Vindicating Vague Objects

Michael Lazarou

B.A. (Advanced), B.A. Hons. (Philosophy)

Department of Philosophy, School of Humanities

The University of Adelaide

Thesis submitted in fulfilment of the requirements for the degree of Master of Philosophy

November | 2019

Table of Contents

Abstract 3

Thesis Declaration 4

Acknowledgements 5

Introduction 6

Chapter 1 – Vague Objects and the Problem of Vague Identity 9

1.0 Preamble 9

1.1 Evans' Argument and Proof 9

1.2 Lewis and the Fallacious Modal Equivalence 16

1.3 Logical Commitments of a Vague-Objects View 22

1.4 On the Determinacy of Distinctness 31

1.5 Taking Stock 38

Chapter 2 – Individuating Vague Objects: The Problem of the Many 40

2.0 Preamble 40

2.1 The Problem of the Many 40

2.2 Responses to the Problem of the Many 44

2.3 Vague Objects: From Many to One? 56

2.4 Weatherson's Objection: Many Problems 63

2.5 Responding to Weatherson: Vague Objects and Vague Parts 70

2.6 Barnes and Williams' Response to Weatherson: Vague Objects and Vague Identity 71

2.7 Taking Stock 75

Chapter 3 – Vague Objects and Metaphysical Vagueness 77

3.0 Preamble 77

3.1 Vague Objects and Metaphysical Vagueness 77

3.2 Precisificational Accounts of Metaphysical Vagueness 79

3.3 Lessons from Precisificational Accounts 91

3.4 Non-Precisificational Accounts of Metaphysical Vagueness 92

3.5 Lessons for a Defence of Vague Objects 100

3.6 Taking Stock – Looking Forward 101

Conclusion 103

Bibliography 104

Abstract

Until recently, vagueness has been seen as a product of representation alone; a byproduct of the limitations of our language or our ability to know the truth. To endorse the contrasting idea that vagueness can come from the world – and to endorse the existence of vague objects – has often been a maligned enterprise. Indeed, proponents of the view have been charged with mistaking features of the world for features of our language or minds, mistaking a question of language and epistemology for a question of metaphysics. Further, even granting the plausibility of such a view, the thesis that vague objects can exist has been laden with commitment to problematic notions such as vague identity and vague existence. This thesis examines the prospects of defending the idea that vague objects exist, vindicating the cogency of such a view and decoupling it from these problematic notions.

Chapter 1 begins by examining Gareth Evans' seminal *reductio* against vague objects, in which the existence of vague objects is tied to the fate of vague identity. Engaging with the literature that Evans' paper generated, we show that vague objects need not be committed to a contradictory notion of vague identity; rather, one can defend an account of vague objects without the need for revisionary logics or gerrymandered notions of identity. Chapter 2 extends the investigation of vague objects by considering the Problem of the Many, a powerful paradox which appears to undermine seemingly well-founded mereological principles and intuitions. After evaluating existing solutions to the problem, we show how vague objects can be used to develop a novel solution that is couched within the logical apparatus defended in Chapter 1. We then demonstrate how the novel solution offers a fruitful means of responding to the problem while retaining desired mereological principles. Chapter 3 draws on the discussion in the preceding chapters to defend the cogency of vague objects in the context of contemporary views of metaphysical vagueness. Specifically, we show that defending the existence of vague objects has genuine utility for supporting a defence of the intelligibility of metaphysical vagueness. We end with a brief examination of the prospects of extending the novel account of vague objects developed, and consider how this view may be applied to future investigations in metaphysics.

Thesis Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Signature:

Date: 07-11-2019

Acknowledgements

I am lucky to have had so much support from so many colleagues, friends, and family – too many to give due credit in so few words! To begin, particular thanks to the University of Adelaide Philosophy Department and my fellow postgraduate students. The collegiality, friendliness and camaraderie has been a highlight of my candidature. Particular thanks to Robert Farquharson, Riley Harris, Victoria Troitiño, Dook Shepherd, Brigitte Everett, James Vlachouilis, Atheer Al-Khalifa, Bec Randell, Danny Wardle and Ted Burston for all of the lively conversations!

Special thanks, of course, to my supervisors Antony Eagle and Garrett Cullity. Particular thanks to Antony for your wisdom, encouragement, and tremendous support as my primary supervisor, and for your ability to tease out the best of what were often half-complete ideas. Our meetings have been central to my development as a thinker. Thank you both for showing me what it means to be a good philosopher. I hope this thesis does justice to you both!

Thank you to all of my friends, particularly Christian, Hayden, and Ben, for the support you have provided as only good friends can. Thank you too to Emily for all of the support at home, and for your unending belief in me in times where it was especially needed.

A tremendous thank you to my family, and particularly Nasi and Helen. You have brought out the best in me, often in ways that I could not have conceived.

Finally, the greatest of thanks to my father, Yiannis. You have taught me what it means to be curious, what it means to be strong, what it means to be good, and what it means to be profoundly human. I owe all of my successes to you. I hope the good that I can produce in the world in some small way stands as a model of thanks for the life you have enabled me to live through your sacrifices as a single father of four. I dedicate this thesis to you.

Introduction

Vagueness and precision alike are characteristics which can only belong to a representation, of which language is an example… 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. – Bertrand Russell, 1923, p. 85.

Vagueness has historically been seen as a product of representation alone; a byproduct of the limitations of our language or our ability to know the truth. Predicates such as “bald” or “tall” appear to lack sharp cut-offs, giving rise to borderline cases which elude precise truth conditions. Even worse, upon close inspection it appears that vagueness proliferates, infecting our everyday language, our naming conventions, and our ability to describe the world. While various “paradigm cases” do exist, providing clear examples for our linguistic conventions, there are also borderline cases, where the polar judgements of truth and falsity seem to misfire. Indeed, the inability of our language to produce sharp cut-offs or precise conditions of application, arising from instances of vagueness, has motivated attempts to reconcile imperfect or unsettled linguistic practice with an otherwise precise world.

Yet, the considerations giving rise to vagueness in language could also be applied to the world. It seems plausible that objects may lack precise spatial boundaries, and that there may be vagueness with respect to properties such as colour, mass, or location. For instance, if we were asked to determine the exact collection of particles comprising the boundaries of a mountain, there is little to suggest that we could provide more than an approximation of this. While our means of representing the world may be imprecise, there is also *prima facie* reason to think that the world may itself give rise to instances of vagueness.

However, to endorse the idea that vagueness can come from the world itself – and to endorse the existence of vague objects – has been an historically maligned enterprise. We see this in Michael Dummett’s charge that worldly vagueness is ‘not properly intelligible’ (1975, p. 314), and in David Lewis’ doubt of having ‘any correct conception of a vague object’ (1993, p. 27). Further, proponents of the existence of vague objects have been charged with committing ‘the fallacy of verbalism’ (Russell 1923), mistaking features of the world for features of our language or minds, and so mistaking questions of language and epistemology for a question of metaphysics. Further, even granting the plausibility of such a view, the thesis that vague objects can exist has been laden

with commitment to revisionary logics, to problematic notions such as vague identity, and to deviant metaphysics – notions which the proponents of representational accounts of vagueness have been able to evade.

This thesis examines the prospects of defending the idea that vague objects exist, exploring the cogency of such a view and decoupling it from the need for extensive logical and metaphysical revision. In short, our goal is to vindicate vague objects and rescue them from their enduringly poor reputation.

Chapter 1 begins by examining Gareth Evans' seminal *reductio* against vague objects, in which the existence of vague objects is tied to the fate of vague identity. Upon engagement with the literature that Evans' paper generated, we show that vague objects need not be connected to a contradictory notion of vague identity; rather, one can defend an account of vague objects without the need for revisionary logics or gerrymandered notions of identity. We then consider Lewis' interpretation of the Evans argument and discuss the issue of *de dicto*/*de re* equivalences in vague contexts. We then defend logical conservatism as a framework for constructing a coherent account of vague objects, discussing related issues such as the problem of higher-order vagueness and the coherence of indeterminate distinctness as a logical notion. We conclude the chapter by arguing that there is good reason to see vague objects as vindicated from their maligned historical reputation; in short, we establish the cogency of vague objects' existence.

Having decoupled vague objects from contradiction, Chapter 2 demonstrates how vague objects can be used to develop a solution to the Problem of the Many, a powerful paradox which appears to undermine seemingly well-founded mereological principles and intuitions. Upon evaluating existing solutions to the problem, including 'supervaluationist' and 'many-one' solutions, we develop a novel vague-objects solution that comports with the logical apparatus defended in Chapter 1. In developing this view, we postulate notions of vague parthood and vague coincidence, establishing their use as part of our novel solution to the Problem of the Many. We then respond to objections, particularly Weatherson's (2003) *reductio* against vague objects, and demonstrate how our novel vague-objects solution can uphold desired mereological principles and our counting intuitions with respect to ordinary objects.

Chapter 3 draws on the discussion in the preceding chapters to defend the cogency of vague objects in the context of contemporary views of metaphysical vagueness. We discuss 'precisificational' and 'non-precisificational' accounts of metaphysical

vagueness, considering their relative merits and challenges in defending the view's intelligibility. We then argue that our novel account of vague objects is comparatively favourable, offering clear identity criteria for ordinary objects. We end with a brief examination of the prospects of extending the novel account of vague objects developed and consider how the view may be applied to future investigations in metaphysics.

Chapter 1 – Vague Objects and the Problem of Vague Identity

1.0 Preamble

This chapter discusses the connection between vague objects and vague identity, as seen in Gareth Evans' seminal paper 'Can There Be Vague Objects?' (1978). I first present Evans' argument before providing an extended response to his view that commitment to vague objects leads to contradiction. I argue that the fate of vague objects can be untangled from the problematic notion of vague identity, exploring the idea that the two concepts need not be inextricably intertwined. To argue for this, I unpack the logical assumptions that purportedly lead to contradiction and consider the role that classical and non-classical logics may play in rescuing vague objects from contradiction.

With a cogent response to Evans' argument in hand, I then unpack the broader logical apparatus needed to construct a coherent account of vague objects. Specifically, I engage with David Lewis' interpretation of Evans' argument, and consider broader issues prevalent in the vagueness literature. These include the issue of *de dicto*/*de re* equivalences in vague contexts, the problem of higher-order vagueness, and the contested coherence of indeterminate distinctness as a logical notion. In discussing these themes, I also respond to worries about the coherence and possibility of the existence of vague objects. I end the chapter by providing direction for how vague objects may be vindicated from their historically maligned reputation.

1.1 Evans' Argument and Proof

Gareth Evans' famous (1978) argument against the possibility of ontic vagueness is one of those philosophical problems that just won't die – Elizabeth Barnes, 2009, p. 81.

Gareth Evans' influential one-page paper, 'Can There Be Vague Objects?' (1978), offers a concise argument against the existence of vague objects. According to Evans, accepting the existence of vague objects carries a commitment to the notion of vague identity. By presenting a proof demonstrating how vague identity leads to contradiction, Evans argued against the tenability of vague identity, and by extension, the existence of vague objects. Following Garrett (2014), Evans' argument can be set out as follows:

(A) There can be vague objects only if there can be true vague identity statements.
(B) There can be no such things as true vague identity statements, lest we face accepting a logically inconsistent notion.
Thus
(C) There cannot be vague objects.

Evans' argument generated a remarkable amount of secondary literature. Much of this literature responds to Premise B – the idea that identity is a non-vague, and so determinate, notion. Now, while this is undoubtedly a central debate in the vagueness literature (indeed, one to which we will return), it is first important to determine how Evans tied the existence of vague objects to the fate of vague identity, as in Premise A. Here Evans' conception of vague objects plays an important role. While brief, his account of vague objects is telling:

> It is sometimes said that the world might itself *be* vague. Rather than vagueness being a deficiency in our mode of describing the world, it would then be a necessary feature of any true description of it. It is also said that amongst the statements which may not have a determinate truth value as a result of their vagueness are identity statements. Combining these two views we would arrive at the idea that the world might contain certain objects about which it is a *fact* that they have fuzzy boundaries. But is this idea coherent? (1978, p. 208, emphasis in original).

Evans' conception of vague objects rests on two ideas; we will unpack each in turn.

First, for vague objects to exist, the world itself must be vague so as to render vagueness a 'necessary feature of any true description of it'. This first idea is a broad definition of ontic vagueness, namely, vagueness that arises from entities or features of the world (whatever they may be) rather than from imprecision in our representation of the world. More specifically, and following Evans' referral to objects with 'fuzzy boundaries', this idea is a "vague-objects view" (VO): a view propounding the idea that, in at least some cases, vagueness arises 'for no other reason' except that there are objects that are themselves vague.[1] For instance, if I were asked to precisely specify the boundaries of my partner's cat Pounce, I would only be able to give an approximation of this. Though I have good working knowledge of determining where Pounce the cat is (after all, I can pick her up in my arms, albeit at my own risk), I could

[1] See Burgess (1989) and Noonan (1990) for similar presentations.

not specify Pounce's exact boundaries if I needed to pinpoint them down to the last hair, or even further, down to the last atom. This is because I lack the means to determine whether to include or exclude some particular atom over others as being part of Pounce, given that some seem to neither be determinately part of or not part of her. To answer this question, as a proponent of VO seeks to do, would be to provide an account of worldly vagueness – to offer a principled means of providing a precise account of features of the world that appear to lack determination.

Keeping in line with Evans' reasoning, we can compare VO to a semantic view of vagueness (SV). Per SV, vagueness arises from our 'mode of describing' objects. Here vagueness is a product of language, arising because our language is insufficiently precise (by design or nature) to denote or express what are otherwise fully determinate and non-vague features of the world. Hence, SV does not see vagueness as a feature of the world; rather, it instead takes some descriptions of the world to be affected by the cloak of vagueness. If SV is correct, vagueness arises because there are deficiencies in how we describe and represent the world, a direct contrast to VO's characterisation of vagueness.

Though VO and SV offer differing accounts of vagueness, they share a common trait: both views characterise the phenomenon of vagueness by identifying certain features of the phenomenon to be determinate, and others to be vague (and so lacking determination). Where the views differ is with respect to what they see as determinate, and what they see as vague. Either the world is vague, or our representation of the world is vague. The contrast between the views is particularly significant when we consider the second idea present in Evans' conception of vague objects, namely, that identity statements may have indeterminate truth-values in virtue of the source of their vagueness. As VO holds that vagueness is derivative from objects or features "of the world", while SV holds that vagueness stems from semantic representations of the world, it follows that a response to Evans' argument against vague identity is dependent on whether one supports VO or SV. These background ideas are important, as Evans intended to show that VO is committed to accepting a knockdown proof to which SV is immune. Indeed, this proof leads to the powerful conclusion that vague identity statements are logically inconsistent.[2] The steps of the proof, as presented by

[2] Or, more precisely, that there can be no true vague identity statements (per Premise B).

Evans, run as follows (here we follow the orthodox interpretation of '∇' as 'it is indeterminate whether'):

(P1) $\nabla(\mathrm{a} = \mathrm{b})$

(P2) $\lambda x[\nabla(\mathrm{a} = \mathrm{x})]\mathrm{b}$

(P3) $\neg\nabla(\mathrm{a} = \mathrm{a})$

(P4) $\neg\lambda x[\nabla(\mathrm{a} = \mathrm{x})]\mathrm{a}$

(P5) $\neg(\mathrm{a} = \mathrm{b})$ [from using Leibniz's Law and from (2) and (4)]

To briefly summarise in words: P1-P5 prove that from a purportedly vague identity statement, assumed for *reductio* (P1), the attribution of the property of indeterminate identity via lambda abstraction (P2, from P1), the determinacy of self-identity (P3), and from the property abstraction of this fact (P4, from P3), one can show by application of the contrapositive of Leibniz's Law that identity is in fact *not* a vague matter (P5, from P2 and P4).[3]

Evans also extended the proof to generate a line that is 'straightforwardly inconsistent with (1) (1978, p. 208)':

> If 'Indefinitely' and its dual 'Definitely' ('Δ') generate a modal logic as strong as S5, then (1)-(4) and presumably Leibniz's Law can each be strengthened with a 'Definitely' prefix, enabling us to derive:

(P5') $\Delta\neg(\mathrm{a} = \mathrm{b})$

In other words, using the modal logic S5, Evans claimed that the distinctness (or non-identity) of two statements can also be proven to be a determinate fact (P5', from P5) – a direct contradiction of P1. Yet, Evans did not explicate how the contradiction is generated (his brief remarks only hint at a full derivation). To see the contradiction, it is important to consider which unstated assumptions are required for the proof to validly proceed. Unpacking these assumptions will demonstrate how Evans intended to derive his contradiction.[4]

[3] A near-identical argument to Evans' P1-P5 is offered by Salmon (1981); indeed, so much so that the argument is sometimes referred to as the 'Evans-Salmon' argument. However, where Evans uses singular terms in his proof, Salmon instead uses variables ('x' and 'y') to represent a 'pair of entities' (p. 241). This enables Salmon to avoid the issue of *de dicto*/*de re* equivalence, to which Evans' proof is susceptible. We discuss this issue in more detail in **1.2**.

[4] Much of the subsequent discussion follows Richard Heck's (1998) lucid analysis of the Evans proof.

1.1.1 – Unpacking Evans' Proof of P5'

Broadly, we need to unpack three main auxiliary assumptions in the Evans proof. First, we need to consider the principles underpinning an understanding of 'Δ' and '∇' as duals (as Evans mentioned, albeit briefly). Second, we need to consider Evans' equivocation of 'Δ' with a distinct but closely related operator "it is definite that" (which we symbolise as 'δ'), as well as its dual "it is indefinite that" (which we symbolise as '■'). Third, as shown by Heck (1998), we need to articulate analogous principles of S5 for a logic of 'determinately'. Explicating these auxiliary principles will show how Evans intended to derive an explicit contradiction in his proof while also revealing how a vague-objects theorist (or 'Indefinitist') may resist the Evans proof.

Given that Evans explicitly stated that 'Δ' and '∇' are duals, he would have likely endorsed the following principles:

(CE) $\Delta A \leftrightarrow \neg\nabla \mathrm{A}$
In words: it is determinate whether A iff it is not indeterminate whether A

(DE) $\Delta A \leftrightarrow \neg\nabla\neg \mathrm{A}$
In words: it is determinate whether A iff it is not indeterminate whether not-A.

(EQ) $\Delta A \leftrightarrow \Delta(\neg \mathrm{A})$
In words: it is determinate whether A iff it is determinate whether (not-A).

The above three principles align with the standard interpretation of 'Δ' as representing the idea that some statement (A) is either true or false, and so has a determinate truth-value. The dual of this symbol would hence represent the idea that some statement *lacks* a truth-value (is neither true nor false, at least on a bivalent framework). These principles seem to match the above interpretation of 'Δ'. Yet, as Heck notes (1998, p. 277), in moving from P5 to P5' Evans equivocated between this interpretation of 'Δ' and a closely related, but distinct, interpretation of the operator.[5] We can shed light on this equivocation by considering Evans' implicit appeal to the following principle:

(T) $\Delta A \rightarrow A$
In words: if it is determinate whether A, then it is the case that A.

[5] This is considered in greater detail in Pelletier (1989). It is worth noting that Evans purportedly retracted this equivocation in personal correspondence with Lewis.

This principle is invalid on the previous interpretation of 'Δ': the statement 'A' having one of either truth-value does not settle *whether* A is true or false, only that A is *either* true or false. In other words, while determinacy in *whether* something is the case establishes the truth of a disjunction (A is either true or false), it does not determine *which* disjunct is the case. Given Evans' reliance on T, then, it appears that he instead wished to utilise a closely related but distinct notion: "it is definite that". This notion stipulates that a statement has a particular truth-value (true or false); in other words, it is definite that a statement has one of the disjuncts of the aforementioned disjunction of truth-values. We can capture this formally in the following manner:

(Tδ) $\delta A \rightarrow A$

In words: if it is definite that A, then it is the case that A.

We can also define the dual of this principle as follows:

(Dδ) $\blacksquare A \leftrightarrow \neg\delta\neg A$

In words: it is indefinite (not-definite) that A iff it is not definitely the case that not-A.

Though Tδ is valid, related analogues of CE and EQ are not valid – if it is definite that A, then it is certainly not true that the negation of A is definitely the case. Evans' equivocation thus gives rise to issues regarding the validity of principles for operators representing "it is determinate whether" and "it is definite that". We cannot conflate the two, nor equivocate, without affecting the validity of any proofs that utilise them.

This issue of Evans' equivocation notwithstanding,[6] the two operators are in fact interdefinable:

$$\Delta A = df \ \ \delta A \vee \delta\neg A$$
$$\nabla A = df \ \ \neg\Delta A$$
$$= df \ \ \blacksquare A \wedge \blacksquare\neg A$$

From the above, we can thus also provide the following definitions:

$$\delta A = df \ \Delta A \ \wedge A$$
$$\blacksquare A = df \ \neg A \ \vee \nabla \mathrm{A}$$

[6] Evans' equivocation between the two operators is a tractable issue, as highlighted by Pelletier (1989, p. 482). Referring to correspondence with Lewis, Pelletier notes that Evans intended his proof to utilise 'Δ'. Given the interdefinability of both operators, Evans' slip can be remedied with sufficient care, as shown in Heck (1998).

Through the careful definition of these distinct operators, we can see that the above definition of "it is indefinite that" provides two distinct instances where a particular statement can be held as indefinite: a statement can be indefinite if it is false (as with the left disjunct), or if it lacks a definite truth-value (as with the right disjunct). By contrast, a statement is indeterminate only if it lacks a definite truth-value. This difference between the two operators plays an important role in contemporary analyses of vague-objects views, and also has implications for claims that certain vague-objects views give rise to problematic instances of "propositional indeterminacy". We will discuss the significance of this issue in more detail in Chapter 3.

The full implications of Evans' proof against vague identity are far-reaching, both for logic more broadly,[7] and for proponents of VO. Here it is important to remember that Evans' conception of vague objects began with the assumption that vague identity is a cogent notion. Following this idea, VO's conception of vague identity thus allows Evans' proof to validly proceed despite the fact that P5' contradicts P1. It follows that we could not endorse VO, and so endorse the existence of vague objects, without accepting a logical inconsistency. Yet, if we follow Evans' proof and accept that vague identity leads to a logical inconsistency, we are left with a pressing question: if identity is a distinctively non-vague notion, then which theses are we committed to hold? Here we can follow Akiba (2014, p. 3559) by formalising Evans' proof into two distinct logical theses. P1-P5 together form a defence of the determinacy of identity (DI), while P5-P5' form a defence of the determinacy of distinctness (DD):

DETERMINACY OF IDENTITY (DI): Identical things are determinately identical, i.e., $\forall x \forall y\ (x = y \rightarrow \Delta x = y)$

DETERMINACY OF DISTINCTNESS (DD): Distinct things are determinately distinct, i.e., $\forall x \forall y\ (\neg x = y \rightarrow \Delta \neg x = y)$

Given that Evans' proof against the thesis of vague identity involved a defence of both DI and DD, it seems that the proponent of vague objects, ostensibly committed to the notion of vague identity, would also be committed to rejecting one of DI and DD. Yet, rejecting the former is widely seen as a dubious endeavour, and rejecting the latter has been rarely attempted (though see **1.4** for a discussion of Akiba's (2014) defence of such a view).

