further investigation. I have sketched a formulation along three-valued lines which I hope has been shown to be promising and coherent, but when we move away from purely classical systems there are always choices to be made about how we construct our systems of reasoning, entailment, and truth, which axioms we accept and which patterns of inference we endorse. Not all of them are necessarily enforced by the existence of primitive metaphysical vagueness, and it is worth digging deeply into the possibilities and the different consequences they

²⁹⁸ Barnes and Williams, ‘A Theory Of Metaphysical Indeterminacy’, p. 114.

would lead to when applied to the various manifestations of vagueness. As I have done throughout the thesis, a cost-benefit analysis should be attempted to find the system which reckons best with pre-theoretic intuitions, pre-existing commitments, and interactions with non-vague parts of metaphysics, epistemology, and other philosophical disciplines.

Another consideration which I have briefly mentioned and which could be the basis of future work is the apparent interactions between metaphysical, linguistic, and epistemic vagueness. As I have stated earlier, I do not mean for my account to be taken as an ‘imperial’ theory of vagueness. That is, I do not believe that all vagueness is metaphysical in nature. Rather, I have argued for the existence of metaphysical vagueness in addition to linguistic and epistemic vagueness, as I think it is clear that one can cause the other. In the most basic case, I take it that our semantics and linguistic representation in general aims at accurately representing the world, to the best of our (sometimes limited) ability. If it is in fact true that the world itself is vague, then it stands to reason that our linguistic and mental representations of the world may inherit some of the vagueness; after all, when there is no matter of fact as to whether some thing is part of an object or not, our concepts would not be doing a very good job of tracking reality if they took a firm stand as to whether the thing was a part or not. The interaction between metaphysical and epistemic vagueness appears to be more complicated on first sight. Williamson’s particular epistemic theory will not necessarily apply in cases of genuine metaphysical vagueness, as the margin of error principle will not be satisfied.

But there is the possibility that epistemic principles may be invoked to deal with the issue of higher order vagueness, another topic which could stimulate much further productive work. In general I am sympathetic to Barnes’ scepticism with regards to the existence of genuine higher order metaphysical vagueness:

[W]hat more could there be to it being, e.g., indeterminate whether it’s indeterminate that p than it simply being indeterminate that p? If it’s indeterminate whether it’s indeterminate that p then, it seems, there is some indeterminacy with respect to p (the facts about p aren’t wholly ‘metaphysically settled’). But the idea that things aren’t quite settled with

respect to p just seems to amount to our basic notion that p is indeterminate.²⁹⁹

But this does not mean that we can ignore higher order vagueness and the intuitions that drive it. Plausibly, there may appear to be higher order vagueness surrounding instances of metaphysical indeterminacy due to ignorance as to whether p is indeterminate; this may be amenable to epistemicist interpretations and explanations.

Overall there are a myriad of areas within the topic of vagueness which await exploration, and which are opened up by taking metaphysical vagueness seriously as a possibility. I hope I have managed to convince the reader that this is a fruitful endeavour.

²⁹⁹ Barnes, 'Ontic Vagueness: A Guide for the Perplexed', p. 618. See also Hyde, 'Why Higher-Order Vagueness is a Pseudo-Problem'; Sainsbury, 'Is There Higher-Order Vagueness?'; Wright, 'Is Higher-Order Vagueness Coherent?'; and Wright, 'The Illusion of Higher-Order Vagueness'.

Bibliography

Akiba, Ken, 'Vagueness in the World', *Noûs*, 38.3 (2004), pp. 407-429.

----- 'How Barnes and Williams have failed to present an intelligible ontic theory of vagueness', *Analysis*, 75.4 (Oct., 2015), pp. 565-573.

Armstrong, David, *A World of States of Affairs* (Cambridge: Cambridge University Press, 1997).

Asher, Nicholas, Dever, Josh, and Pappas, Chris, 'Supervaluations Debugged', *Mind*, 118.472 (Oct., 2009), pp. 901-933.

Barnes, Elizabeth, 'Ontic Vagueness: A Guide for the Perplexed', *Noûs*, 44.4 (2010), pp. 601-627.

----- 'Fundamental Indeterminacy', *Analytic Philosophy*, 55.4 (Dec., 2014), pp. 339-362.