[7] As discussed in Smith (2008a), formal systems in logic often tacitly (if not expressly) rely on notions of identity and distinctness in their construction.

Another available move is to interrogate the premises of the argument that lead to the establishment of the above theses. Using this strategy, we could attack the argument leading to the unwanted conclusion to determine whether a more defensible construction of the desired theses (DI and DD) is available. This seems a worthwhile endeavour. So, to determine if Evans' argument is successful, and to determine whether an alternative argument is tenable, we must first interrogate the logical and metaphysical assumptions that are both explicitly and implicitly endorsed in the argument preceding Evans' defence of DI and DD.

The first assumption we will interrogate is the purported link between vague objects and vague identity. If we are to accept the idea that vague objects can exist only if there can be true vague identity statements, what support is there for this view? It is here that a limitation of Evans' paper becomes apparent. Given the concision of his argument (it is, after all, contained in one page!) it is not immediately clear how Evans intended his proof to undermine the view that there can be (and are) vague objects. Indeed, we are not alone in identifying this: the issue has been pressing enough so as to give rise to various interpretations attempting to elucidate how Evans sought to dispute the existence of vague objects.[8] Hence, to evaluate the link between vague objects and vague identity, we would be well served to first evaluate the merits of such interpretations. Doing so will provide a clearer account of the assumptions that Evans endorses (tacitly or explicitly) in his rejection of vague identity. With clarity on this matter, we will then be able to determine which premises and objections hold force. This will tighten our eventual attack of Evans' argument against vague objects and will ensure that we do not mistakenly ascribe any premises that Evans' argument may not need to endorse to nevertheless be valid.

1.2 Lewis and the Fallacious Modal Equivalence

As misunderstood, Evans is a pitiful figure: a 'technical philosopher' out of control of his technicalities, taken in by a fallacious proof of an absurd conclusion. Rightly understood, Evans endorses neither the bad proof nor the bad conclusion. Instead he is making a good argument in favour of a very different conclusion. – David Lewis, 1988, p. 128.

[8] For instance, see Williamson (2003, fn. 12) for a comprehensive list of papers which offer varied responses to the Evans proof.

Lewis (1988) sought to vindicate Evans' proof from misinterpretation – identifying two problems with the proof that, on first glance, seem to muddy the otherwise powerful conclusion that Evans established. The first problem is that there 'plainly' are such things as vague identity statements. Take Lewis' example: the statement 'Princeton = Princeton Borough'. Here there is no determinate fact of the matter as to whether Princeton denotes Princeton Borough or not. This may be because the name "Princeton" has many candidate referents, none of which can be appropriately identified as the determinately correct referent (for instance, 'Princeton' could denote just the Borough, or perhaps the Borough and surrounding township); or, on another view, because Princeton is a vague object that is indeterminately related to the precise administrative unit that is Princeton Borough.

This first problem connects to the second problem that Lewis identified: a problem relating to a fallacy that is present in the Evans proof. More specifically, the proof makes use of an 'alleged equivalence' between statements of the following forms:

(D) 'It is vague whether … a …', symbolised as ∇ (… a …)
(R) 'a is such that it is vague whether … it …', symbolised as λx [∇(… x …)]a

This equivalence is problematic when we consider two distinct but related issues: the nature, or source, of vagueness, and the rigidity of the terms that feature in vague statements. If the vagueness present in statements like D and R is a result of ontic or worldly vagueness (as is typically attributed to VO), then the names used to determinately designate vague objects permit equivalences between D and R. In cases such as these, an equivalence can be made because the scope of the vagueness operator ('it is vague whether') still captures the relevant aspect of the sentence that would be vague, per VO. For instance, consider the following two statements:

(D') It is vague whether Princeton's boundary includes the Borough and the surrounding Township.
(R') Princeton's boundary is such that it is vague whether the Borough and the surrounding Township are part of it.

In both D' and R', the scope of the vagueness operator includes both the Borough and the surrounding Township – the relevant features of the statement that a proponent of VO would take to be vague features of the world. Relatedly, VO would hold that the name 'Princeton' has a determinate referent: the vague boundary, further contrasting the view that vagueness stems purely from our semantic representation of objects. If

the vagueness present in these cases were to only come from a semantic source, then this would go against a beginning assumption of VO, namely, that the vagueness in question stems (at least in part) from the world itself.[9] For these reasons, proponents of VO take vague identity statements to involve rigid designators: terms that designate the same object (and only that object) in all possible worlds and contexts.[10] Hence, when VO is faced with the alleged equivalence in Evans' proof, it must allow the equivalence to go through, as the scope of the vagueness in both statements is within the vagueness operator present in each statement.

By contrast, if the vagueness is taken to come from a semantic source, then it may not be the case that the referential terms present in both D' and R' are rigid (referring to the same thing). The reason for this, understood broadly, is that semantic indeterminacy (on a supervaluational reading) involves there being multiple precisifications of the vague language used. These precisifications, or alternative ways of understanding the extension (meaning) of a term, have equal claim to be the "correct" or "intended" meaning of a term. This understanding of precisifications is similar to the machinery of possible worlds in modal logic. Possible worlds are ways the world might be; similarly, precisifications are the possible extensions of words in our language. Hence, on a semantic account of vagueness, the terms that feature in vague statements should account for changes in meaning across different contexts (different precisifications). More specifically, the terms featuring in vague statements are shifty; vagueness arises due to imprecision in our referring expressions, as seen with the many possible precisifications for terms being used in vague identity statements.

Applying this idea to the Evans proof, it can be seen that the semanticist can reject a step in the proof to argue that the equivalence should not hold. More specifically, semanticists can reject the step invoking the 'alleged equivalence' (the move from P1 to P2) by stating that the terms involved in each step are shifty (non-rigid). For instance, when considering the sentence D, the scope of the vagueness operator has a domain over a; by contrast, the scope of the vagueness operator in R is of an abstracted property of a, namely, the open formula '…x…'. Such a change in quantificational scope entails a change in the scope of the vagueness in question (or, in other words,

[9] Importantly, this does not necessitate that the VO theorist sees *all* vagueness as derivative from the world. Rather, the thought here is that *at least some* vagueness does not stem from a semantic or epistemic source, and that this can give rise to the vagueness present in cases involving indeterminate identity statements.

[10] See, for instance, LaPorte (2016) for a characterisation and discussion of such terms.

what the vagueness is being attributed to). As the lambda-abstraction involves taking a vague term out of the scope of the indeterminacy operator, this changes the scope of what the indeterminacy operator affects, and by extension, the truth of the sentence. For example, though the indeterminacy operator in D' has scope over a vague referential expression ('Princeton'), there is no guarantee that the sentence resulting from the lambda-abstraction (R') retains the same truth-value, as there is no guarantee that the scope of the indeterminacy operator in the expression includes any vague referring terms (as 'Princeton' appears before introduction of the indeterminacy operator).

Put simply, the proof's contradiction does not hold its force against the semanticist. Semanticists can correctly state that the alleged equivalence in the proof only holds for rigid terms – terms that do not feature in their account of vague statements. While the proponent of VO must hold that terms involved in vague statements are rigid, or non-shifty, the semanticist can quite clearly reject this idea by appealing to semantic indeterminacy (indeed, a calling card of their view).

To make the alleged equivalence in Evans' proof even clearer, we can take the vagueness operator to be analogous or relevantly similar to an operator of contingency, read as 'it is contingent whether'. Lewis himself considered this modal analogue. With reference to Thomason (1982), Lewis highlighted how the equivalence present in Evans' proof is akin to a 'fallacious modal equivalence'. This equivalence can be seen when comparing statements such as 'it is contingent whether Michael is the most caffeinated person in the room' and 'the most caffeinated person in the room is such that it is contingent whether they are Michael'. The former statement is certainly true, as though I drink a lot of coffee, it does not match the impressive imbibing practices of my friend Tom. Depending on whether or not Tom (or another heavy coffee drinker) is in the same room as me, it is a contingent matter as to whether I am the most caffeinated person in the room. On the other hand, Q is false; it implies that object identity is a contingent, rather than necessary, notion.[11] This modal fallacy is important, as it relates to the distinction between the statements D and R.

[11] As proven by Saul Kripke (1971), by David Wiggins (1965), and even earlier, Ruth Barcan Marcus (1947) in the following derivation:
(1) $\forall x \Box \mathrm{x} = \mathrm{x}$
(2) $\forall x \forall y\ (x = y \rightarrow (\Box \mathrm{x} = \mathrm{x} \rightarrow \Box \mathrm{x} = \mathrm{y}))$
(3) $\forall x \forall y\ (x = y \rightarrow \Box \mathrm{x} = \mathrm{y})$

When considering the above equivalence in the context of Evans' paper, we can see that D and R are representative of P1 and P2 in Evans' proof respectively.[12] While the former premise shows a purportedly vague identity statement, the latter premise shows the application of lambda abstraction between a variable (x) and a name (a) – representing a property being attributed to an object respectively. As the move between P1 and P2 tacitly assumes that the use of lambda abstraction is legitimate, we need to consider whether Evans' application of lambda abstraction is legitimate on all readings of the proof.

The main force of Lewis' paper, and his reading of Evans' proof, is that it highlights how the alleged equivalence between D and R is invalid when involving shifty referential terms.[13] If we wished to maintain that the phenomenon of vagueness is semantic in nature (and so vague statements could make use of non-rigid designators), then we could not also legitimately utilise lambda abstraction in these shifty semantic contexts. In other words, on such a reading we would not support the inference that vague identity statements must entail vagueness in a lambda-abstracted property. Further, as the Evans proof does not validly go through for the semanticist, its conclusion (that vague identity statements lead to contradiction) is also invalid on a semantic reading of vagueness.

The significance of this conclusion relates to our earlier analysis of the first problem that Lewis identified, namely, that Evans' proof seems to reject the existence of vague identity statements despite there being a seeming plenitude of potential examples of such statements. Yet, for Lewis this is in fact not a problem at all:

> The correct interpretation is that Evans trusts the reader – unwisely! – to join him in taking for granted that there are vague identity statements, that a proof to the contrary cannot be right, and that the vagueness-in-describing view [semantic view] affords a diagnosis of the fallacy. His point is that the vague-objects view cannot accept this diagnosis (1988, p.129).

So, taken on a superficial reading, Evans' proof seemed to be endorsing the easily refutable view that there are no vague identity statements. However, as Lewis' careful

[12] For the sake of clarity, P1 and P2 are reiterated here:
(P1) $\nabla\ (\mathrm{a} = \mathrm{b})$
(P2) $\lambda x[\nabla(\mathrm{x} = \mathrm{a})]\mathrm{b}$

[13] A point used to motivate a number of responses to Evans' proof, each offering arguments invalidating the Evans proof's use of property abstraction. See, for instance, Barnes (2009), Garrett (1988), Noonan (1982), Thomason (1982).

analysis demonstrates, this is only the case if one takes the Evans proof as valid. While it seems that a proponent of VO must take the argument to be valid (given its commitment to using rigid designators in vague statements), a semanticist can identify the argument as invalidly applying lambda abstraction to shifty semantic contexts. With the Lewisian interpretation in hand, it becomes apparent that while the Evans argument is not a problem for semanticists, it remains a problem for vague-objects theorists.

The question remains: should we fully accept the Lewisian interpretation of Evans' argument? Given that Lewis referred to personal correspondence with Evans to support his claim (1988, p. 130), it seems that his interpretation is, at the least, consistent with Evans' intended argument. Though consistent, the argument is not immune from further scrutiny. To start, it is worth considering the following line offered by Lewis: 'what is in trouble is the vague-objects view combined with the view that vague identity yields identity statements with indeterminate truth value' (1988, p. 129). As one would expect, this conjunction is not a problem for Lewis and Evans; they could simply reject the first conjunct (and by extension, the troublesome conjunction) by adopting a semantic, rather than ontic, account of vagueness. What of the vague-objects theorist? They cannot reject the first conjunct (to do so would go against the very view they seek to defend!) Yet, accepting the first conjunct would seem to carry a commitment to the second conjunct, assuming that the existence of a vague object with fuzzy or indeterminate spatial boundaries would lead to there being true vague identity statements of such objects. Indeed, without a substantive picture demonstrating how a vague object could nonetheless have definite identity conditions, it is plausible to follow Lewis in taking vague objects to entail there being true vague identity statements. So, as the first conjunct plausibly leads to the second conjunct, and as the second conjunct leads to contradiction, it follows that the vague-objects theorist is left with a troublesome conjunction with two conjuncts that she must seemingly accept.

What recourse is left for the proponent of vague objects? To start, Lewis' conjunction (and Evans' argument by extension) follows Premise A in suggesting that the existence of vague objects carries commitment to vague identity (following Premise A, as in **1.1**). Yet, while Premise A has been postulated, no positive argument has yet been made to support it. What validates the thought that an account of vague objects *must* lead to vague identity? Neither Lewis nor Evans offered a positive reason for why a vague-objects theorist ought to be committed to such a view. Rather, what their arguments

establish is that *if* a proponent of vague objects were committed to vague identity, *then* they would be led to a troublesome conclusion.

Here the force of the Evans argument (supplemented by Lewis' analysis) is clear: the argument offers a *restriction* of available options for the proponent of vague objects. Specifically, their options are to either refute one of the steps in the Evans argument (thereby preventing commitment to the troublesome conclusion), or to conceive of vague objects in such a way that does not give rise to vague identity (sidestepping the argument). The former option is what Jonathan Simon calls the "concessivist response" to the Evans argument (2014, p. 42) – a response seeking to identify the assumptions that can be relinquished to successfully sidestep Evans' argument. We will engage with such responses in **1.3**. The latter option is what David Over calls a 'classical way' of talking about vague object and their classical, definite identity'. On this "classical view", which holds that there may be definite identity conditions for vague objects, one would need to accept that 'there is a definite domain of vague objects, that these definitely exist, and that there is a definite number of them, though they may have more or less "fuzzy" edges' (Over 1989, p. 98). The viability of this latter approach rests on one's conception of what a vague object Is – the metaphysical commitments of a vague-objects view – and how these commitments support the principles outlined by Over.

We will explore the prospects of defending the idea that there can be definite identity conditions for vague objects in Chapter 2, where we develop a novel conception of vague objects. Regardless, whether one adopts the 'concessivist response' or 'classical view' (or another view), the burden of proof still rests on the vague-objects theorist. While the semanticist has a tenable reply to the Evans proof, the vague-objects theorist must do extra work to present a viable alternative to the semantic solution. With this in mind, what calls for our immediate attention is how the Evans argument restricts the vague-objects theorist in other ways. If the burden rests on the vague-objects theorist, how heavy is that burden? In other words, given the force of the Evans argument, what are the assumptions that a vague-objects theorist ought to maintain?

1.3 Logical Commitments of a Vague-Objects View

The conclusion of Evans's argument is not an unpalatable one, so one cannot reasonably reject otherwise plausible assumptions purely in order to resist it – Benjamin Curtis and Harold Noonan, 2014, p. 306.

To respond to the Evans argument, the vague-objects theorist needs to unpack the seemingly plausible assumptions that underpin it. These assumptions are akin to theoretical commitments – choices to consider when constructing an account of what a vague object is. Robert Williams (2008b, p. 136) usefully articulates many of the assumptions that viable solutions to the Evans proof ought to uphold, paraphrased as follows:

A1: Classical logic should be preserved.
A2: The logic of 'indeterminately' is to be S5; a consequence of this is that if something is indeterminate, it is determinate that it is indeterminate.
A3: A solution to the Evans argument should target identity statements that are vague in virtue of "ontic" reasons, rather than from semantic indecision or ignorance.
A4: Leibniz's Law (the principle of the indiscernibility of identicals) should be upheld.
A5: The solution should understand properties in a 'thin' or 'merely abundant' sense (as per Lewis (1983)).

Though Williams offers his own arguments supporting these assumptions, their importance calls for a discussion of our own. Our discussion will show that each desideratum ought to be upheld, and further, that they should be incorporated into a defence of vague objects.

To start, A3 can be readily endorsed. A vague-objects theorist cannot rely on identity statements that are vague in virtue of semantic indeterminacy (as shown in **1.2**). We also readily accept A4 and A5, given that they enable the strongest form of Evans' argument to go through (making the task for the vague-objects theorist more difficult, if anything). We will, however, address issues relating to A4 and A5 in Chapter 2, where we discuss the metaphysical assumptions of a workable vague-objects view.[14] Hence, for our current purposes, we will focus our discussion on A1 and A2 to elucidate the "logical" commitments of a vague-objects view.

1.3.1 On Classical Logic (A1)

Classical logic is simple and attractive, and given that many of our best extant theories presuppose classical logic, we have a lot of reconstruction to do if we give it up. For these reasons, an 'innocent until proven guilty' methodology is attractive (that is, you

[14] In saying this, it is pertinent to note that there are instances where logical commitments do affect our understanding of Leibniz's Law, and so agnostically accepting it is not an entirely innocent move. To take one case as an example, Parsons and Woodruff's (1995, 2000) many-valued logic for vagueness gives rise to the possibility that the contrapositive of Leibniz's Law does not necessarily hold.

should keep classical logic unless you're forced not to). – Elizabeth Barnes and Robert Williams, 2011a, fn. 16.

The simplicity and richness of classical logic is attractive, and its effectiveness as a conceptual tool is undeniable. Yet, a tool is only as effective as its use. It would do us no good to be equipped with a powerful hammer for a job requiring a screwdriver. While at first glance the vague-objects theorist would be well served to incorporate the machinery of classical logic into their account (even innocuously, as per Barnes and Williams), we should be prudent and first determine the need for such a logical toolkit, and by extension, the power of classical logic as a tool for explaining vague phenomena.[15]

An example may help to illustrate this point. Supervaluationism, a popular theory in the vagueness literature, offers a substantive and explanatorily powerful account of vague phenomena. Even so, supervaluationism rejects some of the inference rules of classical logic (such as contraposition, conditional proof, *reductio*, and proof by case).[16] Its success and widespread adoption has come in spite of its revisionary nature, as while supervaluational frameworks respect classical logic,[17] they make use of non-classical semantics.[18] Given its popularity, supervaluationism has set a clear precedent for theoretical revision. Yet, some have objected to the revisionary nature of supervaluationism and its incompatibility with certain classical inference rules. For instance, Williamson claims that the aforementioned inference rules are 'vital' to natural deduction, and so their rejection is too expensive a cost for a theory of vagueness.[19]

In response to Williamson, Keefe argues that these revisions are not necessarily problematic or damaging for the supervaluationist, suggesting that we should expect the introduction of "non-classical" notions such as 'definitely' to produce results that do

[15] Broadly, we take classical logic to include extensions of extensional first-order logic (such as the addition of modal-like operators like 'determinately'). We contrast this with non-classical logic in the sense that non-classical logic involves *revisions*, and not merely extensions, of classical logic. So, we are happy to accept extensions (such as new expressions to govern formal inference procedures), so long as this does not involve revision to existing rules.

[16] See Williamson (1994, pp. 151-152).

[17] Supervaluationism accepts the Law of Excluded Middle, as every admissible interpretation of a statement in supervaluationism is either true or false.

[18] Supervaluationism rejects bivalence, as not every statement is either true or false (supervaluationism identifies truth and falsity with supertruth and superfalsity respectively - as some statements come out as neither supertrue nor superfalse, bivalence fails).

[19] See Williamson (1994, pp. 146-153).

not support classical inference rules.[20] Specifically, supervaluationism only needs to revise classical inference rules when 'definitely' operators are present – classical inference rules remain valid in the absence of a definitely operator. As a result, or so Keefe argues, a plausible response to Williamson's objection is to construct a supervaluational apparatus with 'reasonable' modifications to classical inference rules in the presence of definitely operators. One justification for such a modification may be to reject classical inference rules on the grounds that vague statements merit the use of non-classical reasoning, and hence motivate deviant inference rules. Perhaps deviance should be expected when dealing with vague affairs.

Williams offers a different response to Williamson's objection, arguing that alternative supervaluational frameworks can be constructed to altogether avoid the damaging revisions to classical inference rules. On Williams' preferred supervaluational framework, consequence is identified with global consequence (rather than local consequence, which itself would suffice to avoid revisions to classical inference rules). Using a toy model, Williams demonstrates how a 'natural generalisation of the classical characterisation of logical consequence' (2008a, p. 194) can be used to construct a framework that does not require revision to these inference rules.[21]

Regardless of one's stance on the revisionary nature of supervaluational approaches to vagueness, and the purported success of Keefe and Williams' responses,[22] the takeaway from the above discussion is to highlight an important lesson: revisions to classical logic and classical semantics typically come at a theoretical cost. The significance of these costs is a matter of importance for any theory of vagueness, as seen in Keefe and Williams' extended defences against Williamson's criticism. So, to determine the general utility of classical logic in constructing a vague-objects view (and to determine the significance of desideratum A1), it is important to consider whether there is sufficient motivation to accept classical logic as a base theoretical commitment, or whether the costs of alternative, non-classical views are worth their expense.

[20] See Keefe (2000, pp. 174-181)

[21] See Williams (2008a) for a full defence of the view that a suitably constructed supervaluational framework can avoid 'damaging' revisions. *Contra* Williams, see Jones (2011) for an argument against Williams' proposed framework.

[22] For what it is worth, we broadly agree with Williamson in placing the onus of proof on the supervaluationist to justify their revisionary frameworks and their application towards instances of vagueness. A resolution to the debate relies on determining whether the responses offered by supervaluationists succeed, or whether they are vulnerable to further criticism (and whether comparatively viable alternatives are available).

Classical logic is relatively inexpensive in its theoretical costs; however, adopting it as a framework for a theory of vagueness may eliminate some potentially fruitful avenues of theoretical construction. For instance, in accepting bivalence, proponents of classical logic must work with a relative paucity of truth-values – this may be burdensome when tackling borderline cases that, at a glance, seem to require at least a "third" truth-value.[23] Indeed, the rejection of bivalence has allowed various many-valued logics to fruitfully tackle, among other problems, the Sorites paradox.[24] Yet, the diverse toolkit of many-valued logic comes with its own costs: many-valued logics produce what Williams consider '*prima facie* unacceptable results' (2008b, p. 136). One such result is that many-valued logics appear to countenance acceptance of some contradictions as 'less than completely false'. In light of this, Williams calls for the use of classical logic on 'broadly methodological grounds' (2008b, p. 136); classical logic requires no revision to existing and broadly accepted principles, leaving our metaphysics and ontology relatively unscathed by our chosen logic.

On balance, a sufficiently motivated opponent may, quite rightly, point at the similar "theoretical costs" of refusing to consider non-classical logics (what of the explanatory power that many-valued logics offer!) Here the dialectical standoff is apparent: one can use classical logic to minimise theoretical reconstruction elsewhere; or one can wield non-classical logic and its more expansive toolkit, though at the expense of committing to a broader revisionary project. For the purpose of this thesis, which seeks to vindicate the existence of vague objects (a project that has historically been thought to require the use of revisionary logics and metaphysics to succeed), we adopt the former approach: do not revise without necessity. Specifically, our attempt to defend a vague-objects theory will abide by classical logic and will use a logical framework that does not require commitment to theoretical construction elsewhere. While other accounts may fruitfully apply non-classical logics, or non-classical semantics, this thesis will determine the extent to which the existence of vague objects can be vindicated within a classical framework. We therefore endorse desideratum A1 – classical logic is to be preserved.

[23] As seen in Broome (1984).

[24] See, for instance, Smith (2008b) and Hyde (2018).

1.3.2 On the Logic of 'Indeterminately' Being S5 (A2)

I am granting that the formal argument Evans sets out is one which must be accepted by [a vague-objects theorist]. The question is whether [they are] committed to the truth of a contradiction – Richard Heck, 1998, p. 281.

Williams defends the use of S5 on the basis that 'the relevant instances of the characteristic S5 axiom are independently plausible' (2008b, p. 138). While it remains to be seen whether S5 is *the* correct logic for 'determinately', its usefulness in analysing vague phenomena has led to its widespread adoption. Though it may be fruitful to consider instances where determinately vague identity is relevant for analysing vague phenomena, it is more important for our purposes to determine how the characteristic S5 axiom is used in the Evans proof, and also to determine whether it is necessary both for the proof's success and for a vague-objects theorist to accept. By addressing both of these questions, we can clarify how S5 is to be applied to the Evans' proof, and also provide a more substantial defence of desideratum A2.

Evans' move from P5-P5' relies on the use of S5 to produce the desired result (an explicit contradiction in relation to P1), and we have elaborated how Evans intended to achieve his contradiction in **1.1**. Yet, a significant feature of Evans' argument remains absent: while it has been shown that vague-objects theorist ought to accept the formal validity of Evans' proof, Evans does not explicate *why* S5 is the correct logic of 'determinately'. Though strengthening each premise of the proof with a 'definitely' prefix allows a formal contradiction to be made between P5' and P1, no justification is offered (implicit or otherwise) for why the axiom that is characteristic of S5 is one that a vague-objects theorist should endorse in broader contexts outside of responding to the Evans argument. In other words, although S5 has been established as a suitable logic for 'determinately' for the Evans proof, it is not yet clear why S5 is a suitable logic for vague-objects theorists more broadly. Given this, it is important to examine a key consequence of accepting S5 as the logic for 'determinately', namely, that accepting S5 carries commitment to the rejection of higher-order vagueness. As this has significant implications for considering the logical and metaphysical base of any vague-objects view, it merits explicit analysis.

If one wished to resist Evans' proof, then they may also wish to consider his use of S5 as unjustified (as we will see in **1.4**). However, in using S5, it is important to first determine the relevant axiom in the context of determinacy and indeterminacy. Given

P5 and P5', and Evans' reference to S5 in his proof, we can take his proof to require the following principle:[25]

(5Δ) $\nabla A \rightarrow \Delta\nabla A$

This is an analogue of the standard S5 axiom, which stipulates that if something is possibly the case, then it is necessarily possibly the case.[26] Following this, if it is indeterminate whether something is the case, then it is determinately indeterminate whether this is so.

According to Heck, a possible justification for 5Δ comes from consideration of a three truth-valued system (True, False, Neither), where 'by definition, '∇A' must be either True or False' (1998, p. 283). On such a view, '∇A' is False when A is True or False, and is true when A is "Neither". A drawback of accepting 5Δ is that it eliminates the possibility of higher-order vagueness (as all indeterminate facts are determinate, and so are definitely true or false). In other words, accepting 5Δ carries the implication that there is no vagueness with respect to indeterminate matters – all indeterminate states of affairs will not fail to either be definitely true or false. This has a significant upshot: if the vague-objects theorist wants to accept the Evans proof as valid, and if they also wish to accept the relevant S5 axiom, then they cannot countenance higher-order vagueness. Some take this to be simply implausible, and outright reject the axiom present in Evans' argument.[27] The rejection of higher-order vagueness is a purportedly significant blow for the vague-objects theorist – here we would be well-served to consider why this is the case.

One line of thought comes from Williamson (1999), who states that the motivation to accept higher-order vagueness is the same for accepting first-order vagueness – the classificatory difficulties arising from higher-order vagueness are the same that arise for first-order vagueness. In other words, the very issues that 'led us to recognise the problem of vagueness in the first place' (Williamson 1999, p. 127) equally apply to issues of higher-order classification. This is a pertinent worry for semantic and

[25] See Heck (1998, pp. 282-283) for a full articulation of Evans' proof involving S5, including a defence of the relevant distribution axiom required.

[26] Or, following Routley and Montgomery's system (1966), this is a standard truth in contingency logics for S5. The characteristic axiom for the system equivalent to S5 (S51a) is $\neg\nabla\nabla p$. So, given the contingency version of EQ from **1.1** (T1 for Routley and Montgomery), this is equivalent to $\neg\nabla\neg\nabla p$, which is equivalent to $\Delta\nabla p$. By classical logic, this entails (5Δ). So, no facts about contingency are themselves contingent (regardless of the contingency or non-contingency of some p).

[27] See Garrett (2014).

epistemic views, as the imprecise or incomplete means by which language is regimented, or the “unknowability” of sharp delineations of reality, may give rise to higher orders of imprecision in relation to borderline cases. This consideration has given rise to the widespread countenancing of higher-order vagueness. Yet, the same motivation does not seem to apply to views that see the world as being the source of vagueness. Specifically, as vague-objects views can take vagueness to arise from non-referential and non-epistemic sources, they do not share the motivation to accept higher-order vagueness as a general maxim. An alternative motivation would instead come from an answer following question: can it be a vague matter that the world is itself vague? If this were the case, then this would seem to be a matter best explained not by the semantics of logic of a vague-objects view, but rather, from the source of the vagueness, namely, from the relevant metaphysical framework being used. It thus seems plausible that the source of vagueness for any view, particularly for a worldly view of vagueness which countenances vague objects, would need to be established to properly evaluate the claim that indeterminate worldly matters can possibly be vague.

Some may wish to argue that logic and semantics alone should not force the acceptance or rejection of otherwise important notions (such as higher-order vagueness). Indeed, Heck argues that ‘surely’ worldly vagueness is committed to the idea that worldly vagueness cannot be eliminated by the insertion of a logical operator (1998, p. 284). If vagueness is in the world, and is derived from such things as properties, relations, states of affairs, or other worldly items, then what in the logic of S5 warrants the eradication of higher-order vagueness in virtue of a logical axiom? In relation to the supposed ‘ineradicability’ of vagueness, Heck offers the following argument: if vagueness is to come from “the world”, then it seems to follow that semantic and logical principles should not, in and of themselves, be able to “eradicate” vagueness stemming from a worldly source (1998, pp. 282-285). No semantic assumptions or logical assumptions are purportedly strong enough to eliminate the “worldly” vagueness that is characteristic of vague objects. If we are to start with the viewpoint that the “worldly items’ are in and of themselves vague (whatever their metaphysical status may be), rather than statements/sentences representing or referring to these worldly items, then it would follow that independent logical and semantic assumptions should not undermine the starting point of the view (unless they are necessary commitments of the view). Accepting the analogue of S5 in the context of determinacy, 5Δ, implies that any statement can be strengthened with a definitely

prefix and so can “eradicate”, or remove, the vagueness present in what is being predicated. Yet, the starting assumption of the vague-objects theorist is that there is vagueness in that to which a predicate applies (such as an object). This gives rise to a tension between the seeming logical applicability of 5Δ to vague contexts and the starting assumption that objects themselves are vague items.[28]

Accepting S5 as the logic of determinately is a move that can be resisted; why should it be that a vague-objects solution requires the adoption of a particular logical principle? We concede this point, though we note that it is one of a number of options with respect to accepting S5 as a logic for determinacy. Three options come to mind. First, one could accept the use of S5 and take the Evans proof to be legitimate. This option calls for an explanation of the purportedly problematic rejection of higher-order vagueness within the metaphysics of the view. Second, one could accept the use of S5, but attempt to account for higher-order vagueness through the use of a particular logical framework (as seen in Barnes and Williams (2011a)). Third, one could reject the use of S5, and so accept a non-S5 logic for determinacy. This would require a reconstruction of the Evans proof in order to maintain the validity of the idea that there cannot be such a thing as vague identity; alternatively, such a view might require the development of an account endorsing alternatives to “strict” or purely determinate notions of identity.

Perhaps unsurprisingly, we adopt the first option: accept S5 (and the validity of Evans’ proof) while rejecting higher-order vagueness. This provides the most charitable reading of the Evans proof and only makes our project of defending the existence of vague objects more difficult. The cost of accepting this reading, and accepting the validity of Evans’ proof, is twofold: the proponent of vague objects cannot rely on vague identity (lest they must explain away the notion’s apparent inconsistency); second, they must offer a plausible explanation for why a worldly view of vagueness should not countenance higher-order vagueness. In other words, there needs to be a plausible metaphysical reason for supporting the eradication of vagueness beyond the “first-order” in our defence of vague objects. We will discuss this matter in Chapter 2, where we begin construction of a more substantial metaphysical base for the existence of vague objects.

[28] For Heck, this *begs the question* against the vague-objects theorist (1998, p. 284).

With the above in mind, we accept desideratum A2: the logic of 'determinately' is to be S5. Notably, this has not required a revision to our logic or semantics; instead, it has required the concession that much of the explanatory burden has been shifted to the realm of metaphysics to determine how conditions of identity are possibly affected by our metaphysical apparatus. Before we unpack these metaphysical commitments however, we would be well-served to first evaluate a contemporary attempt to resisting the Evans proof that does not rely on an underlying metaphysical framework to do this explanatory work. By considering the upshot of such a view, we will be better positioned to evaluate the virtues of using metaphysics to vindicate the existence of vague objects.

1.4 On the Determinacy of Distinctness

Since, again, this paper's goal is a defense of indeterminate individuation, it should be sufficient to present a plausible countermodel to Evans's argument – Ken Akiba, 2014, p. 3562.

The previous sections considered how Evans defended Premise B, the idea that there are no such things as true vague identity statements. On the Evansian view, this idea is connected to the existence of vague objects (per Premise A). So, if we wish to defend the existence of vague objects, then we should also consider potential avenues of vindication from arguments that explicitly resist Premise B and Evans' *reductio*.

One such argument comes from Akiba (2014), who rejects Evans' proof of DD (the determinacy of distinctness). Specifically, while Akiba accepts Evans' proof of DI (the determinacy of identity), he rejects the logical move from P5-P5', as it relies on the logic of indeterminacy being S5. The main force of Akiba's argument comes from his presentation of a model, itself relying on a precisificational framework, which does not comport with S5. Akiba's model rejects Evans' use of a symmetric accessibility relation for the logic of determinacy; instead, Akiba defends the idea that an anti-symmetric relation can be used for the logic of determinacy. By presenting a model with an anti-symmetric accessibility relation, Akiba seeks to show that S5 need not be "the" correct logic for determinacy. In short, Akiba's argument aims to demonstrate that one can cogently model indeterminate distinctness, and hence that one can resist Evans' proof of the determinacy of distinctness.

This section unpacks Akiba's precisificational framework and considers the reasoning underpinning his use of an anti-symmetric accessibility relation. We then evaluate

Akiba's argument, concluding that while he offers a plausible framework with which to attack Premise B, his argument nevertheless faces problems relating to the domain of objects invoked in the argument. Further, we argue that even if we accept Akiba's argument, it nevertheless achieves a narrow argumentative scope, merely showing that indeterminate distinctness is *not impossible*. Yet, as we will argue, the possibility of indeterminate distinctness does not entail the stronger conclusion that vague objects give rise to, or need to rely on, the notion of indeterminate distinctness. Indeed, we will show that Akiba's argument is not convincing for a vague-objects theorist, as it rests on logical assumptions that the vague-objects theorist has good reason to reject. We will thus evaluate Akiba's argument within the broader dialectic of vindicating a vague-objects view. Specifically, while one can reject Evans' proof on logical grounds,[29] as is the case with Akiba's novel argument, this does not necessitate the adoption of the logical assumptions underpinning such arguments, particularly when preferable alternatives are available. To conclude, we argue that a vindication of vague objects ought to target Premise A of Evans' argument, rather than attempt to repudiate Premise B.

1.4.1 Akiba's Precisificational Framework and Countermodel

… while the proof of DI relies only on logical and/or unproblematic principles and is thus irrefutable, the proof of DD relies crucially on a problematic non-logical, modal principle, Axiom B: $\Delta \neg p \rightarrow \Delta\neg\Delta p$ – Ken Akiba, 2014, p. 3560.

Akiba begins his argument by distinguishing "usual" talk of 'indeterminate identity' from what he calls 'indeterminate individuation'. Akiba's motivation is to emphasise the difference between DI and DD, as briefly discussed in **1.1** (2014, p. 3559):

DETERMINACY OF IDENTITY (DI): Identical things are determinately identical, i.e., $\forall x \forall y\ (x=y \rightarrow \Delta\ x=y)$

DETERMINACY OF DISTINCTNESS (DD): Distinct things are determinately distinct, i.e., $\forall x \forall y\ (x\neq y \rightarrow \Delta\ x\neq y)$

Indeterminate individuation represents the conjunction of DI and DD: if two objects are indeterminately individuated, then they are neither determinately identical nor determinately distinct from one another. This can be formalised as: $\neg\Delta a=b \wedge \neg\Delta a\neq b$.

[29] As others have done in a number of ways. For instance, Parsons and Woodruff (1995) present a three-valued logic as part of their resistance to Evans' proof. For a non-classical "fuzzy" interpretation of identity, see Priest (1998).

This conjunction underpins the basis of Akiba's argument. While Akiba accepts that identity is always a determinate relation,[30] he defends the idea that distinctness may, in some cases, be an indeterminate notion. In essence, and remembering the background argument of the Evans paper, Akiba rebrands the notion of vague identity as the notion of indeterminate individuation. This constitutes a reformulation of Premise B of Evans' argument:

(B') There is no such thing as vague (indeterminate) individuation, lest we face accepting a logically inconsistent notion.

This reformulation of Premise B is what Akiba attacks. If it is possible to present a model demonstrating the logical consistency of indeterminate individuation, then there may be room to resist Premise B'.

Akiba presents a precisificational framework in support of his argument – a framework providing the conceptual space for indeterminate distinctness. His precisificational framework interprets the accessibility relation as 'is a precisification of', which is meant to represent the 'process of precisification' (2014, p. 3559). Accessible worlds include states of affairs that are not present in the actual world; namely, they include precise states of affairs that are not present in the actual, vague world. Hence, w' is accessible from w iff w' is a precisification of w. Strictly speaking, w is a precisification of itself, but it is not a 'proper precisification' as there is no difference in state of affairs (no precisification has occurred).

An example may help to clarify the utility of the framework. Let us imagine in the actual world (@) that Mt Kilimanjaro (MK) is a vague object with a vague spatial boundary, where an electron, Sparky (s), is neither determinately part of, nor determinately not part of, MK.[31] On Akiba's precisificational framework, we can consider an accessible world (w1) where Mt Kilimanjaro determinately includes s – call this precise aggregate MK+. Hence, in w1, MK = MK+; Mt Kilimanjaro is identical to the aggregate of electrons which include Sparky as a determinate part. Here a 'process of precisification' has occurred between the actual world @ and w1. As a result, while s is a borderline part of MK in @, it is determinately a part of MK in w1. Similarly, there may be an accessible world (w2) where Mt Kilimanjaro determinately does not possess s. In this world,

[30] Though there may be instances of indeterminate identity *statements*, as discussed in **1.2.**

[31] Here we present the notion of a "vague object" with a "vague spatial boundary" and Sparky as a "borderline part" agnostically. See Chapter 2 for a more detailed characterisation and disambiguation of the terms.

MK = MK-, where MK- is the precise aggregate of which Sparky is determinately not a part. Our toy example thus involves three worlds: @, the actual world where MK has Sparky as a borderline part; w1, where MK = MK+; and w2, where MK = MK-. While all three worlds contain the precise aggregates MK+ and MK-, Mt Kilimanjaro is identical to, and distinct from, different aggregates in each respective world.

As Mt Kilimanjaro is identical to different sets of electrons (and various other constitutive matter and properties) in different accessible worlds, this entails that Mt Kilimanjaro has modal properties that the precise aggregates themselves lack. For instance, while Mt Kilimanjaro may be identical to MK+ *at a world* (w1), this does not entail that Mt Kilimanjaro is strictly identical with MK+. Quite clearly, Mt Kilimanjaro has different modal properties to MK+, including the property that it is not identical to MK+ in all accessible worlds. While MK+ is always identical to MK+ in other possible worlds (such is the nature of a precise aggregate), Mt Kilimanjaro is not always identical with MK+. Hence, both MK and MK+ possess distinct modal properties, as does MK compared to any precise aggregate. Furthermore, as MK is determinately distinct from MK+ in @, this also entails that MK is distinct from MK+ even at w1; while MK is identical to MK+ *relative to the world* w1, MK does not possess the same modal properties as MK+.

In short, identity at a world is not the same relation as identity *simpliciter*.[32] Even if MK is identical to a precise aggregate at a world (where the "process of precisification" has occurred), MK is nonetheless determinately distinct from any precise aggregate. The modal properties possessed by MK and those possessed by each precise aggregate are distinct, giving rise to clearly different identity conditions. Yet, if MK is determinately distinct from any precise aggregate, then we may wish to ask: wherein lies the conceptual room for indeterminate distinctness?

Akiba carves out conceptual room for indeterminate distinctness using a particular interpretation of the accessibility relation. His interpretation holds that the accessibility relation is transitive, reflexive, and anti-symmetric. Here we are interested in the latter property. While Akiba does not intend to 'determine exactly' the logic of determinately (2014, p. 3563), he nonetheless provides a quite explicit justification for an anti-symmetric accessibility relation. On Akiba's account, if determinacy involves the

[32] This may seem peculiar – if "identity at a world" is not equivalent to identity *simpliciter*, then we may wonder whether this represents a "precisification", given that a precise world should agree with an imprecise world with respect to what is determinate at the imprecise world (and given that the necessity of identity entails the determinacy of identity).

process of precisification, and if w1 is a proper precisification of w, then it follows that w cannot itself be a proper precisification of w1. The process of precisification entails that one world has undergone this process, and so has been made "more precise" than another world. Furthermore, if two worlds have access to the other, then on the current framework this entails that they are the same world, for no 'genuine precisification' would have occurred (2014, p. 3563).[33] These two considerations support an anti-symmetric accessibility relation. As a consequence, Akiba argues that 'there can be a precisification at which for some x and some y, $x \neq y \wedge \neg \Delta x \neq y$' (2014, p. 3564). This can be applied to the Kilimanjaro case; if Akiba's model is correct, then the following statement may hold: $MK \neq MK+ \wedge \neg\Delta\, MK \neq MK+$.[34]

Let us consider the reasons grounding MK's distinctness from MK+, namely, their aforementioned difference in modal properties, as it is these reasons that Akiba considers in defending the cogency of indeterminate distinctness (2014, pp. 3564-3658). For Akiba, while this difference in modal properties holds at @, it does not necessarily hold in other worlds, such as at w1. A more precise example may help. From @, there is an accessible world where MK does not equal MK+ (w1) – a precisification of @. Yet, from w1, it is possible that the only worlds accessible from w1 all agree that MK is identical to MK+ (where access is governed by the 'is a precisification of' relation). In other words, all precisifications of w1 may agree that MK is identical to MK+, and hence from w1 there may be no accessible worlds (precisifications) specifying the relevant modal difference between the two objects, as @ is inaccessible from w1 on this model.

As a result, or so Akiba argues, it is *not* the case that MK is determinately distinct from MK+, as there are some worlds (such as w1) where MK is determinately identical to MK+, rather than determinately distinct. There may be distinctness *at a world*, but the grounds of this distinctness are not available to other worlds such as w1. This purportedly gives rise to the possibility of two objects being indeterminately distinct, as the reasons grounding MK's distinctness from a precise aggregate are not present at all worlds where MK exists. From @, MK is determinately distinct from MK+, yet from the perspective of w1, all accessible worlds agree that MK is identical to MK+. Therefore, given the existence of two worlds with differing interpretations of MK's distinctness from MK+, this purportedly demonstrates that MK can be indeterminately

[33] If two worlds can both access the other, then neither is "more precise" than the other. Hence, given this shared access, on Akiba's framework both worlds are identical to each other.

[34] In words: MK is distinct from MK+, and it is indeterminate that MK is distinct from MK+.

distinct from MK+, as differing worlds disagree with respect to the determinacy of MK's distinctness from MK+.

We see two problems with this argument. The first relates to the above presentation of modal properties, and the second relates more broadly to the dialectic of vindicating vague objects. We will unpack each problem in turn.

To see the first problem, we can follow Smith (2008a, p.7) and interrogate the domain of objects stipulated in the model.[35] Akiba motivates the cogency of indeterminate distinctness by considering MK and its identity to a precise aggregate *at a world*. On Akiba's model, indeterminate distinctness arises when MK is determinately identical to MK+ *at a world*, and determinately distinct from MK+ *at another world*. We can nonetheless argue that the distinctness of MK and MK+ holds at *all worlds*. This is because the domain used in the example still specifies two objects in articulating a supposed identity relation. Even if one world stipulates that these objects are determinately identical (as at this world there may be no accessible world where the objects are distinct), and another world stipulates that these objects are determinately distinct, this itself not give rise to indeterminate distinctness. Rather, the presentation offered by Akiba shows that a world claiming MK is identical to MK+ is in fact offering a *description* of identity, rather than a *genuine relation* of identity. More precisely, if w' holds that MK is identical to MK+, then w1 would only have one object in the domain: MK. Both objects would not exist separately in the world if they were truly determinately identical.

So, in the world where MK ≠ MK+, there are in fact two objects, not one. This holds even if there is an identity-like relation that can hold between them – perhaps a relation such as 'x is the mountain constituted by y' on some precisification. This is not identity, and indeed it could not be, since MK has a property which MK+ lacks, namely, the property 'is always precisified by MK'. Considered more broadly, determinate identity *at a world* is not determinate identity *simpliciter*; rather, it is a sort of "relativised" determinate identity as pertains to identity *at and from a world.* Therefore, this relative determinate identity does not support Akiba's view that MK is not determinately distinct from MK+: MK's inclusion as a separate object in the domain with MK+ illustrates their

[35] Interestingly, Akiba does address Smith's paper (2014, pp. 3569-3571). However, Akiba focuses on countering Smith's criticism of views that attempt to 'make sense' of vague identity (or of indeterminate individuation, per Akiba's terminology). He does not quite consider the criticism we present here, though he comes close in discussing his 'analogical argument' (2000, 2004) from his earlier work (2014, pp. 3567-3568).

determinate distinctness. Even granting that MK may be determinately identical to MK+ relative to the worlds accessible from w1, this itself does not supplant the characteristics distinguishing MK from MK+. Ultimately MK and MK+ two objects are not indeterminately distinct; rather, they are determinately distinct, even if at some worlds there is the apparent guise of determinate identity.