Barnes, Elizabeth, and Cameron, Ross, 'Are There Indeterminate States of Affairs? No', in *Current Controversies in Metaphysics* ed. Elizabeth Barnes (Abingdon: Routledge, 2017), pp. 120-132.

Barnes, Elizabeth and Williams, J. Robert G., 'A Theory Of Metaphysical Indeterminacy', in *Oxford Studies in Metaphysics: Volume 6* (Oxford: Oxford University Press, 2011), Karen Bennett and Dean W. Zimmerman, eds., pp. 103-147.

Burgess, J. A., 'Vague Objects and Indefinite Identity', *Philosophical Studies*, 59.3 (1990), pp. 263-28.

----- 'The Sorites Paradox and Higher-Order Vagueness', *Synthese*, 85.3 (1990), pp. 417-474.

Carmichael, Chad, 'Vague Composition Without Vague Existence', *Noûs* 45.2 (Jun., 2011), pp. 315-327.

Carnap, Rudolf, 'Testability and Meaning', *Philosophy of Science*, 3.4 (Oct., 1936), pp. 419-471.

Dummett, Michael, 'Wang's Paradox', *Synthese*, 30 (1975), pp. 301-324.

Edgington, Dorothy, 'Indeterminacy de re', *Philosophical Topics*, 28.1 (2000), pp. 27-44.

Eklund, Matti, 'Characterizing Vagueness', *Philosophy Compass*, 2/6 (2007), pp. 896-909.

Evans, Gareth, 'Can There Be Vague Objects?', *Analysis*, 38.4 (Oct., 1978), p. 208.

- Feferman, Solomon, 'Reflecting on Incompleteness', *The Journal of Symbolic Logic*, 56.1 (Mar., 1991), pp. 1-49.
- Fine, Kit, 'Vagueness, Truth and Logic', *Synthese*, 30.3/4 (Apr. – May, 1975), pp. 265-300.
- Frege, Gottlob, *Basic Laws of Arithmetic: Derived using concept-script* (Oxford: Oxford University Press, 2013), transl. and ed. P. Ebert and M. Rossberg.
- Geach, P. T., *Reference and Generality* (Ithaca: Cornell University Press, 1980).
- Greenough, Patrick, 'Vagueness: A Minimal Theory', *Mind*, 112.446 (Apr., 2003), pp. 235-281.
- Grice, H. P., 'Logic and Conversation', in *Syntax and semantics Vol. 3: Speech acts* eds. P. Cole, and J. Morgan (New York: Academic Press, 1975), pp. 41-58.
- Hawley, Katherine, 'Vagueness and Existence', *Proceedings of the Aristotelian Society*, 102 (2002), pp. 125-140.
- Horsten, Leon, *The Tarskian Turn: Deflationism and Axiomatic Truth* (Cambridge: MIT Press, 2011).
- Horwich, Paul, *Meaning* (Oxford: Oxford University Press, 1998).
- 'The Sharpness of Vague Terms', *Philosophical Topics*, 28.1 (Spring 2000), pp. 83-92.
- Hyde, Dominic, 'Why Higher-Order Vagueness is a Pseudo-Problem', *Mind*, 103.409 (Jan., 1994), pp. 35-41.
- van Inwagen, Peter, *Material Beings* (Ithaca: Cornell University Press, 1990).
- Keefe, Rosanna, and Smith, Peter, 'Theories of Vagueness', in *Vagueness: A Reader* eds. Rosanna Keefe and Peter Smith (Cambridge: MIT Press, 1997), pp. 1-57.
- Keefe, Rosanna, *Theories of Vagueness* (New York: Cambridge University Press, 2000).
- Khatchadourian, Haig, 'Vagueness', *The Philosophical Quarterly*, 12.47 (Apr., 1962), pp. 138-152.
- Kleene, *Introduction to Metamathematics* (Amsterdam: North Holland Publishing Co., 1952).
- Kripke, Saul, 'Outline of a Theory of Truth', *The Journal of Philosophy*, 72.19 (Nov., 1975), pp. 690-716.
- Leech, Jessica, *Thinking of Necessity* (Oxford: Oxford University Press, 2023).
- Lewis, David, *On the Plurality of Worlds* (Oxford: Blackwell, 1986).