The second problem relates to the broader dialectic of vindicating vague objects. While Akiba does not explicitly engage with this dialectic, it is nonetheless important in the context of the logical conservativism outlined and defended in **1.3**. Akiba's model rests on two key assumptions: first, the model offered utilises a precisificational framework (so as to involve the 'process of precisification' that motivates the framework underpinning the model); second, this precisificational model utilises an anti-symmetric accessibility relation. Both assumptions are independently plausible and may be fruitful for resisting Evans's argument (as Akiba intended to do). However, both assumptions are most certainly not necessary, nor desirable, for the proponent of vague objects to accept. As we will discuss in Chapter 3, other accounts have been constructed in defence of vague objects. Some use precisificational models that rely on non-anti-symmetric accessibility relations, while other accounts forgo the machinery of precisifications and instead utilise other conceptual tools. So, on both counts, even if successful, Akiba's model remains one among various options. Its force as a counter to Evans' argument needs to be contextualised within the broader field in which the account is being offered. Furthermore, given the highly specific logical commitments of the model, it remains unnecessary for the vague-objects theorist to accept to nonetheless attempt a defence of vague objects. Starting from a particular set of logical assumptions for the purpose of responding to Evans' *reductio* seems unnecessary and goes against our motivation to construct a logically conservative (classical) account of vague objects, as discussed in **1.3**.

To recap: what has this section established? For the purpose of vindicating vague objects, Akiba's defence of the possibility of indeterminate distinctness, at best, offers the following picture: if vague objects are objects that are indeterminately distinct from other objects, and if the notion of indeterminate distinctness helps to elucidate the nature of vague objects, then we may have a means of vindicating vague objects from Evans' line of attack. However, the above argument does not necessarily tie the fate of vague objects to the possible cogency of indeterminate distinctness. Indeed, it seems that we have good reason *not* to do so. Unless one was ostensibly committed to Akiba's precisificational framework, then there is little immediate utility to the idea.

On the other hand, if our goal were to rescue vague objects from Evans' line of attack, then Akiba's rejection of Premise B could have utility (assuming its plausibility). However, we wish to argue for a stronger thesis. We intend to defend the actual existence of vague objects, rather than defend the mere possibility of their existence through a rescue from Evans' *reductio*.

Hence, this section highlights the proper target of our enquiry: Premise A of Evans' argument. While the logical consistency of vague identity is important, for our purposes it is more pressing to determine whether vague identity is a necessary part of a defence of vague objects. As we have seen, there are many lines of attack against Premise B; however, far fewer responses have targeted Premise A. The arguments we have unpacked have demonstrated the various logical principles that a proponent of VO ought to uphold. These principles serve to delineate fruitful accounts of vague objects - fruitful approaches to attacking Premise A. Our attention will now turn to applying these principles towards constructing a metaphysical framework, and towards responding to an infamous paradox: the Problem of the Many.

1.5 Taking Stock

This chapter has responded to Evans' *reductio*, articulating and defending various logical assumptions which underpin the broader project of the thesis. We first discussed the major assumptions underpinning Evans' argument in **1.1**, including the (in)determinacy operators utilised in its construction. We then discussed Lewis' interpretation of the proof in **1.2**. We demonstrated how the semanticist can respond to Evans, and determined that while the vague-objects theorist cannot utilise the same argumentative moves, they nonetheless have plausible options available in response. We also defended the idea that logical revisionism is unnecessary for the purposes of vindicating vague objects in **1.3**. We demonstrated that other views (such as supervaluationism) often face problems when making such revisions, even to meta-logical assumptions. In short, we have defended logical conservatism as a methodological framework for the remainder of the thesis.

We then extended our analysis by considering the possibility of higher-order vagueness, finding it to be an unnecessary step on the path to vindication. As a consequence, if we are to have vague objects, then we are to begin by building from a single level of vagueness. We then considered a potential rejoinder to the Evans proof in **1.4**: Akiba's argument defending the possible cogency of indeterminate distinctness. While a novel and admirable argument, we found little reason to accept

Akiba's argument given the aims of this thesis. Chapter 1 has thus motivated the idea that if a cogent metaphysical picture of vague objects can be provided, then one can decouple vague objects from problematic instances of vague identity. Given that much of the extant literature has tied the fate of vague objects to the cogency of vague identity, their separation merits an in-depth discussion with a sufficiently robust metaphysics in hand. We set to the task of doing just this in Chapter 2, where we respond to an infamous challenge to seemingly well-founded mereological principles: the Problem of the Many.

Chapter 2 – Individuating Vague Objects: The Problem of the Many

2.0 Preamble

This chapter applies the logical principles and lessons learned from Chapter 1 to respond to Peter Unger's influential Problem of the Many (1980), and to Brian Weatherson's argument against the existence of vague objects (2003). I articulate the Problem of the Many, evaluate the virtues and drawbacks of existing solutions to the problem, and develop a novel "vague-objects solution". I then set-out and respond to Weatherson's objection, considering the metaphysical costs of endorsing the existence of vague objects. Finally, by elucidating the principles that underpin our novel vague-objects solution, I aim to vindicate the use of vague objects as a means of providing a useful response to the Problem of the Many.

2.1 The Problem of the Many

… it seems clear that no matter which relevant concrete complex is deemed fit for cloudhood, that is, is deemed a cloud, there will be very many others each of which has, in any relevant respect, a claim that is just as good … either all of them make it or else nothing does – Peter Unger, 1980, p. 415.

The Problem of the Many (PM) arises whenever there is an object with fuzzy or indeterminate spatial boundaries. For instance, following the example in McGee and McLaughlin's (2000) version of PM,[36] let us consider the part of the material world comprising Mt Kilimanjaro (MK): a collection of rocks and various other bits of matter. A close inspection of MK's boundary reveals various bits of matter that seem less clearly part of the mountain than others (a loose rock or particle at the mountain's boundary, perhaps). Some bits of matter are clearly part of the mountain, while others appear to be questionable or borderline parts. Let us assume there are a large number of these borderline parts. Given this, there may be a corresponding number of entities that have all the clear parts of Mt Kilimanjaro, as well as one of these questionable

[36] Other presentations include Unger's original presentation (involving clouds) and Geach's closely related 'Paradox of the 1001 Cats' (1980). There is also Lewis' (1999) presentation of PM as one involving vague or indeterminate parts – we return to this presentation in **2.2**. See Jones (2011) for detailed discussion of whether Unger and Lewis' presentations should be treated as two distinct problems, or as variants of the same puzzle.

parts respectively. Call these entities MK1, MK2, MK3, …. MKn, etc. Each of these entities are closely overlapping fusions of matter, for they differ only marginally in what is included as part of their boundary. While they closely overlap, each entity is distinct from each other. Further, it seems that MK1, MK2 etc. are mountains iff Mt Kilimanjaro is; after all, it would seem astonishing that the mere inclusion or exclusion of one rock could be the difference between being a mountain and non-mountain. It seems clearer still that Mt Kilimanjaro is a paradigm case of a mountain (if anything is). So, taking stock of this brief picture, where we expect one mountain, we now have a plethora of distinct but closely overlapping mountains. This result seems incredible, if not counter-intuitive; yet, a close look at the principles underpinning this conclusion seem independently plausible. From plausible assumptions we have come to an incredible conclusion. We are left with a problem that merits further investigation.

Part of PM's set-up is the idea that every "candidate" mereological fusion – every mereological sum that appears a suitable fusion of the object and a satisfier of the sortal predicate or object kind in question (MK1, MK2, and so on) – is itself an object of that kind (Weatherson 2014, S7). For instance, in our mountain case we take it that each of the closely overlapping mountains is a satisfier of the property 'is a mountain' (or of the sortal predicate 'mountain'); this is, after all, why the presence of such fusions gives rise to a problem in the first place. Importantly, this feature of PM distinguishes it from other cases involving vagueness such as the Sorites paradox or the Ship of Theseus. In these latter cases, there is typically a point where it is vague whether a property or sortal predicate, such as 'heap' or 'Ship of Theseus', obtains. These examples of vagueness call for a means of appropriately classifying objects in line with sortal predicates, what Geert Keil (2013) calls 'vagueness of classification' in which vagueness arises from an absence of precise satisfaction criteria for determinate sortal predicates. This is quite unlike the problem posed by PM, where each candidate fusion of parts is itself an object of the relevant kind. So, if one wished to maintain the intuition that there is one mountain despite the presence of many candidate mereological fusions, then this would involve determining which of many candidate boundaries of an object is *the* single correct one, rather than an attempt to do the work of explaining whether, and why, the candidates are in fact mountains. This is what Keil calls 'vagueness of individuation', and it is this that will be our focus for the chapter.

To help distinguish different solutions to PM, and to better consider the assumptions and intuitions that feed into the paradox (any good paradox, of course, results in a clash between independently plausible assumptions), let us borrow Woods'

(forthcoming) tripartite set-up to the problem. For statements involving objects of a kind K (mountains, cats, tables, humans, etc.), we can consider PM as involving the following three principles (Woods forthcoming, p. 3):

Conservatism: There is just one K where we take there to be a single K.
Abundance: There are many *K-candidates*, i.e. things the part of which seem to be arranged K-wise, in the vicinity of an object of kind K (such as MK1 and MK2 for Kilimanjaro).
Parity: If any one of [these candidates] is a K, then each [candidate] is a distinct K.

As one might expect, Conservatism, Parity, and Abundance form an inconsistent triad. Conservativism holds there is one K. Abundance holds there are many K-candidates. Parity holds that if a K-candidate is a K (which is granted), then there are many K's, *contra* Conservativism. Seemingly innocuous assumptions have instead led to a paradox. To see the paradox more clearly, we would be well served to unpack the assumptions that lead to it.

Conservatism captures the idea that when there is an object satisfying a sortal predicate (a mountain, cat, person, cloud, etc.), there is just one unique object and one unique mereological sum in that vicinity. This principle represents our 'counting intuition' (Lopez de Sa 2014); the idea that we typically do not include a plethora of closely overlapping fusions when counting ordinary objects. For instance, when we count how many people typed this thesis, or when we count how many Mt Kilimanjaro's there are, in non-PM contexts we typically state that there is one, not many.

Abundance states that for any object of a kind K, there appear to be a large number of only minutely different objects in the vicinity that similarly satisfy K-ness. Parity captures the idea that if Abundance holds, and if any K-candidate satisfies the criteria for being a K, then all K-candidates are of that kind (while distinct in having slightly different parts, they are still themselves K's). Here one may baulk at the suggestion that the difference of one small rock (or a water droplet for a cloud, or a hair for a cat) would fail to meet similar satisfaction criteria, absent some clear principle. In other words, rejecting Parity and/or Abundance would require a principled reason for excluding the many candidates from the purportedly one "genuine" object.

Yet, we seem to do just this in a variety of contexts. Consider the idea of a 'maximality principle' (see Sider 2001, 2003a), where sortal predicates are taken to be *maximal* in the sense that candidates satisfy these predicates if they are not a large proper part of

a satisfier of the object of a same kind. So, if there are two satisfiers of the predicate ‘mountain’, then the satisfier that is the large proper part of the other would be “excluded” in this sense. Perhaps the many candidates should be seen as a plethora of large proper parts of K, rather than as a series of distinct K’s. Using maximality, we could respect our counting intuition and seemingly do away with the many.

However, the sense in which the K-candidates are objects of the same kind (perhaps due to shared possession of some relevant intrinsic properties) may give us pause. Though maximality enables us to select one from the many, it excludes many entities which themselves are tremendously similar to the “maximal” entity. Does lacking large proper parts of the same object kind satisfactorily explain away the apparent existence of the many candidates?

Perhaps not. Consider Lewis’ (1993) example of Tibbles the cat. While there may be many fusions of cat-cells in the vicinity of Tibbles, all of these cat-candidates would seem to purr, meow, scratch, etc. Even with the maximality principle in hand, there is a strong sense in which the many have very good claims to being existent things. This relates to the idea of a ‘grounding intuition’ (Lopez de Sa 2014). What constitutes one candidate would provide similar grounds for other candidates to satisfy the relevant object kind. Even granting something like the maximality principle, the elimination of candidates requires explaining away their apparent similarity, and to explain how the apparent special properties that the one possesses *also* gives us ground to disregard the existence of the many candidates in its vicinity. In short: we must ensure selecting one object as *the* unique satisfier of some object kind is not an arbitrary differentiation. Further, we ought to ensure that such a selection considers the intuitions which lead us to contemplate the possibility of there being K-candidates in the first place.

If we wish to respond to PM, and so rescue ourselves from the paradox, then an analysis of existing solutions is in order. Through a discussion of these solutions, we will see the force of the problem raised by PM and elucidate the independent plausibility of each background assumption. With this in mind, the purpose of the next section is not to present an exhaustive account of existing solutions,[37] nor is the purpose to present knockdown objections to these solutions. Our purpose instead is to show that solutions to PM each have their virtues, and their costs. In attempting to retain intuitions and seemingly plausible assumptions, there are associated costs in

[37] Comprehensive surveys of various solutions can be found in Weatherson (2014) and Loux and Crisp (2017, pp. 250-278).

the need to reconsider other seemingly well-founded ideas. By presenting the views as such, we will motivate further inquiry into the possible merits of a vague-objects solution to PM in **2.3**. As we wish to consider the merits of a vague-objects response to PM, we must first do due diligence and show that existing solutions do not come cost-free.

2.2 Responses to the Problem of the Many

To deny that there are many cats on the mat, we must either deny that the many are cats, or else deny that the cats are many. We may solve the paradox by finding a way to disqualify cats for cathood … Or else, if we grant that all candidates are truly cats, we must find a way to say that these cats are not truly different from one another – David Lewis, 1993, p. 167.

As with any good paradox, responses to PM vary. We will briefly consider a number of these responses before moving to a detailed discussion of Lewis' response, as well as two contemporary solutions to PM. After engaging with these responses, we will then elucidate the role that vague objects may play in resolving PM and evaluate various criticisms of such a view.

2.2.1 – Bite the Bullet: There are Many or None

Unger (1980) provides two responses to PM. The first holds that there is an abundance of ordinary objects where we normally count one. Without a principle motivating the elimination of all but one K-candidate, each candidate mereological sum seems to satisfy the conditions for being the relevant object kind. So, if we lack a clear means of "excluding" these candidates, then perhaps the lesson to draw from PM is that there are, in fact, many mountains, cats, and persons. To many, this result would appear astonishing. A plenitude of (sometimes arbitrary) mereological sums – K-candidates – have now quickly become a dizzying array of mountains, cats, and persons. An answer that can explain the grounds for such a surprising conclusion, or provide reason to avoid it, is clearly desirable.

Per Unger's second response, an eliminativist (or nihilist) view, there are no mountains, cats, or composite objects. On such a view, Abundance does not hold, for there are no

mereological sums.[38] If there is no clear means of distinguishing one object as uniquely having a property of mountain-hood (or cat-ness, or personhood, etc.) which the other candidates lack, then perhaps none of the candidates satisfy the sortal predicate over and above the others. Perhaps the lesson to draw from PM is that there are, in fact, no composite objects; instead, what exists are bits of matter arranged mountain-wise and table-wise. On a first reading this view also seems astonishing; if there are no mereological sums, then Conservatism also seems under duress. Yet, detailed and plausible accounts of eliminativist views have been offered,[39] though these views may not be able to fully escape PM's grasp (one could possibly recast PM in terms of bits of matter "arranged mountain-wise" – perhaps a parallel plenitude can be constructed in a manner compatible with eliminativist views).[40] Alternatively, we may wish to retain the view that there are such things as composite objects; this too, may be desirable.

Both of the above views reject at least one of the principles that comprise the inconsistent triad. This is to be expected when responding to a paradox. However, in doing so, both produce astonishing results. It would be good to avoid such conclusions in resisting the challenge posed by PM. Yet, as we will see, attempts to maintain intuitions that seem tenable outside of the context of PM lead to puzzling conclusions of their own.

2.2.2 – Supervaluationism: Determinate That, Indeterminate Which

Another view, famously put forward by David Lewis (1993), involves a rejection of Conservatism and the idea that there is just one cat on the mat. Lewis' solution combines two views: a supervaluational response to PM, and an 'almost-identity' response. We will evaluate each response in turn before considering how Lewis combined them to form his "two-part" solution to PM.

Recall our earlier description of supervaluationism in Chapter 1, in which there are multiple interpretations (or precisifications) of the target language. Couched in an understanding of vagueness as semantic indecision, Lewis' supervaluational response

[38] Or, on a more conservative reading, there are no mereological sums of things that do not constitute a "life" (consider van Inwagen 1990).

[39] For more detailed accounts of such views, with differing criteria as to when or how composition occurs (if ever), see Unger (1980), van Inwagen (1990), and Merricks (2001).

[40] See Merricks (2017) for an interesting argument defending the possibility of 'worldly vagueness', or vagueness in the absence of language and thought. In his paper, Merricks briefly states that his arguments can plausibly be 'recast in eliminativist terms'. Such arguments could, perhaps, be applied to cases involving PM (though Merricks' paper itself does not address this idea).

to PM holds that there are different interpretations of the exact extension of the word “mountain”, or “cat”, or any sortal predicate. On different interpretations there will be a different candidate “selected” as the precise extension. Extended to PM, we have a case where it is super-true that there is one mountain-candidate that is “the” correct Mt Kilimanjaro (every precisification has one candidate where this is the case – MK1 on one precisification, MK2 on another, etc.), but where it is left undecided “which” candidate is in fact the mountain (the precisifications disagree with respect to which candidate comes out as true on their reading). Indeed, on this view there is no candidate (nor fusion) such that *it* is determinately the mountain, the cat, etc. While it is determinate that there is a cat, mountain etc., it is indeterminate which of the many candidates is the satisfier of the property. So, while supervaluationism provides a cogent response to PM (we have a means of retaining our counting intuition that there is just one object), it does not provide us with a means of non-arbitrarily selecting any particular candidate as “the one”. In short, there is a gap in determining what principled difference there is between candidates in different sharpenings, other than the idea that these sharpenings are admissible but distinct ways of precisifying the extension of a term. Supervaluationism accounts for the grounding intuition, but in doing so it excludes a number of otherwise admissible candidates on different precisifications without providing an explanatory mechanism justifying this exclusion.

Lewis candidly accepts this, stating that ‘we cannot deny the arbitrariness. What we can deny, though, is that it is trouble’ (1993, p. 172). So, while it may be peculiar that the supervaluationist picture leaves it indeterminate *which* candidate is *the* correct object, one may argue that this is not a drawback of the view. Perhaps PM gives rise to an inherent explanatory gap in selecting the correct candidate; following this, perhaps a strict sense of Conservatism should be loosened to capture our semantic indecision regarding the exact spatio-temporal boundaries of objects. Here though, we may wish to interrogate the means by which supervaluationism leads to the selection of different admissible candidates on different precisifications. In other words, how does supervaluationism account for the grounding intuition’s rejection? This is the very criticism that Lopez de Sa (2014) raises with the view. Specifically, supervaluationism seems to run against the ‘principle of minute differences’ (pmd):

Pmd: if something is a *paradigm* case of an f, and something else is very similar to the former with respect to the features relevant for something being an f, then the latter is also an f.[41]

Where supervaluationism runs into trouble is in violating pmd. Each candidate in PM seems to be "very similar" to the other in terms of the relevant features they possess; yet, on different precisifications these candidates are determinately excluded in favour of one. While this gives us a cogent response to PM, it does not provide us with a non-arbitrary means of stating what enables this selection in the first place. If a candidate is admissible on one precisification, and not admissible on many others, then it seems to be desirable to have an explanation for why we can do away with pmd or other related principles. If there is a relevant difference between different candidates (other than the precisifications in which they feature), then in virtue of what is this the case? Though arbitrary selection itself may not be troublesome for supervaluationists like Lewis, the absence of a non-arbitrary articulation of the relevant differences between different candidates (differences which enable selection and exclusion in different precisifications) seems to present a greater concern.

To be clear: the main point we are emphasising here is not that a rejection of the grounding intuition or pmd is faulty; rather, as Lopez de Sa highlights, we are ideally looking for accounts that can offer a principled, non-arbitrary explanation for why the grounding intuition should or should not be upheld. Indeed, as we will see, the drawbacks of supervaluationism led Lewis to supplement the view with another means of responding to PM.[42]

2.2.3 – Identity: From Many to Almost One

The second part of Lewis' two-part solution involves rejection of Conservatism and the idea that there is just one cat on the mat. On Lewis' view, our counting intuition holds force if we count by *strict* identity; yet, the intuition may be weakened if we consider

[41] Lopez de Sa (fn. 6) distinguishes pmd from McKinnon's (2002) closely related ideas regarding 'Non-Arbitrary Differences' (nad). The key difference between the two is that while nad similarly speaks of 'principled differences' between closely related entities (candidate-objects and objects seem to not differ in a principled manner), pmd only includes 'paradigm' cases in its antecedent. Lopez de Sa does so to protect pmd from sorites-style objections. As nad does not restrict itself to paradigm cases, it leaves itself open to a sorites-style case of a series of objects that do not differ from their "neighbours" in a principled manner.

[42] Or, as Lopez de Sa argues (2014), these drawbacks should give us reason to reject the supervaluational response in favour of accepting the almost-identity solution as a standalone solution, given that it circumvents the issues relating to the supervaluational view.

counting through use of the weaker notion of *partial* identity. In brief, on Lewis' view none of the assumed "cat-candidates" are strictly identical to the cat. Rather, each is *almost* identical to the cat. For Lewis, using the notion of *partial* identity allows us to respond to PM by saying that the many are all *almost* one: the many candidates are all *almost* the cat. Where strict identity and strict distinctness sit at opposing ends of a spectrum, partial identity captures a middle ground between the two. Consider the mereological analogues of coincidence (or complete overlap), partial coincidence or (partial overlap), and disjointedness (or complete distinctness).[43] Here partial identity is the analogue of partial overlap. Extended to PM, where the many candidates significantly overlap one another, it may seem appropriate to say that all candidates are *almost identical* given that they all *almost overlap* given the many parts that they share. Here the candidates are neither completely distinct nor completely identical. In Lewis' words:

> strictly speaking, the cats are many. No two of them are completely identical; their differences are negligible … we have many cats, each one almost identical to all the rest (1993, p. 178).

Lewis relates his use of partial identity to conversational contexts in which strict ('absolute') identity is not always needed (such as in regular conventions of counting). For Lewis a 'blameless approximation' of 'almost identity' will do, as 'for most contexts it's true enough' (1993, p. 178). In cases where strict identity is needed, then for Lewis the relevant concept can be utilised in these contexts. So, while strict identity has value and utility in many circumstances, in many ordinary conversational contexts we do not use strict identity. In special cases such as PM, partial identity also enables us to maintain Parity and Abundance (the grounds for there being K-candidates in the first place) while replacing our counting intuition captured by Conservatism with a view of 'almost identity'. While intuitions hold there is one cat, for the almost-identity solution there are instead many almost identical cats.