----- 'Vague Identity: Evans Misunderstood', *Analysis*, 48.3 (Jun., 1988), pp. 128-130.

----- 'Many, But Almost One', in *Papers in metaphysics and epistemology* (Cambridge: Cambridge University Press, 1999), pp. 164-182.

MacFarlane, John, 'In What Sense (If Any) Is Logic Normative for Thought' (presented at the 2004 Central Division APA symposium on the normativity of logic), accessed at <<https://johnmacfarlane.net/normative.pdf>>.

McGee, Vann, 'Applying Kripke's Theory of Truth', *The Journal of Philosophy*, 86.10 (Oct., 1989), pp. 530-539.

----- *Truth, Vagueness, and Paradox* (Indianapolis: Hackett, 1990).

McGee, Vann, and McLaughlin, Brian P., 'Review of *Vagueness* by Timothy Williamson', *Linguistics and Philosophy*, 21.2 (Apr., 1998), pp. 221-235.

----- 'The Lessons of the Many', *Philosophical Topics*, 28.1 (Spring 2000), pp. 129-151.

McKinnon, Neil, 'Vague Simples', *Pacific Philosophical Quarterly*, 84 (2003), pp. 394-397.

Merriam-Webster, 'Vague', <<http://www.merriam-webster.com/dictionary/vague>>.

Miller, J. T. M., 'Are All Primitives Created Equal?', *The Southern Journal Of Philosophy*, 56.2 (Jun., 2018), pp. 273-292.

Moline, Jon, 'Aristotle, Eubulides and the Sorites', *Mind*, 78.322 (Jul., 1969), pp. 393-407.

Morreau, Michael, 'What Vague Objects Are Like', *The Journal of Philosophy*, 99.7 (Jul., 2002), pp. 333-361.

Noonan, Harold W., 'Are There Vague Objects?', *Analysis*, 64.2 (Apr., 2004), pp. 131-134.

Oddie, Graham, 'Armstrong on the Eleatic Principle and Abstract Entities', *Philosophical Studies*, 41.2 (Mar., 1982), pp. 285-295.

Priest, Graham, *An Introduction to Non-Classical Logic* (Cambridge: Cambridge University Press, 2008).

Quine, W. V., 'Ontology and Ideology', *Philosophical Studies*, 2.1 (Jan., 1951), pp. 11-15.

Raffman, Diana, 'Vagueness Without Paradox', *The Philosophical Review*, 103.1 (Jan., 1994), pp. 41-74.

- *Unruly Words: A Study of Vague Language* (Oxford: Oxford University Press, 2014).
- van Rooij, Robert, 'Vagueness and Linguistics', in *Vagueness: A Guide* ed. G. Ronzitti (Dordrecht: Springer, 2011), pp. 123-170.
- Rosen, Gideon, and Dorr, Cian, 'Composition as a Fiction', in *The Blackwell Guide to Metaphysics* ed. R. Gale (Oxford: Blackwell, 2002), pp. 151–174.
- Rosen, Gideon, and Smith, Nicholas J. J., 'Worldly Indeterminacy: A Rough Guide', *Australasian Journal of Philosophy*, 82.1 (Mar., 2004), pp. 185-198.
- Russell, Bertrand, 'Vagueness', *Australasian Journal of Psychology and Philosophy*, 1.2 (1923), pp. 84-92.
- Sainsbury, R. M., 'What Is a Vague Object?', *Analysis*, 49.3 (Jun., 1989), pp. 99-103.
- 'Is There Higher-Order Vagueness?', *The Philosophical Quarterly*, 41.163 (Apr. 1991), pp. 167-182.
- 'Vagueness, Ignorance, and Margin for Error', *The British Journal for the Philosophy of Science*, 46.4 (Dec., 1995), pp. 589-601.
- 'Concepts Without Boundaries', in *Vagueness: A Reader* eds. Rosanna Keefe and Peter Smith (Cambridge: MIT Press, 1997), pp. 251-264.
- Salmon, Nathan, *Reference and Essence* (Oxford: Blackwell, 1982).
- Sattig, Thomas, 'Mereological Indeterminacy: Metaphysical but Not Fundamental', in *Vague Objects and Vague Identity: New Essays on Ontic Vagueness* eds. Ken Akiba and Ali Abasnezhad (Dordrecht: Springer, 2014), pp. 25-42.
- Schiffer, Stephen, 'The Epistemic Theory of Vagueness', *Philosophical Perspectives*, 13 (1999), pp. 481-503.
- 'Skepticism and the Vagaries of Justified Belief', *Philosophical Studies*, 119.1-2 (2004), pp. 161-184.
- Shapiro, Stewart, *Vagueness in Context* (Oxford: Oxford University Press, 2006).
- Sider, Theodore, *Four-Dimensionalism: An Ontology of Persistence and Time* (Oxford: Clarendon, 2001).
- Smith, Nicholas J. J., 'A Plea for Things That Are Not Quite All There: Or, Is There a Problem about Vague Composition and Vague Existence?', *The Journal of Philosophy*, CII.8 (Aug., 2005), pp. 381-421.
- Sorensen, Roy, 'Vagueness, Measurement, and Blurriness', *Synthese*, 75.1 (Mar., 1988), pp. 45-82.