As Lopez de Sa (2014) notes however, many find the almost-identity solution to be 'implausibly counterintuitive'. For one, rejecting Conservatism is a costly manoeuvre, particularly given the strength of the counting intuition which it represents. This is not to say that doing away with the intuition is always going to be a doubt-riddled affair; rather, it raises the need to explain how our counting practices are to work if we reject

[43] Lewis also briefly presents a similarly analogous discussion (1993, p. 177).

the idea of there being one object where we expect one object.

How, then, does Lewis' view justifiably dispose of Conservatism? It does not; rather, what the view entails instead is a weaker sense of Conservatism. We may determinately have a mountain or cat without it being determinate "which" entity is "the" object (per supervaluationism), or we may be able to postulate that the entity is *almost identical* to various candidates (the almost-identity solution). In neither case do we capture Conservatism "proper", in the sense in which there is one truly distinct and unique object that is non-arbitrarily the proper target of our counting intuitions. Conservatism is modified – our counting intuitions are accounted for by considering how we may use identity in strict and conversational senses. In contexts involving PM, perhaps our counting intuitions are better mapped to relations of partial rather than strict identity.

For some this may be too costly a modification. Even so, this is not necessarily a knockdown argument against Lewis' view (and other corollaries); rather, this should provide motivation for evaluating views that attempt to capture a stricter sense of Conservatism and of our counting intuitions. The counter-intuitiveness of modifying or rejecting Conservatism should perhaps not be surprising when faced with a paradox; the more pressing question is whether such revision is implausible. Paradoxes rarely leave us comfortable. Indeed, as we will soon see, while other views succeed in maintaining Conservatism and our counting intuition, doing so will not come cost-free.

2.2.4 – Constitution: From Many to One

A different line of response to PM is to reconsider Parity: to reject the idea that the many candidates are, themselves mountains (or cats, persons, etc.). On the view espoused by Lowe (1982, 1995) and Johnston (1992), a solution to PM requires invocation of the difference between what constitutes an object, and the object itself. While the many are mountain-constituters and cat-constituters (different collections of physical matter that are, themselves, sufficient to *constitute* a mountain, cat, etc.), they are not themselves the actual object itself. In a slogan: constitution is not identity. What such accounts need to address is to explain the relation of constitution, and to demonstrate how an object can be distinct from its spatial (and possibly temporal) constituents. In short, constitution accounts need to explain the apparent duality between the constitutive matter of an object and its identity conditions.

Lowe (1995) argues that only one of the many "candidates" is the actual constituter:

the others *would be* the constituter if particular parts were altered. If we include borderline parts as per Lewis, then Lowe argues that an object can have 'just *one* constituter, but that it is *indeterminate* whether this is [K] or a certain [K-candidate]' (1995, p. 180). The basis of this indeterminacy, for Lowe, is vagueness in designation: 'the definite description 'the constituter of Tibbles' is a *vague designator*' (1995, p. 180, emphasis in original). Endorsing a loosely supervaluationist view, Lowe grants that there may be semantic indecision or unsettledness with respect to "constituter". Importantly however, Lowe does not extend this supervaluational gloss to *cats*, only to their *constituters*. Hence, Lowe intends to show that there is only one cat, but that there is a plenitude of candidates for what constitutes it (many constituters). One of the constituters is the cat, but it is indeterminate which. This indeterminacy is due to semantic indecision, grounded in a broadly supervaluational account of how to account for such semantic indecision (where the extension of "constitutes" may vary across sharpenings, and where a different constituter counts as constituting Tibbles on each sharpening).

An objection to raise here is that it is unclear what work constitution is doing to help resolve the problem. If constitution was introduced to help explain the apparent uniqueness of the one above the many, then what is it that *metaphysically* distinguishes the one constituted cat from the many candidate constituters? If one cat is 'significantly more cat-like than any other' on each sharpening (hence how Tibbles can determinately refer to this cat on each sharpening), then in virtue of what is this the case? Given that we are looking for non-arbitrary reasons for selecting one cat among many, this seems unsatisfying. The explanatory gap is also present in supervaluationism (though it is accepted with qualification); here the gap remains, but we are left with an additional piece of ontological machinery that does not fulfil its purported explanatory purpose.[44]

Even worse, another objection can be raised: speaking of a duality of constituters and objects merely modifies the problem. Where previously there were 1001 cats, now there are 1001 cat constituters. Lowe responds that the same intuitions which ground our counting intuitions about cats may not apply so easily to constituters – intuitions can shift depending on the context in which they are applied. Even granting this, it is

[44] This criticism can also be put to Markosian's (1998) "brutalist" solution to PM, in which there are "brute facts" for when things compose. Markosian's account provides a rather clear response to PM: there is one mountain, and it is a brutal fact that there is only one, not many. This avoids the issue of relying on exotic ontological machinery; however, the account bites the bullet in being opaque about explaining how or why such composition occurs, and why this occurs for one candidate over and above others.

still unclear how thinking of PM in terms of a vague constitution relation assists in selecting the correct cat. Perhaps there is a salient difference between Tibbles and her constituters, and perhaps it is indeterminate which candidate constitutes Tibbles on differing precisifications. Even granting all of this still, a constitution-based account does not itself offer a selection principle – we have no special means of picking one over many. This is not to say that constitution cannot play such an explanatory role; rather, the above highlights that such work needs to be for the ontological machinery to be of use in responding to PM. As we will see in **2.25**, recent work has been done to do just this.

2.2.5 – Multiple Constitution: The Many Constitute the One

A more contemporary answer using the machinery of constitution comes from Jones (2015). Broadly following Lowe and Johnston's line of argument, Jones responds to PM by presenting an account of 'multiple constitution', and more importantly, presents an explanation of how the concept explains away the apparent many. Jones' response also develops a broader ontology of ordinary objects with which to justify the use of multiple constitution – his account explicitly sets out to present a cogent metaphysics to accompany his response to PM.[45]

Interestingly, Jones' solution to PM seeks to retain all three of the principles underpinning PM (albeit a modified form of Parity).[46] He wishes to retain Conservatism (what he calls Solitude), Abundance, and a weaker form of Parity, what he calls Equality. Jones presents Equality as follows (2015, p. 221):

Equality: If many collections of particles are equally good, and good enough, candidates to constitute cats, then each of those collections constitute a cat.[47]

[45] This metaphysics sees objects as 'fundamentally subjects of change: the changes an object can survive are explanatorily prior to its constitution' (2015, p. 218). For Jones there are 'characteristic changes' which objects survive – cats survive losing hair and nuzzling their owners' hands, for instance. For Jones these characteristic changes are too 'coarse-grained' to enable us to distinguish between different candidates in PM-contexts, given that the characteristic changes (not changes *simpliciter*) would apply to each candidate. Jones develops this idea in full in his paper; it is beyond the scope of this paper to evaluate the tenability of such a view.

[46] For the sake of clarity, we restate the principles here:

Conservatism: There is just one K where we take there to be a single K.
Abundance: There are many *K-candidates*, i.e. things the part of which seem to be arranged K-wise, in the vicinity of an object of kind K.
Parity: If any one of [these candidates] is a K, then each [candidate] is a distinct K.

[47] The relevant difference between Parity and Equality is that while Parity relates to each candidate comprising a *distinct* cat (or object), Equality makes the weaker claim that each candidate (collection of

Given that Jones accepts Conservatism, Abundance, and Equality, how does he aim to resolve the problem arising from their supposed joint inconsistency? On his view, there is a fourth assumption underpinning the set-up to PM which enables the paradox to go through – Unique Composition:

Unique Composition (UC): 'Tibbles, like every other cat, is constituted by exactly one collection of particles.'[48]

Jones' response to PM rejects UC; the one is not composed by a single collection of particles; rather, it is multiply constituted by the various candidate constituters of the object. Hence, Tibbles the cat is constituted by many collections of particles, not one. Jones (fn. 6, fn. 12) quickly discharges the potential concern that any view of constitution can be construed as involving "multiple constitution" in the sense that objects can be readily seen to be constituted by a collection of particles and other physical matter. Where the "multiple constitution" view differs is that it takes entities to be constituted by 'many collections of particles that do the constituting' rather than merely 'many particles'. The commitment to objects being constituted by multiple collections of many particles, rather than one collection of many particles, distinguishes 'multiple constitution' from the closely related view of 'unique composition' that Jones explicitly rejects. Hence, Multiple Constitution can be described as follows:

Multiple Constitution (MC): Tibbles is constituted by many collections of particles.

How does Jones justify consideration of MC? His insight comes from the joint consideration of Conservatism, Equality, and Abundance. Taken at face value, Equality and Abundance imply that many collections of particles (many candidates) all constitute a cat, while Conservatism implies that there is one cat. Jones' response holds that the many collections implied by Equality and Abundance are precisely what constitute the one. While such a reply goes against the widely held notion of UC (which may have tacitly undergirded prior responses to PM), it is quite amenable if one accepts MC.

The result may seem counter-intuitive at first; there is one cat, but it is constituted by

particles) constitutes *a* cat, without this necessarily entailing that they constitute a *distinct* cat. We unpack the significance of this distinction below.

[48] Jones notes the worry that if Tibbles is composed of a set of Fs, and if the Fs are themselves composed of a set of Gs, then Tibbles is constituting by Gs and Fs, giving rise to banal kind of multiple constitution. With this in mind, Jones restricts UC to 'absolutely fundamental particles, or to 'elements of relatively fundamental decomposition of reality into non-overlapping particles' (2015, p. 221).

each of the 1001 cat-candidates (or cat-constituters).[49] If one does not already buy into the concept of multiple constitution, then Jones' response may seem perplexing. Even granting this, a close inspection reveals the many virtues of Jones' view. Importantly, it respects our counting intuition, for there is one (multiply constituted) cat. The view also respects our grounding intuition, as the collections of particles constituting each candidate are also part of the multiply constituted singular object. Further, a gerrymandered selection principle is not needed, as the legitimate grounds of the many are explained as their being part of what constitutes the one. The use of many-one constitution is also doing clear explanatory work in underpinning the solution to PM, unlike Lowe's use of constitution.

Here one could comment that while Jones' solution abides by Equality (for if the candidates multiply constitute one cat, then they each constitute at least *a* cat), this goes against Parity (which requires each candidate to compose, or constitute, a *distinct* cat). Point taken, though here the criticism levelled against modifying one of the assumptions that give rise to the very paradox may be too quick. We could envisage Jones replying that the solution he offers provides justifiable grounds for disposing of Parity, given his use of Equality, and given his use of multiple constitution. Surely, then, we cannot dig our heels in and enforce the acceptance of Parity for all proposed solutions of PM, particularly if there is good reason to do away with the assumption.

Of course, however, all of these virtues rely on the acceptance of two weighty ideas. First, we need to accept MC as a cogent piece of onto-ideological machinery – far from an inexpensive cost. Second, we also need to accept Jones' particular ontology of ordinary objects; a picture which he uses both to motivate the idea of multiple constitution and also as a means of couching it within a consistent metaphysical framework. Many could resist Jones' view at this stage, as they may not wish to accept the cost of buying into the metaphysics entailed by accepting these ideas.

Yet, the many virtues of Jones' view are apparent, and much of his view's appeal comes from offering a cogent response to PM in a way that respects many of the

[49] This may give rise to another worry: if the 1000 cat-constituters suffice to collectively constitute a cat, then there would be good reason to think that 999 cat-constituters would similarly suffice to constitute a cat (had only 999 existed, they surely would still constitute a cat!) Generalising this idea, perhaps there are many collectively-constituted cats where we had one – maximality may be needed for the multiple constitution view. Note however, that this would require a maximality principle construed for constitution (for example, the Fs only *constitute* a G if there is not larger collection of Fs which constitute a G). Yet, such a principle seems less obvious than the original presentation of the maximality principle (involving large proper parts).

intuitions we held prior to grappling with the paradox. In saying this, there is room to determine whether a less ontologically costly view is on the offing. If we can respond to the problem, respect our counting intuition, and offer a means of sustaining our grounding intuition (and so respect the many candidates' relation to the one) without necessarily relying on the notion of multiple constitution (and the ontological view that enables it), then this could produce an even more fruitful account. Given this, then, could we take motivation from elsewhere, such as from Lewis' views regarding identity?

2.2.6 – Many-One Identity: The Many are Identical to the One

A recent response building on Lewis' almost-identity solution comes from Woods (forthcoming). Woods offers a 'many-one identity' solution to PM. Instead of using partial identity, per Lewis, Woods instead argues that many things can *collectively be identical* to one thing – an identity parallel to Jones' multiple-constitution view. For instance, on Woods' view a cat is collectively identical to a set of legs, hair, whiskers, and so on. Woods grounds this idea in "portion of reality" talk, whereby multiple entities may collectively comprise a single portion of reality (such as my legs, torso, head, etc.) and so be identical to it. So too, then, could one say that 'the many candidates are the same portion of reality as the cat and are thus identical to the one cat on the mat and compose it' (Woods forthcoming, p. 6).

Applying this view to PM, Woods offers a "many-one solution" as follows:

Many-One Solution: 'in instances of the problem of the many, the many candidates for being an object of kind *K* are collectively identical to a single *K*, and no candidate is identical to any other object of kind *K*' (Woods forthcoming, p. 7).

Leaving issues regarding the tenability of many-one identity to the side,[50] what we are interested in here is that assuming many-one identity to be a cogent notion, does it present a useful and intelligible solution to the problem? Recall the inconsistent triad of Parity, Abundance, and Conservatism. Per Woods' view, Conservatism is retained as there is one object: that which is identical to the collective of candidates.[51] Similarly.

[50] Woods notes this, and 'plead[s] for toleration' in using the notion of many-one identity (forthcoming, p. 4). He does not deny that there are criticisms regarding many-one identity (though he thinks responses can be levelled at these criticisms); rather, his paper aims to determine that '*if* there are many-one identities, then many-one identity can offer a solution to the problem of the many' (forthcoming, p. 5). Accordingly, we will leave these objections and responses to the side.

[51] As Woods notes (forthcoming, p. 9), there would only be one such object, for there is only one that is collectively identical to *all*, and not merely *some*, of the other candidates. This distinguishes the one as

Abundance is also retained, given the respect of the existence of the many candidates. Absent a (non-arbitrary) selection or exclusion principle, there is still the need to respect the candidates as being seemingly equally good satisfiers of the relevant object-kind (the very means by which PM arises in the first place). The many-one identity solution respects this by treating each as part of the identity conditions of the one object; each candidate plays a role in forming the collective identity of the object, while nonetheless being distinct mereological sums.

What Woods rejects, interestingly, is Parity. Each candidate is not a *distinct* K; rather, they are collectively identical to a single K. So, the many candidates are not each different mountains, cats, etc.; instead, the candidates are collectively identical to one mountain, cat, etc. This picture relies on the use of many-one identity, and in particular, the broad idea that multiple "portions of reality" are collectively identical to a larger portion of reality. Here Mount Kilimanjaro is collectively identical to the many candidates (MK1, MK2, etc.), for they comprise the "portions of reality" to which Mt Kilimanjaro is identical.

Woods' solution has a number of similarities to Jones'; indeed, Woods notes that 'the main disagreement between [the two] is ideological, viz. whether one accepts the ideology of constitution or many-one identity' (forthcoming, p. 16). Both views are similarly motivated and have similar virtues in producing a desired answer of sustaining our counting intuitions, and in offering an explanation with respect to our grounding intuitions. Both also rely on the acceptance of "many-one" relations, a byproduct of the explanatory utility their views offer.

In offering "many-one" solutions to PM, both Jones and Woods move the dialectic forward in sustaining counting intuitions while grounding these solutions within a clear (albeit controversial) ideological apparatus. Compare these many-one solutions to our earlier discussion regarding solutions relying on supervaluationism and orthodox constitution. While these latter views were able to capture our counting intuitions, they were unable to provide the means by which to sufficiently ground the distinction of the one from the many. By contrast, Woods and Jones' many-one solutions were able to sufficiently ground this distinction and explain the notion underpinning this ground. The former views were relatively inexpensive in their ideo-ontological cost; however, they were unable to sufficiently furnish their views to explain the apparent similarity of the

the "best" candidate in that no other candidate has the feature of both being a K and being collectively identical to (and being composed of) all other K-candidates.

grounds between the many and the one. Alternatively, Jones' and Woods' views were able to do this explanatory work, albeit at the expense of requiring the acceptance of heterodox ideo-ontological concepts. Considering this, then, there is motivation to attempt to sustain the virtues of both approaches; could we offer a relatively inexpensive but explanatorily substantial response to PM? This seems an ambitious move; however, the prospects of developing such a solution are worth investigating. We do just this in **2.3** by exploring the prospects of a vague-objects solution to PM.

2.3 Vague Objects: From Many to One?

No treatment of ontic vagueness will be satisfactory that fails to explain how ontic vagueness is related to the phenomenon of fuzzy spatial boundaries and to the problem of the many – Geert Keil, 2013, p. 162.

As with many of the solutions discussed, our vague-objects solution to PM seeks to retain Conservatism and our counting intuition. It does so by defending the idea that there is only one object – the vague one! Specifically, by taking the many candidates to be merely precise mereological sums, the view intends to offer a means of distinguishing one vague object from the many candidates. Without further substantiation however, this strategy seems faulty: the reasons why we might stipulate that there is one vague object (possessing both precise and vague parts) would provide similar reason to include various *vague* candidates. Where previously there were many precise candidates, now there is also a plethora of vague candidates. Where there were 1001 cats, now there are 1002. Even worse, where there were 1001 precise cat candidates, now there is grounds to posit that there are 1001 *vague* cat candidates, for the reasons grounding there being one vague object would, absent a further principle, provide similar grounds for there being many vague cats. Call this the problem of 1001 vague cats. Surely an explosion of candidates is not a desirable consequence of the view. Here, then, we see it is insufficient to merely posit that there is one vague object amidst a plenitude of precise candidates: further work needs to be done to retain Conservatism. Woods and Jones succeeded in offering reasons for positing the one as constituted by, or identical to, the many; could there be a similar view in the offing for a proponent of vague objects?

Some initial criticisms suggest pessimism. Lewis (1993, pp. 169-171) provides two pointed critiques of the vague-objects view: the first is that it seems to require commitment to the untenable notion of vague identity. While we may have discharged

vague objects from vague identity in Chapter 1, the objection in the context of PM may have particular force, considering that notions other than vague identity (such as vague composition or vague existence) may nonetheless lead to problematic instances of vague identity (or give rise to other problems).[52]

The second criticism levelled by Lewis pertains to parsimony: given the existence of less ontologically inflationary solutions (such as his preferred two-part solution), unless the vague-objects view offers significant benefit, the ontological costs seem quite expensive by comparison. Other accounts are relatively metaphysically austere, and if successful, present a cleaner ontological picture – a point against the proponent of vague objects. The relative weight of this cost hinges on how one constructs the vague-objects view, and how we weigh this against its utility in response to PM. Let us begin our attempt to construct such a view.

To provide an initial idea of what a vague object is, let us again consider Mt Kilimanjaro (MK) – recall our presentation of a similar example in **1.4**. Assume that Mt Kilimanjaro is comprised of a set of definite parts, with the exception of one vague part: a rock on the borderline of being part of, or not part of, Kilimanjaro boundary. Call this rock "Rocky" (r).[53] Given Rocky's borderline status, we can stipulate that it is indeterminate whether r is a part of MK.[54] At first blush, this indeterminacy would seem to run against an instance of Excluded Middle (LEM). Broadly applying LEM, it follows that either Rocky is part of Kilimanjaro, or that Rocky is not part of Kilimanjaro. Hence, if it is indeterminate whether r is a part of MK, then it seems we have two candidate fusions for MK: one that definitely includes Rocky as a part, and another that definitely does not include Rocky as a part. Call the body of land constituted by the definite parts of Kilimanjaro, with Rocky definitely included as a part, 'Kilimanjaro+' (K+). In short, K+ is the fusion of the definite parts of MK, and r. Similarly, call the body of land constituted by the definite parts of MK, with Rocky definitely not included as a part, 'Kilimanjaro-'

[52] For instance, see Sider (2003b, 2009) for a critical discussion of vague existence. See Smith (2005) for an attempt to 'make sense' of vague composition and vague existence. See also Hawley (2003) for a defence of "modest" vague existence, and Barnes (2013) for a defence of a certain type of metaphysically indeterminate existence.

[53] See Weatherson (2003) and Barnes and Williams (2009) for similar presentations as part of related discussions. Here we use the vague part "Rocky" rather than "Sparky" (as in **1.4**) to assist with distinguishing each respective discussion.

[54] Others have used phrasing such as "questionable part", "vague part", or "indeterminate part" to capture the sense of an object such as Rocky in related discussions of vague objects. See, for instance, Morreau (2002) and Tye (1990).

(K-); the fusion of the definite parts of MK that do not overlap with r. Given classical extensional mereology (and unrestricted composition), both K+ and K- exist.

Now, in putative instances of mereological indeterminacy (where there is, ostensibly, vagueness with respect to some part of some ordinary object), a common assumption is that there is a fact of the matter as to whether an object definitely has some part or not: it is simply vague or unsettled which.[55] Following this, MK would either have a boundary identical to K+ or K-. If Rocky is included as a part of K's boundary, then MK is identical to K+ (MK = K+). Similarly, if Rocky is not included as a part of MK's boundary, then MK is identical to K- (MK = K-). Yet, given the stipulation that Rocky is on the borderline of K's boundary, there appear to be equally good reasons for including or excluding Rocky as a genuine part of Kilimanjaro. After all, if there were a decisive argument in favour of Rocky being included or not included as part of Kilimanjaro, then PM would not be such a problem! In short, there appear to be many equally good candidates for the mereological fusion that constitute what we take to be Mt Kilimanjaro: K+ and K- (and presumably, possibly many candidates for each possible borderline or questionable part). Taken broadly, we have a case where there is a disjunction that we take to be true (MK is either identical to K+, or K-), but where neither disjunct seems true (we cannot say it is determinately the case that MK is identical to K+, or identical to K-).[56] So far this mirrors the supervaluational view; determinately, there is a mountain, but there is indeterminacy with respect to its boundary (the mereological sum to which it is identical).

Here it would be worthwhile to briefly retrace our earlier steps. If we take this stipulation seriously, then it would seem that MK is definitely non-identical to K+ or K-; recall our argument from Chapter 1 that vague objects are determinately distinct from any precise aggregate. Perhaps Mt Kilimanjaro just *is* a precise aggregate and it is indeterminate which precise aggregate it is. This would align with the above application of LEM. On the other hand, if we treat Kilimanjaro as a vague object that is determinately distinct from any such precise aggregate, then it is not the case that MK = K+ or that MK = K-. We will return to this issue in **2.4**.

[55] This reading appears (if only tacitly) in Weatherson (2003), Barnes and Williams (2009), and Paganini (2016), among others.

[56] Or, as Simon (2014) puts it, it seems *determinate that* Kilimanjaro is identical with a precise mereological fusion, but *indeterminate which* precise fusion this is.

Another important aspect of this vague-objects view, as put here, is the very notion of a vague part. In virtue of what is a part vague? Here we need to tread carefully. We do not mean that the parthood relation is itself a vague relation, nor do we wish to state that a precise relation of parthood is to be vaguely instantiated, as to speak of properties being vaguely instantiated produces unwanted consequences for their truth-conditions (though see **3.4.1** for a discussion of Rosen and Smith's (2004) view that there may be vague or borderline instances of certain maximally specific "point properties", using the apparatus of fuzzy logic in support of this view). So, if we do not have a vague parthood relation nor the vague instantiation of a precise parthood relation, how do we still have a sense of a vague part?

A common view, articulated earlier, is that there are vague parts insofar as they give rise to indeterminacy with respect to whether an object has some part. A statement representing this might be 'it is indeterminate whether r is part of MK'. Remembering our discussion of the logic of indeterminacy in **1.1**, this would entail both that it is not the case that r is a part of MK, and that it is not the case that r is not a part of MK (*contra* LEM). This entails a particular reading of vague objects as involving unsettledness with respect to their vague parts. However, this is not the only reading available to capture the idea of vague parthood, for it is one with a particular view of vague objects in mind (more on this in **2.4** and Chapter 3). Indeed, given our goal of vindicating an account of vague objects, we can motivate a different, novel account of vague objects and vague parts. Our intent here is exploratory; we do not intend to provide a fully systematic account of such vague objects. Rather, our intent is to explore the prospects of developing a plausible and coherent account of vague objects in response to PM. If this initial presentation is successful, then there is motivation to extend the work in future research. In short, our discussion so far has created conceptual space for the possibility of a vague-objects view; let us determine if a proposed account can fill that space.

Let us introduce 'vague parthood' as a precise relation by which an object may have a part that is distinct from parthood or non-parthood *simpliciter.* Rather than rely on a vague instantiation of a property, or there being unsettledness with respect to whether an object has some part, we instead posit that a vague part is an instantiation of the precise relation of 'vague parthood'. As with any new ideological tool, an explanatory

burden arises to elaborate what this relation of vague parthood is, and to justify its inclusion in our conceptual toolkit.[57]

The proposal is as follows: we postulate two relations that are not clearly distinguished in our ordinary notion of parthood, such that either relation is sufficient for x being part of y. The language of indeterminacy and parthood interact as follows:

- It is determinate that x is part of y iff $x \leq_d y$;
- It is indeterminate whether x is part of y iff $x \leq_v y$.

Here, vague parthood (VP, symbolised as $\leq_v$) is among our basic mereological relations, alongside precise parthood (PP, symbolised as $\leq_d$). These are the two kinds of parthood relations. Since if it is determinate that p, it is not indeterminate whether p, and *vice versa*, vague and precise parthood are mutually exclusive kinds of parthood. In other words, if x is a vague part of y, then it is not the case that x is a precise part of y.[58]

The above biconditionals need not be understood as real definitions; rather, we can take (in)determinate parthood to be primitive. In this case, the invocation of VP and PP is something like a proposal for how to model indeterminacy in mereology. So, we need not take these relations with full ontological seriousness (we might take this something in the way that the modal primitivist takes talk of possible worlds to be a useful fiction, justified because it models how 'possibly' behaves). Alternatively, we could take the right-hand side of the biconditional to be basic, and so take VP to be explanatorily prior to glosses in terms of the parthood relation holding indeterminately. For the sake of convenience, I will proceed as though we should take vague and precise parthood seriously, so that an object will be associated with a set of vague parts and a set of precise parts.

Given extensionality, we then have the following: x is identical to y iff x and y share all vague parts and share all precise parts.[59] As a result, on our view an object is individuated by its set of precise parts *and* by its set of vague parts.

[57] With thanks to Antony Eagle for the following presentation of formalisms and proposal regarding vague and precise parthood.

[58] More formally, $x \leq_v y \leftrightarrow \neg x \leq_d y$.

[59] More formally: $x = y \leftrightarrow \forall z(z\ VP\ x \leftrightarrow z\ VP\ y) \wedge \forall z(z\ PP\ x \leftrightarrow z\ PP\ y)$.

To build on this idea of vague parthood, we can adopt a "precisificational" framework, itself adapted from a broadly supervaluational semantics, whereby an object x is a precisification of an object y iff x's parts are the same as y's parts, and y's precise parts are a subset of x's precise parts. Here we do not take precisification to be a means of making the extension of vague terms more precise; rather, we take precisification to instead involve worlds that determinately have particular candidates as the candidates of the object.[60] We can use the idea of precisifications to further clarify our notion of vague parthood: x is a vague part of y when x features as part of some precisifications of the world (or when x features as a part of some K-candidates), but not as a part of other precisifications (K-candidates). In short, our view of vague parthood broadly analogous to contingency, albeit applied to the case of determinacy.

Further, x is a precise part of y when x features as part of *all* precisifications in the world – when x features as a part of all K-candidates (this is broadly analogous to necessity). With the notions of vague parthood and precise parthood clarified, we can now consider a further idea: a vague object is composed of a *maximal* set of precise parts (here taken as the parts on which all K-candidates agree) and vague parts (here taken as the parts on which some but not all K-candidates agree). Adopting this idea ensures that a single vague object (as desired by our counting intuition) is individuated.

We will continue building on this initial conception of vague parts and vague objects in the remainder of the thesis. What calls for our immediate attention, however, is a reason as to why our response to PM can claim there is a single vague object (as opposed to many), and how this one vague object is relevantly different from the other K-candidates in question.

Unlike Woods and Jones, who take the "many" candidates to *constitute* or *be identical to* one object respectively, our view intends to capture the idea that a vague object is composed of the maximal set of relevant vague parts with respect to some K (where these are parts of some K-candidates, and non-parts of other K-candidates), and the maximal set of relevant precise parts with respect to some K (where these are parts of all K-candidates). We can introduce some new terminology to help unpack this. Specifically, we take it that some X is a *determinate-candidate* (d-candidate) of MK iff all its parts are among the parts of MK, and iff all the precise parts of MK are among

[60] Remaining agnostic with respect to what these precisifications are. See **3.2** for a discussion of various precisificational accounts of metaphysical vagueness, each with a different conception of what a precisification is.

X's precise parts. So, the "many-one" relation we are using here is not one of constitution or identity, but rather, one between MK and its many determinate candidates (its many d-candidates), where MK and the various d-candidates all share the same "core" of precise parts. MK is the unique object composed of this core, as well as the maximal set of vague parts, each of which feature as parts among some but not all of the many d-candidates.

A beneficial consequence of this idea is that it involves a singular identity relation (rather than a many-one identity relation as in Woods); the vague object is simply identical to the precise and vague parts which compose it. Further, our view upholds LEM. We accept that objects either have or do not have parts; it is simply that when an object has a part, this will in some cases be a *vague* part rather than a precise part. In the case of Rocky and MK for example, our view respects the idea that MK either has Rocky as a part or does not have Rocky as a part. However, in accepting that MK has Rocky as a part, we also need to specify whether Rocky is a precise part of MK, or a vague part of MK. Determining whether Rocky is a vague or precise part of MK is achieved using the precisificational picture outlined above, namely, as being dependent on whether Rocky features among all or merely some of the precisifications of MK.

Our view also respects the grounding intuition. The use of vague parts provides us with a means of distinguishing the one vague object from the many d-candidates. At the same time, these d-candidates play a role in determining whether an object's parts are precise or vague (in virtue of featuring in all, or some but not all, d-candidates' extensions). As a result, we can avoid the aforementioned problem of 1002 cats or the problem of 1001 vague cats. The selected vague object uniquely has a relevant feature distinguishing it from the many d-candidates: it is the only object that is composed of *all* precise and vague parts for some K. While other d-candidates may ostensibly exist (comprised of some smaller subset of precise and vague parts), only one object is *maximally* composed of *all* such parts, and so is the unique object selected in response to PM. This is justified through application of the maximality principle. Per maximality, there is one object: that which has no large proper parts of the same object-kind. This selection nonetheless respects the apparent similarity of the many d-candidates, for

the “maximal” vague object is composed of all parts held by all possible d-candidates for the relevant object-kind.[61]

We can build on this view by considering its resemblance to those discussed in **2.2**. Here we are utilising precisifications (motivated by a broadly supervaluational framework) to capture a sense in which d-candidates are more precise, non-maximal objects; further, we make use of a many-one relation (to be elucidated in greater detail in **2.4**) to respect our counting intuitions. Our novel view also respects the logical commitments outlined in Chapter 1, particularly classical logic and the need to have a vague object with a definite set of parts, so as to ensure this vague object can be definitely designated.

An immediate worry: in building this account of vague objects, are we making our view vulnerable to issues regarding problematic instances of vague identity, proper parthood, vague coincidence, or broadly revisionary notions which we explicitly attempted to avoid with our logical conservatism defended in Chapter 1? Even further, some might (aptly!) question whether we can maintain broadly accepted mereological principles and classical logic with the inclusion of the revisionary notion of vague parthood. It is to these questions that we now turn in the remainder of this chapter. **2.4**, **2.5** and **2.6** discuss a pressing objection from Weatherson, considering how vague parts may lead to unwanted revisionary principles and other ideological consequences. We will also elaborate on our novel use of vague parts in relation to vague coincidence, defending the use of such a notion as part of “standard” mereology.

2.4 Weatherson’s Objection: Many Problems

Accepting that Kilimanjaro is a vague material object distinct from both [K+] and [K-] has either metaphysical or logical costs – Brian Weatherson, 2003, p. 491.

Weatherson (2003) presents a powerful argument against the existence of vague objects. He argues that maintaining both ‘classical logic’ (CL) and ‘standard mereology’ (SM) is inconsistent with the existence of vague objects. The argument rests on four

[61] This assumes some agreement with respect to some set of precise parts, whereby the maximal object is the largest object with those precise parts. This aligns (if tacitly) with various descriptions of PM whereby all candidates agree with respect to some “core” set of precise or non-borderline parts (for not every part of objects in PM is considered borderline). For example, consider Lewis’ version of PM involving Tibbles the cat – using his terminology, this “core” would be the set of non-borderline or questionable hairs that have not begun the gradual process of loosening from Tibbles. On our view, this “core” is comprised of those precise parts which are precise parts for all candidates – all candidates agree with respect to some set of precise parts (hence why they are candidates in the first place!)

assumptions that underpin the views above. We will first set out these assumptions before then discussing how our view can uphold these assumptions. This discussion will ensure our view appropriately engages with Weatherson's argument, helping to clarify our use of vague parthood in relation to SM and related notions. We will then elaborate how Weatherson justifies and uses the assumptions to construct his argument, before then defending our conception of vague objects against Weatherson's pointed critique.

To begin, let us precisely state Weatherson's assumptions, as below:

VOD: Vague objects are (determinately) distinct from precise objects (or precise aggregates).
CL: Classical logic, and in particular the Law of Excluded Middle, holds.
Coincidence: If two objects share all the same parts, excluding Rocky, then the objects coincide if Rocky is part of both objects, or if Rocky is not part of both objects.[62]
Coincidence as Identity (CAI): If two objects coincide**,** then they are identical.

We readily accept VOD and CL, as per our earlier analysis from Chapter 1. We take it that vague objects are distinct from precise objects, given that precise objects lack such vagueness (making them determinately distinct from any vague object with, for instance, a vague part).[63] Similarly, we readily accept classical logic and LEM.

The assumptions requiring our immediate attention are Coincidence and CAI. In effect, to accept Coincidence and CAI is to adopt an extensional mereology, and we are happy to remain "classical" in this respect. These assumptions are two aspects of what Weatherson takes to be a part of the package of views comprising SM.[64] Presumably, coincidence refers to complete or strict coincidence, rather than the sense of partial coincidence implied in Lewis' almost-identity solution to the problem.

However, the view we outlined in **2.3** (involving vague parts and vague objects) may seem to give rise to a variety of different kinds of coincidence and non-coincidence; specifically, a possible type of *vague coincidence*. A worry, then, is that we are rejecting standard mereology (here a package of views including Coincidence and CAI) to accommodate our account of vague objects. If this were true, then Weatherson's

[62] In Weatherson's words: 'If for all x other than [Rocky], x is a part of y iff x is a part of z, then if [Rocky] is part of both y and z, or part of neither y nor z, then y and z coincide' (2003, p. 491).

[63] Though see **2.6** for a discussion of Barnes and Williams' (2009) rejection of VOD.

[64] Though there is live debate regarding what can be taken as "standard" mereology – see Varzi (2017).

critique of vague objects would hold, and it would reinforce the idea that one cannot maintain CL, SM, and VOD while holding that vague objects exist. Clearly this is an undesirable outcome for the friend of vague objects. However, we will argue that the notion of vague objects used in our novel view does *not* require the rejection of SM. While this will require the articulation of a cogent notion of *vague coincidence*, we will see that the notion can be cashed out in a manner consistent with the assumptions underpinning Coincidence, CAI, and Weatherson's view of SM.

To help with evaluating Weatherson's argument, let us first consider a series of cases to clarify how our view treats different kinds of coincidence and non-coincidence. This will clarify how our view comports with CAI and Coincidence.

Case 1: Vague Coincidence

Suppose there are two objects x and y. Let us assume both x and y have a number of vague and precise parts – they are vague objects, per our view. Let us further suppose that x and y agree with respect to the parts they hold; in other words, x and y share all parts. At first blush this would seem to be a simple case of coincidence (in a classical sense), given x and y's agreement with respect to what parts they hold. This too would entail that x and y are identical, given CAI. However, let us suppose that x and y disagree with respect to *how* they hold these parts; namely, while x and y agree with respect to *what* parts they have, they disagree with respect to which parts are precise parts, and which are vague parts. In short, they disagree with respect to how they hold their parts. On our view, then, x and y are strictly non-identical, for they differ with respect to which of their parts are vague and which are precise parts. Given CAI, then, it follows that x and y are in fact non-coincident.

Problem: have we beset ourselves with a contradiction in claiming x and y to be both coincident and non-coincident? Not quite. Rather than posit that x and y are coincident *simpliciter* (and so leave ourselves open to a contradiction), we can instead posit that x and y are *vaguely coincident*. Here we can supply a notion of *vague coincidence*, whereby two objects are vaguely coincident when they share the same parts, but disagree with respect to *how* they respectively possess these parts (the manner in which they hold these parts). Coincidence *simpliciter* on our view, then, is agreement both with respect to parts held *and* with respect to how the parts are possessed by each object. Importantly, this last point is consistent with Weatherson's portrayal of Coincidence – we do not wish to modify any of the four principles he has used to set-

up his argument. Our view rather makes use of the new notion of vague coincidence to capture the respect in which objects can possess parts in particular ways (vaguely or precisely). A crucial upshot of this, then, is that we are able to maintain the original idea of coincidence *simpliciter* when objects are identical, and so our view also upholds CAI.

This posit of vague coincidence also enables a clear articulation of when two objects are *non-coincident* with one another, namely, when objects disagree with respect to *what* parts they hold. Our view also respects the intuitive idea that if two objects are non-coincident, then they are determinately distinct from one another. Further, if two objects are vaguely coincident from one another, then we hold that they too are determinately distinct. In cases of vague coincidence however, the source of this distinctness is not due to a difference in what parts these objects hold; rather, it is due to the fact that while both objects agree with respect to what parts they hold (they share all parts), they disagree with respect to which of these parts are precise and which are vague.

An analogy may be useful here. Recall our discussion of Lewis' (1993) view in **2.2.2**, in which we made use of a similar analogy: consider the difference between two objects completely overlapping, two objects partially overlapping, and two objects being completely disjoint from one another. The first idea relates to the sharing of all parts, the second to the sharing of some but not all parts, and the last to the sharing of no parts. Our above articulation of coincidence, vague coincidence, and non-coincidence (respectively) correspond with these ideas. The idea doing the work here is the respect in which parts are possessed by the objects (precisely or vaguely), and how this can give rise to vaguely coincident objects.

With the above articulated, we can move to discussion of a second case. This will further clarify our understanding of coincidence and identity when considered in the context of producing an account of vague objects in response to PM. As with our discussion above, explicating how coincidence and identity work in this context will allow us to better clarify our novel view, while also considering how this understanding of coincidence and identity may respect our understanding of these notions in non-vague contexts.

Case 2: Extra Vague Part

Suppose there are two vague objects x and y, and suppose they are comprised of a set of precise and vague parts. Let us assume that x and y agree with respect to their precise parts – they share all precise parts. However, suppose x has a vague part that y definitely lacks as a part. Besides this one vague part, x and y otherwise agree with respect to all parts they have, and how they hold these parts. So, the only respect in which x and y disagree pertains to one vague part which x possesses that y does not possess. This example gives rise to a possible worry: if the only difference between x and y is the possession of a *vague* part, then how can we be confident that this entails that x and y are *determinately* distinct, and so determinately non-coincident? Consider that the only difference between the two involves a part which is putatively vague – would this perhaps lead to the objects being vaguely coincident? To accept Case 2 as an instance of vague coincidence would be to reject CAI, and to require revision to our strict use of identity as articulated above.

However, we need not accept Case 2 as an instance of vague coincidence. Rather, we instead take Case 2 as an example where two objects are determinately *non-coincident*. The reason for this is, simply, that x and y disagree with respect to the parts they hold. This is how we envisage non-coincidence in regular contexts – if an object has a part that another lacks, then they are non-coincident and so distinct.

Case 2 is quite different from Case 1, as in Case 1 all parts were shared by x and y; the only difference between the two objects was the manner in which the parts were held, namely, which parts were vague and which were precise for each object respectively. By contrast, Case 2 illustrates a situation where one object (x) determinately possesses a vague part that the other (y) determinately lacks. Though this additional part is a vague part, we can still maintain that this is not a case of vague coincidence, for as we stipulated earlier, vague coincidence pertains to disagreement *only* with respect to how parts are held. Vague coincidence still requires two objects to agree with respect to *what* parts there are, lest this be a case of determinate *non-coincidence*. As Case 2 involves disagreement with respect to a part possessed by one object compared to another, even as a vague part, per CAI the two objects are determinately distinct.

The takeaway from the above two cases is that we have clarified a sense of how objects may coincide even when considering vague parts. This has been developed

by thinking about coincidence with respect to what parts objects possess, and also with the additional idea of *how* objects hold these parts (the manner in which these parts are held). This latter idea is key to our novel view of vague objects and vague parts; this is where our view differs from usual accounts of coincidence *simpliciter* (where, presumably, in non-vague contexts it is held that objects can only have or lack parts precisely or determinately). We can summarise these ideas, motivated by our novel account of vague objects, as follows:

If x and y do not share all parts, then they are non-coincident, and so are determinately distinct.

If x and y share all parts, and agree with respect to how these parts are held (the same parts are precise and vague for both x and y), then they are coincident, and so are determinately identical.

If x and y share all parts, but disagree with respect to how these parts are held (different parts are precise and vague for x and y respectively, and so they disagree with respect to which parts are precise and vague), then they are *vaguely coincident*, and so are determinately distinct (per CAI).

What is the upshot of this discussion? To begin, we can now see that our novel conception of vague objects plausibly enables us to accept the four assumptions underpinning Weatherson's argument while still maintaining the cogency of the existence of vague objects. In short, we have provided an initial framework for a novel account of vague objects which can respect Coincidence, CAI, VOD, and CL (and presumably SM more broadly).

So, given that our view respects Coincidence and CAI, it is now prudent to evaluate how Weatherson wielded these assumptions to argue against the existence of vague objects. As we agree with Weatherson's assumptions, we can engage with the strongest form of his argument and fully consider its implications for our novel view of vague objects as a response to PM.

Following Barnes and Williams (2009, p. 180), and Paganini (2016, p 460), we can reconstruct Weatherson's argument into the following step argument:[65]

1) Either Rocky is part of MK, or it is not a part of MK.
   (Assumption: LEM)
2) If Rocky is part of MK, then MK coincides with K+.
   (Assumption: Coincidence)
3) If MK coincides with MK+, then MK = K+
   (Assumption: CAI)
4) If Rocky is part of MK, then MK = K+
   (From 2, 3)
5) If Rocky is not part of MK, then MK coincides with K-
   (Assumption: Coincidence)
6) If MK coincides with K-, then MK = K-
   (Assumption: CAI)
7) If Rocky is not part of MK, then MK = K-
   (From 5, 6)
8) Either MK = K- or MK = K+
   (From 1, 4, 7).
9) If there is a vague object MK, then MK ≠ K+ and MK ≠ K-,
   (VOD)
10) There is no vague object MK
   (8, 9, Modus Tollens)

We agree that Weatherson's argument holds, and so presents a strong counterargument against the existence of vague objects. It seems that one cannot endorse the existence of vague objects while maintaining VOD, CL, and SM, on pain of a *reductio*. It appears that we are in trouble if we wish to sustain our defence of the existence of vague objects. Yet, if we accept Weatherson's argument, one may wonder how we intend to sustain our defence of vague objects. It is to such a response that we now turn.

[65] Strictly speaking, this presentation differs from Barnes and Williams', as it includes steps 9 and 10. Barnes and Williams do not include these steps, for their argument rejects VOD. Nonetheless, they show that the view can be problematic even without this assumption in hand.

2.5 Responding to Weatherson: Vague Objects and Vague Parts

Weatherson's argument is effective on a standard view of coincidence, whereby the sharing of parts is sufficient for coincidence. On this view, Weatherson's argument is valid. Though we accept this, what we intend to show is that its conclusion holds only if one accepts the *same* conception of vague objects. The important detail here is that Weatherson's argument is effective against a *particular conception* of what vague objects are. Specifically, as stated earlier, Weatherson's argument makes use of a notion of vague objects and vague parts, whereby it follows from an object's having a vague part that there is unsettledness with respect to whether an object possesses a particular part. This was motivated by LEM, and the idea that an object either has or does not have a part in any given instance. Given the desirability of retaining LEM, this seems an important assumption to uphold.

As our above discussion of our view of vague parts demonstrated however, we can maintain LEM while defending a different account of vague objects. Specifically, we can argue that while Rocky is a part of MK (and so sustain LEM), it is a *vague* part possessed by MK in a particular manner, namely, as one that features in some but not all d-candidates of MK. As **2.4** demonstrated, the sharing of parts is *insufficient* for coincidence – objects also to share the manner in which these parts are held for there to be coincidence. So, our response to the Weatherson argument is simple: we reject Premise 2. If Rocky is part of MK, then it is not the case that MK coincides with MK+, for Rocky may be a vague or precise part of MK. Importantly, this is a view that comports with LEM – the initial motivation Weatherson had for creating Premise 2 remains on our view, even as we reject Premise 2. Equipped with a notion of vague coincidence (built from a notion of vague parthood), there is a clear means by which we can reject the Weatherson proof.

Here one may worry that we have changed the target of the Weatherson proof, given that it targets a particular conception of vague objects. We can resist this thought by recalling our discussion of the Evans proof in **1.1** and **1.2**, whereby we accepted the force of the Evans proof, but argued that what it ultimately demonstrated was that *if* a vague-objects theorist were committed to vague identity, *then* they would be led to a troublesome conclusion. This led us to conclude that the Evans proof *restricted* the options available to the vague-objects theorist, namely, to constructions of vague objects that avoided the use of vague identity, or to otherwise find a means of resisting the proof itself.