- 'Process Vagueness', *Linguistics and Philosophy*, 13.5 (Oct., 1990), pp. 589-618.
- *Vagueness and Contradiction* (New York: Oxford University Press, 2001).
- Tappenden, Jamie, 'The Liar and Sorites Paradoxes: Toward a Unified Treatment', *The Journal of Philosophy*, 90.11 (Nov., 1993), pp. 551-577.
- Thomason, Richmond H., 'Identity and Vagueness', *Philosophical Studies*, 42.3 (Nov., 1982), pp. 329-332.
- Tuggy, David, 'Ambiguity, Polysemy, and Vagueness', *Cognitive Linguistics*, 4.3 (1993), pp. 273-290.
- Tye, Michael, 'Vague Objects', *Mind*, 99.396 (Oct., 1990), pp. 535-557.
- 'Fuzzy Realism and the Problem of the Many', *Philosophical Studies*, 81 (1996), pp. 215-225.
- 'On the Epistemic Theory of Vagueness', *Philosophical Issues*, 8 (1997), pp. 247-253.
- 'Vagueness and Reality', *Philosophical Topics*, 28.1 (2000), pp. 195-209.
- Unger, Peter, 'The Problem of the Many', *Midwest Studies in Philosophy*, 5.1 (Sept., 1980), pp. 411-67.
- Weatherson, Brian, 'Vagueness as Indeterminacy', in *Cuts and Clouds: Vagueness, its Nature, and its Logic* eds. Richard Dietz and Sebastiano Moruzzi (Oxford: Oxford University Press, 2010), pp. 77-90.
- Williams, J. Robert G., 'Multiple Actualities and Ontically Vague Identity', *The Philosophical Quarterly*, 58.230 (Jan., 2008), pp. 134-154.
- Williamson, Timothy, *Vagueness* (New York: Routledge, 1996).
- 'Reply to Commentators', *Philosophy and Phenomenological Research*, 57.4 (Dec., 1997), pp. 945-953.
- Wilson, Jessica, 'Fundamental Determinables', *Philosopher's Imprint*, 12.4 (Feb., 2012), pp. 1-17.
- 'A Determinable-Based Account of Metaphysical Indeterminacy', *Inquiry*, 56.4 (2013), pp. 359-385.
- 'Are There Indeterminate States Of Affairs? Yes', in *Current Controversies in Metaphysics* ed. Elizabeth Barnes (Abingdon: Routledge, 2017), pp. 105-119.

Wittgenstein, Ludwig, *Philosophical Investigations* (Oxford: Blackwell, 1958).

Wright, Crispin, 'On the Coherence of Vague Predicates', *Synthese*, 30.3/4 (Apr. – May, 1975), pp. 325-365.

----- 'Is Higher-Order Vagueness Coherent?', *Analysis*, 52.3 (1992), pp. 129-139.

----- 'On Being in a Quandary', *Mind*, 110.437 (Jan., 2001), pp. 45-98.

----- 'The Illusion of Higher-Order Vagueness', in *Cuts and Clouds: Vagueness, its Nature, and its Logic* eds. Richard Dietz and Sebastiano Moruzzi (Oxford: Oxford University Press, 2010), pp. 523-549.