Applied here, the moral to take from Weatherson's argument is similar: *if* a vague-objects theorist wished to sustain the view of vague parts as producing unsettledness between precise states of affairs (as on Weatherson's view), *then* they could not sustain CL, VOD, and SM. In other words, if one wanted to uphold such an account of vague objects with these assumptions in hand, then they would need to find a means of blocking a step in Weatherson's argument. Given our novel conception of vague objects however, we need not follow this line of thought. Instead, our response to Weatherson's argument is to simply reject Premise 2, given that it is reliant on a standard view of coincidence, which does not hold on our conception of vague objects. Therefore, while Weatherson's proof demonstrates that *one* conception of vague objects is in trouble, this does not entail that all conceptions of vague objects face a similar issue.

Here it is important to not move too quickly. It is prudent to also consider the prospects of the other option, namely, of blocking a step in the proof. We will now set to this task in **2.6**, where we will consider Barnes and Williams' response to Weatherson's proof that accepts his conception of vague objects.

2.6 Barnes and Williams' Response to Weatherson: Vague Objects and Vague Identity

Barnes and Williams resist Weatherson's argument by arguing that the existence of vague objects need not lead to vague identity (or, rather, vague identity that is susceptible to Evans' *reductio*). They do so by accepting CL and SM, but by dropping VOD, the assumption that vague objects are determinately distinct from precise objects. Their argument pertains to steps 1 through 8 of Weatherson's argument (which is presented as such in their paper). For Barnes and Williams, Weatherson's argument is a clear indication that vague identity 'seems inescapable' (2009, p. 180) for the vague-objects theorist who wishes to retain CL and SM while invoking the notion of vague parts.[66] Given that Barnes and Williams reject VOD, their argument needs to

[66] Paganini (2016) offers a different response to Weatherson's argument, claiming that one can endorse the existence of vague objects while accepting CL and SM, and without the need to induce vague identity. Importantly however, Paganini's view relies on two points which we have limited reason to accept: first, her argument relies on reality being gunky (which she claims SM remains silent on); second, the argument relies on a modification of Weatherson's 'Coincidence' to instead read as follows: 'if two things have the same parts, setting aside Rocky, then they coincide if Rocky is part of both or if Rocky *is disjoint from both'* (2016, p. 463, emphasis in original). Even if Paganini's argument holds, it relies on a particularised reading of SM as involving gunk, and more significantly, relies on interpreting Weatherson's argument in a manner distinct from his own initial presentation. As we intend to defend

instead rely on a method of blocking Evans' *reductio* against vague identity to sustain their view of vague objects.

Barnes and Williams do just this; specifically, they resist the Evans *reductio* against vague identity by blocking the step in the proof that uses lambda abstraction (from P1 to P2). This is typically the move put forward by the semanticist, who takes the terms involved in the proof to be referentially shifty. They normally argue that one can have referential indeterminacy with shifty designators in the proof (typically grounded in a view of vagueness as semantic indecision – recall our discussion in **1.2**). Barnes and Williams instead posit the idea that one can have referential indeterminacy with a *metaphysical* source, as opposed to a semantic source (we will unpack their account in more detail in **3.2.1**). Broadly, Barnes and Williams hold that the source of referential indeterminacy can come from the world itself, namely, that there may be cases where the world does not settle the reference relation of terms. As a result, or so they posit, there may be referential indeterminacy from the "worldly" aspect of reference-fixing functions.[67] Let us accept this at present. Here, then, Barnes and Williams' view is not susceptible to the Evans-Salmon *reductio*; they can avoid the unwanted conclusion that the existence of vague objects leads to the acceptance of a contradiction by blocking the Evans proof in a manner similar to the semanticist.

From this, Barnes and Williams draw a specific lesson: while an account of vague objects that makes use of vague parts may give rise to vague identity (per Weatherson's argument) this need not give rise to problematic *de re* vague identity, so long as one is willing to reject VOD and rely on an ontic sense of referential indeterminacy. While at first glance this may seem a satisfactory conclusion for the friend of vague objects, as Barnes and Williams note, this strategy leads to the following thought: 'we cannot have a vague object *together with* a plenitude of precise objects' (2009, p. 183). This has significance for the vague-objects theorist more broadly, and also for our novel response to PM – indeed, Barnes and Williams explicitly state that their argument 'casts doubt on whether vague objects can serve as a solution to the problem of the many' (2009, fn. 19). Why is this so?

Per Barnes and Williams, and assuming the rejection of VOD, the existence of vague objects would lead to there being objects that are *vaguely identical* to precise objects

vague objects while adopting the strongest interpretation of Weatherson's argument, we will not adopt Paganini's strategy (tenable or otherwise) in response to Weatherson.

[67] See Barnes and Williams (2009, pp. 181-182) for the presentation of such an example.

in such a way that is susceptible to the Evans *reductio*. After all, if we reject VOD, then we cannot assume that vague objects are determinately distinct from precise objects. As vague objects certainly seem determinately non-identical to precise objects, it appears the only option remaining is to bite the bullet and accept that they must be vaguely identical to precise objects, if they are to exist at all. Given that Barnes and Williams provide a viable response to the Evans *reductio*, this may not seem problematic. Yet, what runs into trouble is not merely the cogency of vague identity, but rather, the implications this has for our metaphysics. Given that the view of vague objects here makes use of vague parts, it is potentially susceptible to rejecting the principle of 'De Re Determinate Universal Fusion' (Barnes and Williams 2009, p. 182):

De Re Determinate Universal Fusion (DUF): 'Whenever there are some objects, there is an object which is *determinately* their sum' (Barnes and Williams 2009, p. 182, emphasis in original).

The combination of DUF and the acceptance of *de re* vague identity leads the proponent of vague objects to be vulnerable to the Evans proof, specifically in attempting to provide a means of distinguishing a single vague object from the putatively vaguely identical precise candidates involved in PM. Per DUF, any vague object has a determinate sum. However, per Barnes and Williams' argument, the vague-objects theorist needs to rely on the use of vague identity to resist Weatherson's argument. Combining these commitments, then, we have a view ostensibly committed to objects with determinate sums, yet one where such objects are also vaguely identical to any (or perhaps many) other objects, such as those invoked in examples of PM. We are left with endorsing determinate identity criteria, but also the ostensible need for vague identity. Thus, the upshot of the Barnes and Williams argument is that the vague-objects theorist either needs to reject DUF, or to reject the use of *de re* vague identity in their account of vague objects.[68] Barnes and Williams ultimately defend the use of *de dicto* vague identity, and as we will see in **3.2.1**, construct a plausible view of vague objects that relies on just this. Our view instead rejects the need to rely on vague identity (*de re* or *de dicto)* as part of a view of vague objects;

[68] Barnes and Williams conclude by noting that there is a 'stand-off' between those defending a 'plenitude of vague objects', and those defending a 'plenitude of precise objects' (2009, p. 183). For our purposes however, we are interested in the implications of their argument specifically for the vague objects-theorist, and so we leave resolution of the stand-off between precise and vague ontology to the side.

recall that we explicitly decoupled our view of vague objects from *de re* vague identity in Chapter 1.

Barnes and Williams' argument is certainly significant for our proposed view of vague objects, given that some may hold we have tacitly defended DUF in endorsing the possibility of there being definite designators for vague objects. The important lesson to draw from this argument pertains to our vague-objects solution to PM. Specifically, while we have proposed a means of selecting one vague object (such as MK) from a plenitude of d-candidates, the need to maintain definite identity conditions remains as important as ever. The grounds with which we select the one vague object from the many need to sustain the idea that there is one determinate sum from purportedly many objects. Given that we have embedded the notions of vague parts and vague coincidence into our view, and as these produce definite criteria for one maximal object to satisfy, we can present a view that comports with Weatherson's argument and with the lessons drawn from Barnes and Williams.

This is not to become complacent however, as this does lead to the upshot that our view endorses that vague objects which have determinate sums. This may seem counter-intuitive, particularly given the central role that vague parts and vague coincidence play in our account of vague objects. An apt criticism here may be that, in holding a vague object to have a determinate sum, we are merely naming relations as "vague" parts and "vague" coincidence and characterising entities as "vague" objects while addressing a different, non-vague phenomenon altogether. Indeed, part of the typical understanding of vagueness involves borderline cases, or cases where there is no "fact of the matter" as to the determinate status of some phenomenon. Here then, we must clearly retrace our steps in constructing the view. We have not merely posited relations and categories in a manner distinct from the debates regarding PM and vague identity; rather, the very characterisation of vague objects used in both debates have motivated our conception of vague parts. Recall that the use of "borderline parts" appears in various conceptions of vague objects, and recall that Lewis' set-up to PM (and his criticism of a vague-objects solution) invokes the same notion. The notions used to build our view connect with the dialectic established in the literature, and so hold with respect to capturing a sense of vague objects and vague parthood.

However, this is a distinct view from upholding the existence of vague objects as a genuine instance of the broader notions of "metaphysical vagueness" or "worldly indeterminacy". While we have linked vague parts to vague objects within a broadly

classical logical setting, we have yet to establish a connection between vague objects and the broader literature surrounding the construction of cogent views of mereological indeterminacy and metaphysical vagueness. We set to this task in Chapter 3, and in doing so, engage with a variety of contemporary views whom have also presented accounts seeking to present responses to PM and cogent accounts of metaphysical vagueness.

2.7 Taking Stock

This chapter has presented a novel vague-objects solution to the Problem of the Many. We first articulated the problem in **2.1** and then evaluated various existing solutions in **2.2**, including discussion of "supervaluationist" and "many-one" solutions. We elucidated the merits of each view, and considered the ways in which they attempted to capture our counting and grounding intuitions regarding material objects. We then constructed a novel vague-objects solution in **2.3**, arguing that such a view respects these intuitions while upholding classical logic and standard mereology. We then responded to Weatherson's argument against vague objects in **2.4** and **2.5**, and also considered Barnes and Williams' alternative rejoinder to Weatherson in **2.6**. While responding to Weatherson's argument, we considered how notions of vague parthood and vague coincidence can provide a plausible means by which we can select one vague object from many candidates. Importantly, these concepts were constructed using the machinery of precisifications and other independently plausible notions. We did not develop gerrymandered notions in response to PM; rather, we built these notions both from the considerations that led to PM in the first place, and also from independently plausible machinery used in other contexts involving vagueness. In short, we have defended an account of vague objects that retains desired principles of logic and metaphysics, while also respecting the intuitions about ordinary objects that PM challenged. Given that much of the literature previously disregarded a vague-objects solution on the grounds of ontological profligacy, or due to the view merely leading to a new "vague" example of PM, our novel view has made genuine progress.

However, work remains. We have offered a broad overview of this novel view, rather than a systematic account, for our task so far has merely been to vindicate the plausibility of vague objects' existence, and to offer a means of presenting a viable solution to PM. If successful, this would motivate the development of a more systematic metaphysics, mereology, and logic to underpin the proposed view. So, to assess the plausibility and viability of our novel account, our immediate next step is to first evaluate

how precisificational machinery has been used in other contemporary contexts. This will allow us to determine the theoretical virtues and costs of precisifications as used in our novel account of vague objects, and in its use as a solution to PM. Further, by analysing precisificational accounts of metaphysical vagueness as well as non-precisificational accounts, we will be in a better position to see how vague objects can support the defence of the intelligibility of metaphysical vagueness more broadly. We can then use this analysis to consider how precisifications can be used to present a cogent account of vague objects. By expanding on the initial gloss presented in this Chapter, we can consider the prospects of our account as a potentially fruitful avenue for further research. It is to this task that we now turn in Chapter 3.

Chapter 3 – Vague Objects and Metaphysical Vagueness

3.0 Preamble

This chapter contextualises the thesis (and the novel solution to PM proposed in Chapter 2) in relation to the contemporary literature and existing dialectic pertaining to "metaphysical vagueness". I discuss contemporary accounts of metaphysical vagueness, which can be broadly categorised into two classes of views: those defending "precisificational" accounts, and those defending "non-precisificational" accounts. I then evaluate the relative merits and strengths of these accounts, particularly the manner in which they conceive of vague objects and their ability to respond to PM. In doing so, I compare these views to the novel account of vague objects presented in Chapter 2. I end the chapter by considering how the novel account of vague objects can be extended in relation to broader issues pertaining to vague and indeterminate phenomena. In short, this chapter ends by situating the thesis within the contemporary discourse on metaphysical vagueness.

3.1 Vague Objects and Metaphysical Vagueness

A number of philosophers have sought to explain the phenomenon of vagueness in metaphysical terms by providing intelligible accounts of metaphysical vagueness (or metaphysical indeterminacy).[69] By considering these accounts, we hope to draw out the strengths of their proposed explanations, and in particular, how their accounts of metaphysical vagueness apply to instances of PM.

Here it is important to clarify the purpose of this chapter. Recall that we are explicitly engaged in a project seeking to vindicate the existence of vague objects, and with providing a cogent response to PM. However, many accounts of metaphysical vagueness begin with quite different theoretical motivations. While these accounts engage with considerations arising from PM (and often present a picture of what a

[69] Hereafter we will use the terms "metaphysical vagueness" and "metaphysical indeterminacy" interchangeably, so as to capture the phenomena addressed by those in the literature without marking a strong distinction between the two (see Barnes (2010a) for a similar view). In saying this, some, notably Eklund (2008, 2011, 2013), mark a clearer distinction between the two phrases as representing distinct phenomena (typically due to vagueness being seen as intimately connected to sorites-susceptibility).

vague object is), in many cases their scope of inquiry often involves attempting to explain a broader range of phenomena relating to the development of intelligible accounts of metaphysical vagueness.[70] Our target is narrower: we wish to defend the existence of vague objects. If we offer a coherent picture of vague objects, then we have provided an intelligible example of metaphysical vagueness. By defending the plausibility of the former, we demonstrate the cogency of the latter.[71]

Further, we do not intend to offer a fully reductive characterisation or explanation of what vague objects are *simpliciter*; we are not attempting to characterise *all* instances of vague objects or metaphysical vagueness. Rather, we follow the orthodoxy and accept that indeterminacy with respect to an object's spatial boundary and parts – what is sometimes called *mereological indeterminacy* – is a sufficient condition for being a vague object.[72] This orthodoxy is highlighted by Abasnezhad and Hosseini (2014), who note that 'among the various indeterminacies, the lack of a precise spatial boundary seems to be a common feature of *ordinary* vague objects' (p. 241, emphasis added). Here, then, our view aligns with the orthodox conception of vague objects as objects possessing imprecise spatial boundaries; in our novel view, this is a result of a vague object being (partially) composed of vague parts.

To better engage with the views to be discussed, let us first consider a general formulation of metaphysical vagueness. Barnes and Williams offer a general characterisation of 'metaphysical indeterminacy' as arising from a 'generic concept of indefiniteness', where: 'it is metaphysically indeterminate whether p iff (1) it is indefinite whether p, and (2) the source of this indefiniteness is the non-representational world' (2011a, p. 108). Here we can see that while this characterisation remains silent with respect to the *nature* or *grounds* of metaphysical indeterminacy (what produces examples of metaphysical indeterminacy), it nonetheless highlights a key feature of the idea: indefiniteness with a non-representational source. We take this as exemplifying a broad feature of metaphysical vagueness. So, if we show that vague objects have their vagueness non-representationally, or that there could be vagueness

[70] Examples include indeterminacy of colour, the open future, or even certain types of indeterminacy arising from quantum mechanics.

[71] Though the proponent of metaphysical vagueness may not be fully satisfied by this; they may have explanatory aims which our narrower defence of metaphysical vagueness, derived from our vindication of vague objects, does not achieve.

[72] Of course, mereological indeterminacy may entail indeterminacy in other related properties, such as that with location, shape, size, etc.

that would remain ‘in the absence of language and thought’ (Merricks 2017, p. 225), then we would have a clear example of metaphysical vagueness.

Finally, we do not intend to “resolve” once and for all whether vagueness itself should be best understood as a worldly phenomenon (in all or some cases). Rather, our intent is to consider the cogency of various accounts of vague objects as examples of metaphysical vagueness, and to determine their applicability to the pertinent issue of responding to PM. Succeeding in this regard would suffice to motivate the development of a more systematic account of metaphysical vagueness that engages with deeper questions relating to the nature of vagueness itself.

3.2 Precisificational Accounts of Metaphysical Vagueness

Our intent in distinguishing precisificational from non-precisificational views is due to the use of precisifications as a means of explaining vague phenomena. While some have divided the field into other categories,[73] the purpose of our distinction here is to highlight the differing conceptual machinery used by the views to be discussed. So, our immediate focus is on accounts of metaphysical vagueness that rely on precisificational machinery; we will engage with non-precisificational accounts in **3.4**.

Broadly, precisificational accounts treat vagueness using a metaphysical analogue of supervaluationism’s “precisifications”. Perhaps unsurprisingly, different precisificational accounts in the literature define and use precisifications in quite distinct ways – we will consider two such views.[74] Specifically, we will evaluate these accounts and their diverse attempts to account for the existence of vague objects, and consider how this may apply to the considerations that arose from our earlier discussion of PM. In short, this section teases out important lessons to use in a sustained defence of vague objects; we will apply these lessons to reinforce our novel

[73] For instance, Wilson (2013) distinguishes between “meta-level” views and her own “object-level” view to distinguish different accounts regarding where vagueness is situated. As our focus will be to discuss each view themselves as a whole, our broad categorisation should be reasonably innocuous for our present purposes.

[74] One precisificational account with which we do not engage is Abasnezhad and Hosseini’s “supervaluational” account of metaphysical vagueness (2014). They adopt a view which has structural similarities to Barnes and Williams’ account, though they use an ontological variant of a broadly “Finean”, rather than “Lewisian”, approach to vagueness.

account of vague objects, which makes use of precisifications, but shares a similar structure to non-precisificational views.[75]

3.2.1 Barnes and Williams' "Unsettledness" View

On Barnes and Williams' view (hereafter "B&W"), vagueness is a result of the world being unsettled between distinct, yet precise, states of affairs.[76] For B&W these states of affairs are ersatz possible worlds; different representations of ways the actual concrete world (the one we inhabit!) might actually be. We can see this as a metaphysical offshoot of Lewis-style supervaluationism. The difference between the two is that B&W's possible worlds represent different precise ways the world could be, rather than different precise ways the extensions of terms could be. B&W also adopt an "ersatzist" theory of possible worlds, on which possible worlds are 'abstract objects which represent (classically complete) ways the world might be' (2011a, p. 114). These possible worlds are accessible from the actual world via the addition of another "division" among accessible worlds: those which are 'precisificationally possible' (2011a, p. 115). B&W treat these precisifications as actual possible worlds; precisifications are accessible from the 'actual, vague world' insofar as they represent more precise ways the world might be (2011a, p. 115). B&W also make use of the notion of the "actualised world": the abstract world (the precisificationally possible world) that represents things as they are in the concrete, actual world, rather than as a representation of things as they actually are not.

Using these concepts, B&W develop an account of metaphysical vagueness in which vagueness is a product of the world being unsettled between different precisifications – a product of unsettledness between maximally specific and precise possible states of affairs (possible ways the world might be). So, while we exist in the concrete, actual world, there is unsettledness with respect to which precisificational possibility is the actualised world, namely, the precisification that represents things as they are in the actual world. In the words of B&W:

> We think that metaphysical indeterminacy consists in a fundamental kind of unsettledness in the world. When p is metaphysically indeterminate, there are two possible (exhaustive, exclusive) states of affairs – the state of affairs that p and the

[75] So, we consider our view to be broadly non-precisificational despite our use of precisifications in developing the notions of vague parthood and vague coincidence. We will discuss this in greater detail in **3.5**.

[76] Barnes and Williams (2011a). See also Barnes and Williams (2009, 2011b), Barnes (2010b) and Williams (2008b, 2008c).

> state of affairs that not-p – and it is simply unsettled which in fact obtains. No further explication is possible or needed (2011a, p. 112).

A concrete example may help to clarify this picture. Recall our Mt Kilimanjaro (MK) case in Chapter 2, where we discussed the mountain's possession of a borderline or fuzzy part. Per B&W, in this example there would be two precisificationally possible worlds, each with a different maximal and precise state of affairs in which MK's possession of the part is settled. In one world there is a precise state of affairs where MK definitely has the part in question, while in the other world there is a precise state of affairs where MK definitely does not have the part. Here the actual world is unsettled between these two states of affairs; importantly, there is no further "vague object" that determinately possesses a "vague part". There is also no "vague world" above and beyond the two worlds which realise MK in these different ways. In each world, MK is identical to one of the objects comprised entirely of precise parts; it is simply unsettled which precisification represents the actualised world, and so which precise object correctly represents MK as it is in the actual world. So, B&W's view gives rise to cases where a disjunction is settled (an object is identical to some precise aggregate), but where each disjunct is unsettled (it is unsettled, and so indeterminate, the precise aggregate to which the object is in fact identical). If there were no vagueness in the world, then there would be just one precise state of affairs: the actual world. Per B&W, then, multiple precisificationally possible worlds only exist due to the presence of vagueness; it is the presence of vagueness that gives rise to the possibility of there being different precise ways the actual world can be represented.

Here metaphysical vagueness is a product of the world being unsettled between different possible precisifications of reality. While the actual concrete world we inhabit is comprised of various concrete entities and objects, these abstract possible worlds (precisifications) merely represent different ways the concrete world might be. Of note is that each world is taken to be maximal and fully complete: each precisified world is "settled" and fully precise. There is no vagueness *within* a world or precisificational possibility – there is only vagueness as a product of unsettledness in terms of which precisificational correctly represents the actual world. We will return to this idea later in this section.

A strength of B&W's view is its ability to avoid instances of "propositional indeterminacy", or the need to introduce a "third" category or truth-value.[77] As each precisificationally possible world is fully determinate, and as the actual world is ultimately identical to one of these precise states of affairs, there is no extra "indeterminate" category required to explain this vagueness. Further, B&W treat the unsettledness characteristic of the view as a kind of conceptual primitive. As such, there is 'no reductive characterisation' (2011a, p. 108) or 'more fundamental story' (2011a, p. 112) that needs to be told regarding metaphysical vagueness. Their view enables one to retain a fully bivalent system with truth and falsity about various matters. There is no "vague" or "indeterminate" state of affairs; rather, indeterminacy arises due to different worlds disagreeing with respect to how the actual world is.

Another virtue of B&W's account is that it enables one to remain silent with respect to their metaphysical commitments.[78] There is no "extra" type of parthood, object, or metaphysical category that is added to their ontology, nor is there a need to adopt a deviant or revisionary understanding of ordinary objects. The unsettledness used in B&W's account indicates that the world simply does not meet the demands for some state of affairs to be determinately true or false. This lack of commitment to additional ontological or ideological machinery allows B&W to avoid undue commitment to deviant metaphysics – a point in their favour. So, if B&W's explanatory utility and ability to respond to objections is comparable to that of other views, then the view will be quite favourably situated in the contemporary literature.

However, there are two pressing objections to B&W's view which may give us pause. Indeed, while their account is admirable in its aim to intelligibly model metaphysical vagueness, it has also generated a number of critical responses. First, B&W's view has been charged with not delivering a "genuine" account of vague objects. Second, B&W's response to PM runs into similar issues as supervaluationism did in producing peculiar results with respect to our intuitions regarding ordinary objects. We will unpack each criticism in turn, and then consider how the view compares to our novel view of vague objects.

Abasnezhad and Jenkins (2018) argue that B&W do not deliver a genuine account of vague objects, for their view relies on 'unsettledness of reference, *while maintaining*

[77] See Barnes and Cameron (2016) for an extended discussion of these virtues.

[78] Barnes and Cameron (2016, p. 125).

that all worldly candidate referents are perfectly precise' (p. 281, emphasis in original). For Abasnezhad and Jenkins, B&W's vagueness is derivative from the 'interface between representation and reality' (2018, p. 281), in the sense of being derived from unsettledness between different precisifications (the admissible possible worlds which represent ways the actual concrete world might be). They argue this is not itself vagueness derived *from* an object, and so it is not, strictly speaking, a theory of vague objects (2018, p. 280). Even if B&W offer a theory of metaphysical vagueness that can *concern* objects, for Abasnezhad and Jenkins this is not itself a theory *of vague objects* in the sense of the vagueness being derived from features of the object.

Abasnezhad and Jenkins press this, noting that B&W's account generates metaphysical vagueness that only involves *precise* states of affairs and objects (2018, p. 280). Unsettledness between precise states of affairs is not vagueness in, or derived from, an object; B&W's objects themselves ultimately have a definite boundary and a definite set of parts. Indeed, this aligns with a more recent defence of B&W in Barnes and Cameron's 'Are There Indeterminate States of Affairs? No' (2016). In this paper, Barnes and Cameron explicitly detach their "Unsettledness View" of vagueness from the idea of defending the existence of some kind of indeterminate states of affairs (vague objects being a corollary of this). We can clarify this with an example: consider the proposition 'it is indeterminate whether a is F'. Here B&W do not set out to endorse metaphysical vagueness as arising from there being vague mereological structure with respect to a (taken to be, for instance, an object or state of affairs); rather, they are primarily concerned with providing metaphysical grounds for indeterminate property possession (the 'F' in this proposition') as grounds for metaphysical vagueness.

Here B&W could respond by noting that not all theories of metaphysical vagueness need deliver vagueness "from an object" – the dialectic does not force the friend of metaphysical vagueness to accept vague objects as existent things. Yet, some may hold it as intuitive that metaphysical vagueness would involve the brute fact that the world *itself* (the actual world) just *is* imprecise, and that any "precise" representation would be a mere abstraction of an actually imprecise world. Surely, then, an intelligible view of metaphysical vagueness ought to maintain such an intuition.

We could imagine B&W replying that their view, if suitably interpreted, could deliver a theory of vague objects. Consider the proposition 'a is F' – where a is some object, and F is some property. Per B&W, if a proposition is vague, then it is true according to one precisification and false according to a second precisification. Here B&W could

perhaps defend the idea that they capture a relevant sense in which a just *is* a vague object – it is just a picture that involves an unsettled reference (given that the actual world and all precisifications refer to the *same object*; it may just be indeterminate whether the object to which we precisely refer has a particular property). This seems a feasible reply, for B&W could argue that their view makes room for there being "vague objects" with settled referential features, even if the view does not comport with there being additional "indeterminate states of affairs" or imprecision "at a world".

Even accepting Abasnezhad and Jenkins' critique however, it is far from a knockdown argument against B&W. A thesis of vague objects can be separated from an account of metaphysical vagueness, with B&W presenting a clear example of an account that (arguably) defends the latter without explicitly endorsing the former. Yet, while this in itself may not be problematic, a problem arises when we extend such a view to issues concerning PM. Specifically, the method by which B&W block the Evans *reductio*, and the way in which they treat objects (and states of affairs), leads to a criticism of how their view responds to PM.

Recall our discussion in **2.6,** where we noted that B&W block the Evans proof by rejecting the step involving lambda abstraction (the move from P1 to P2);[79] let us expand on this discussion in view of their account of metaphysical vagueness. This method of blocking the Evans proof is similar to Lewis' supervaluationism; one can hold that the terms involved in the Evans proof are (analogically) shifty or non-rigid, enabling appeal to referential indeterminacy to escape Evans' problematic conclusion. The key with B&W is that this shifty reference is not due to semantic considerations; rather, it is due to the world being unsettled between different referents. Their picture of metaphysical vagueness argues that the world is unsettled with respect to different precisifications – different admissible worldly representations of how a state of affairs may be. The upshot of this, for B&W, is that the Evans argument can be blocked if there is referential indeterminacy with an *ontic* source. This is precisely the result they intend, for they remark (with reference to the treatment of vague identity offered by Williams (2008b): 'The Evans-Salmon argument targets only *de re* indeterminate identity, not all metaphysically vague identities' (Barnes and Williams 2009, p.181). For example, if the statement 'MK has Rocky as a part' has a metaphysically indeterminate extension (being well-defined when Rocky is part of MK, and failing to refer to an object

[79] For the sake of clarity, P1 and P2 are reiterated here:
(P1) $\nabla\ (\mathrm{a} = \mathrm{b})$
(P2) $\lambda x[\nabla(\mathrm{x} = \mathrm{a})]\mathrm{b}$

otherwise), then the term is (analogically) shifty. This leads to the following being true: 'it is indeterminate whether MK is the mountain that has Rocky as a part'. So, Barnes and Williams endorse the idea that there can be metaphysically vague identity that is resistant to Evans-style reasoning so long as it is not *de re* indeterminate identity, but rather, *de dicto* indeterminate identity, i.e. referential indeterminacy.[80] In a phrase: B&W defend the cogency of referentially vague identity statements that have a metaphysical source.[81]

Given B&W's similarity to supervaluationism (both in its structure and in its method of blocking the Evans *reductio*), it is pertinent to recall our earlier critique of supervaluationism in responding to PM in **2.2.2**. Supervaluationism fails to offer us a non-arbitrary means of selecting one candidate over another. B&W is susceptible to a similarly motivated critique, for their view does not deliver the result of a particular precise aggregate with which to respond to PM. While there is presumably a single object present at each precisification, and a single object in the actual world, B&W do not offer a means of producing a determinate referent for some unique object above and beyond the many candidates. For example, as B&W accept that their account delivers indeterminate identity statements, they would hold that MK is *not* determinately distinct from each precisificationally possible state of affairs; rather, MK is identical to one, it is simply indeterminate to which it is identical. While accepting this sense of metaphysically indeterminate identity enables resistance to the Evans view, it does not produce a response to PM which generates a unique object as the correct "one" presumed by our counting intuitions. Perhaps this is to be an expected result of an account invoking a fundamental kind of unsettledness. Indeed, perhaps the lesson to learn is that accounts of metaphysical vagueness, when applied to PM, will not easily allow for a "settled" answer. Nonetheless, this should motivate a closer look at views which more strongly respect our counting intuitions, such as the novel account of vague objects we have proposed. It is surely desirable to respect our counting intuition and provide a means of determinately selecting one object from the purportedly many

[80] See also Barnes (2009) for a counterpart-theoretic response to the Evans argument – a different means of blocking the Evans proof's problematic conclusion. We will leave discussion of this strategy to the side.

[81] Though see Akiba (2015b) for an argument that Evans' *reductio* against vague identity can be reformulated to still present trouble for B&W. Consider that Salmon's (1981) version of the Evans argument uses variables, rather than singular terms; Akiba utilises this to generate problematic instances of vague identity that do not invoke referential indeterminacy. B&W may be susceptible to such examples, though see Abasnezhad (2016) for a defence of their view. Abasnezhad notes that B&W can still account for 'indeterminate identity between referentially determinate objects [RDI]' (2016, p. 197) which are still resistant the Evans proof, though he argues this comes at the expense of explanatory utility in characterizing vague objects which do not generate RDI (2016, p. 201).

generated in instances of PM, something which B&W ultimately cannot achieve – a point in our favour. Yet, how does B&W compare to our view in other respects?

First, one may aptly consider B&W to be a relatively metaphysically austere view compared to our proposed novel view (which involves the ostensibly revisionary notion of vague parthood). Indeed, B&W's ability to remain largely "neutral" with respect to their metaphysics is a virtue they themselves propound. Yet, there is reason to think that our proposed view is not unfavourably revisionary. While it appears that our view relies on extra onto-ideological machinery, note that our notion of vague coincidence comes "for free" from our notion of vague parthood. While our notion of vague parthood is revisionary (given that it involves concretising indeterminate parthood into a determinate mereological structure with two distinct kinds of parthood relation), this itself is not radically revisionary.[82] Further, our novel view of vague objects is developed using the familiar notions of precisifications and coincidence (the former of which also features in B&W, and the latter of which is a part of standard mereology more broadly). So, our "added" onto-ideological machinery involves the use of otherwise largely acceptable machinery. If you believe that the use of precisifications is valid, then our notions should be similarly tractable as well.

Second, while B&W uses precisifications to treat vagueness as something holding in virtue of a relationship *between* worlds, our novel view uses precisifications to ground vagueness in the *actual* world. Consider: while B&W ground vagueness in the *unsettledness between* different properties an object might have had (or unsettledness between different states of affairs), our view instead grounds vague objects in the actual world. We ground vagueness in a vague-part relation. While B&W entails that the actual world is in fact precise (and it is merely unsettled which way it is precise), our view captures the intuitive aspect of metaphysical vagueness as involving the brute fact that the world *itself* just *is* imprecise, and that any "precise" representation would be a mere abstraction of an actually imprecise world. If this intuition has force, then retaining the intuition is another point in our view's favour.

None of this is to present a knockdown argument against B&W as an account of metaphysical vagueness;[83] rather, it is intended to demonstrate some of the virtues of

[82] Consider the comparable move involved in modelling metaphysical incompatibility as a different kind of relation compared to logical incompatibility.

[83] Particularly given that the cogency of vague objects is of secondary importance for their view – consider that the vague objects which their view would generate are "representationally vague" (so as to avoid the Evans proof), and so are of a different kind to the type of vague objects we defend.

our account, as highlighted by our consideration of PM. Indeed, B&W's idea of being "balanced" or "unsettled" between different possible ways things could be is not the guiding thought behind the vague objects used in response to PM (which themselves actually do seem to exist in addition to the many determinate candidates). Consider too that B&W's view is primarily situated towards the development of an intelligible view of indeterminate properties more broadly, rather than the presentation of a response to PM. This may illustrate the difference in focus between our accounts and the respective explanatory strengths of each view.

This discussion of B&W has sought to highlight how our novel view retains the virtues gained from using precisificational machinery, but also possess the virtue of having the ability to ground vagueness in the actual world while respecting our counting intuitions in relation to PM. We will discuss these virtues in greater detail in **3.4.2**, where we compare our view to Wilson's non-precisificational account of metaphysical vagueness, and also in **3.5**, where we consider the prospects for developing our novel view further. Yet, in talking about the virtues of using precisifications, it is worthwhile to consider how another precisificational account fares in accounting for metaphysical vagueness, and in responding to PM.

3.2.2 Akiba's Precisificational Aggregation Account

On Akiba's precisificational account of metaphysical vagueness,[84] precisifications are not abstract representations of the world; rather, precisifications are part of reality itself. Akiba endorses the idea that there is a "precisificational dimension" along which objects run (analogous to spatial and temporal dimensions).[85] Vagueness arises when objects partly coincide along the precisificational dimension, representing the precisificational extensions of objects – ways these objects may be made precise. The motivation here is that we are already familiar with ways in which objects may coincide spatially, and how they can perhaps coincide temporally (temporarily or otherwise).[86] Akiba extends this idea by positing that objects can also coincide precisificationally. So, on his view, precise objects retain their spatio-temporal profile across different precisificational dimensions (precise objects cannot be made *more* precise, after all!)

[84] Akiba (2000, 2004). See also his (2015a) for comparisons between his view and B&W.

[85] Consider Akiba's words: 'in whatever sense we believe in the first five dimensions, we may also believe in the precisificational dimension. You may take the existence of these dimensions literally, or you may take it instrumentally, as a *façon de parler*.' (2004, p. 408).

[86] Or, perhaps, familiar with how interlocuters argue that objects may coincide in various respects.

However, vague objects do not share this fundamentally precise nature, and so vague objects 'coincide with different precise objects in different precisifications' (Akiba 2000, p. 366). While both vague and precise objects act as 'transworld objects' on Akiba's view (both extend over various precisificational dimensions), it is only vague objects that have different spatio-temporal extensions in different precisificational dimensions (given that they are able to be made precise in a number of ways). For Akiba, an object is as much composed of spatial and temporal parts as it is an 'aggregation of its precisifications', represented as different slices of the precisificational dimension. Further, precisifications are taken as discrete parts of reality (running along particular spatial, temporal, and precisificational dimensions). As vague objects may coincide with different precise areas of sets of constitutive matter on different precisifications, there is vagueness insofar as objects may coincide in some but not all precisificational dimensions. Thus, Akiba can hold that vague objects are determinately distinct from other objects, for they have unique extensions across spatio-temporal and precisificational dimensions.

Akiba's view of vague objects shares a similarity with our view: both views separate the existence of vague objects from reliance on a sense of vague identity. For Akiba a vague object 'coincides with different precise areas in different precisified worlds; it is not identical with any of [the precise objects]' (2004, p. 417). Indeed, this enables his view to avoid the issues faced by B&W in replying to PM while endorsing some notion of indeterminate identity. Yet, Akiba does note that his view involves 'admittedly heavy onto-ideological machinery' (2004, p. 414). He defends this cost with the thought that it has the virtue of retaining bivalence and classical logic – a reasonable point, on balance. What of Akiba's application of his view towards PM?

Akiba's response to PM begins in a familiar fashion: he distinguishes the one vague object from the various candidates comprising the many. Akiba's example involves the vague cat Tibbles, and the many candidate "p-cats", or "precise cats" which do not have any borderline parts. Remember that Tibbles and the many p-cats share the feature of being transworld objects; they all extend across various dimensions (the various different ways the world can be made precise). The key difference comes in how the vague object differs in each precisified world (understood in analogy with metaphysically possible worlds):

> But while each p-cat occupies transworld same space-time region in all precisified worlds, Tibbles occupies slightly different space-time regions in different precisified

> worlds, coinciding with each p-cat in each world. This is what it means to say that the cat Tibbles has vague boundaries, and that *none of the hairs h1, h2, …, and h1000 determinately does, or determinately does not, belong to Tibbles* (Akiba 2004, p. 418, emphasis added).

We have emphasised the latter aspect of Akiba's description of vague objects, as it is where the view differs from our novel response to PM. Akiba holds that the spacetime region which Tibbles the vague object occupies is variable, changing depending on the precisificational dimension in which she occupies. Vague objects have "indeterminate" boundaries in the sense that 'their locations are slightly different in different precisified worlds' (Akiba 2004, p. 411). This seems satisfactory with respect to distinguishing a vague object from a precise object, but it does not explain why the p-cats are not themselves cats. Akiba grants this, stating in a footnote that his account would need to be supplemented with an 'exclusion principle to explain why the p-cats are not themselves cats despite sharing very close spatio-temporal profiles to that of the vague cat. Akiba briefly proposes a general exclusion principle, which he puts as 'something like this: if there is an object o of kind k (e.g. cat) in a given region, then there is no object of the same kind that significantly overlaps o' (2004, fn. 22). Here Akiba is offering a principle resembling that of a maximality principle, albeit one applied to his particular conception of what a vague object is (where overlap is indexed to a particular region). Akiba also takes into account the possible need to hedge the principle for cases such as Theseus' ship, where there may be multiple things at the same place for a certain time, for they may differ in modal/temporal properties. Even with this qualification in mind, does the combination of Akiba's view of vague objects, and his attempt at presenting an exclusion principle, satisfactorily explain our grounding intuitions within the explanatory apparatus of his view?

I argue there is reason to be cautious; we can see this through comparison to our novel vague-objects response to PM. Where Akiba's view differs with our account is with respect to what is considered to be determinately and indeterminately part of the vague object itself.[87] Recall the example of the vague mountain Mt Kilimanjaro. Akiba's vague mountain is 'a transworld object that coincides with different precise areas in different precisified worlds' (2004, p. 411). Akiba uses this idea to construct a sense of the "penumbral area" of Mt Everest; the analogue on our view is the set of vague parts of the mountain. For Akiba, this "penumbral area" is neither determinately a part nor

[87] Moving beyond the obvious ideological and ontological differences between our views.

determinately not a part of the mountain (2004, p. 410). Hence, if we were to ask of Akiba's view whether we could determine the space-time region with which the vague mountain coincides, Akiba's view would admit of indeterminacy in responding to this query. Expanding on this further, Akiba's view entails that such a penumbral area is indeterminate, as its inclusion as part of the mountain's spatio-temporal profile shifts depending on which precisified world the mountain occupies. This has implications for responding to PM, as while Akiba's view provides the means of selecting one object from many, it does not provide a means of selecting one vague object from purportedly many. On Akiba's view this is exemplified by there being many slightly different precisificational extensions present from world to world. For example, if we had to select one specific spatio-temporal boundary for MK in response to PM, this would differ depending on our location along the precisificational dimension (the boundary would overlap with MK1 in one world, MK2 in another world, etc.). So, for a vague object there is no single spatio-temporal boundary that coincides across all worlds. Indeed, this was part of Akiba's very definition of a vague object – they are those that have variable spatio-temporal boundaries across different worlds (whereas precise objects have static spatio-temporal boundaries when extended across the precisificational dimension). Given that there is no determinately vague boundary which we can select, then, Akiba's view is susceptible to the paradox of 1001 vague cats. While we have vague objects, we do not have a single object boundary to distinguish between many vague candidates.

Here one may baulk at the suggestion that PM should require the need to select a determinate boundary from the many candidates posed by the paradox. Yet, our proposed novel view can do just this; we can determine the set of parts to which we can identify with the vague mountain. Rather than rely on a "penumbral area" of indeterminate parts, our view instead grounds this in the various vague parts determinately possessed by the vague mountain. We instead hold that it is *determinate* that this area of vague parts is determinately part of the mountain's space-time region. This is enabled on our picture of vague parts and vague coincidence – one constructed using precisificational machinery so as to retain classical logic and avoid propositional indeterminacy, but also a view which extends this machinery to provide greater explanatory power with respect to PM.

So, the broad issue with Akiba's view is that we are left with an instance where there is no determinate spatio-temporal boundary to which we can identify with a vague object – there is an absence of grounds for determinately selecting one particular

boundary for the vague object. The ability to provide these grounds is a key strength of our novel account; we can state that the vague object has a definite set of parts, and that some of these parts are vague parts, understood via use of precisificational machinery. We can present a single vague object with a clear set of parts comprising its boundary, whereas Akiba's view leaves us with a plenitude of many possible boundaries. So, despite its many virtues, Akiba's view, too, falls short in responding to PM.

3.3 Lessons from Precisificational Accounts

As we have seen, a clear strength of precisificational views is their ability to broadly accept classical logic, and their ability to avoid the posit of a "third" vague truth-value or instantiation of a property. The reliance on precisificational machinery enabled these accounts of metaphysical vagueness to maintain similar virtues to supervaluationism, as applied to worldly cases. However, as we saw above, each precisificational view discussed came with various costs. For B&W, their view was challenged in delivering a genuine account of vague objects, and more significantly, had relatively limited explanatory power with respect to PM. On the other hand, Akiba's view was susceptible to issues surrounding the paradox of 1001 vague cats, and also carried relatively heavyweight onto-ideological costs in requiring the acceptance of a "precisificational dimension" existing in addition to dimensions of space and time.

In saying this, both views were shown to be viable options in their own right; there is clear merit in using precisifications to construct an account of metaphysical vagueness. From the discussion of the relative strengths of these precisificational accounts, we are now in a better position to see the merit in our novel account's use of precisificational machinery. Further, in discussing the limitations of these views, we can also appreciate the novel account's extension of precisificational machinery as a means of developing cogent notions of vague parthood and vague coincidence so as to produce a satisfactory response to PM – particularly in respecting our counting intuition and in providing a tenable response to the paradox of 1001 vague cats. By using precisifications in this way, our novel solution produces a clearer means of demarcating one vague cat from purportedly many vague cats by providing a single, clearly delineated boundary for the one vague object.

Here it would be prudent to also consider non-precisificational accounts of metaphysical vagueness. In eschewing the use of precisifications, these accounts offer

comparatively robust metaphysical pictures in developing their distinct accounts of metaphysical vagueness. It is to these views that we now turn in **3.4**.[88] Yet, as we will see, in presenting unique, novel, and substantial conceptions of what vague objects may be, these non-precisificational accounts also face issues of their own. A further comparison between our novel account, and these non-precisificational views, will demonstrate the need to be prudent when revising one's metaphysics.

3.4 Non-Precisificational Accounts of Metaphysical Vagueness

3.4.1 Rosen and Smith's Fuzzy View

Rosen and Smith (2004) (hereafter R&S) defend there being vagueness in the world by considering how there may be vague properties, relations, and 'objects in a certain respect' (p. 185). Specifically, their view seeks to make sense of the idea that 'the world contains vague objects: not just fuzzy properties but fuzzy *things*' (2004, p. 187, emphasis in original). To present an account of 'fuzzy things', R&S distinguish between properties *simpliciter* and maximally specific 'point properties' (2004, p. 189). For instance, an example of the former is the colour 'blue', while an example of the latter is some particular shade of blue, perhaps named 'B-17' (Rosen and Smith 2004, p. 188). Per R&S, an object is vague – a fuzzy thing – if it vaguely instantiates a 'point property' – a maximally specific property. An elaboration of some representative cases may be helpful here.

First, we can think of a think of a ball that is definitely blue and definitely a particular shade of blue. Second, we can think of a ball that is vaguely blue and vaguely a particular shade of some colour. Third, we can think of a ball that is definitely blue but is vaguely a particular shade of blue – it vaguely instantiates a point property such as 'B-17'. R&S would treat the third case as an instance of a "vague object", but not the second. What distinguishes the two cases is that only the vague instantiation of maximally specific point properties qualifies vagueness to be of the object, whereas vagueness with respect to more general properties (such as "blue") is not constitutive of a vague object (even if it may still be a genuine instance of worldly vagueness).

[88] A non-precisificational account we do not discuss comes from Sattig (see his 2010, 2013, 2014, 2015). Broadly, Sattig's view endorses the existence of vague objects, and of 'mereological indeterminacy *de re'*, in order to construct a plausible response to PM. His account relies on acceptance of his 'quasi-hylomorphic ontology', and also relies on a certain type of 'formal indeterminacy *de re*' to produce a cogent account of vague objects. Due to reasons of scope, we leave his view to the side.

Importantly for R&S, properties such as location, mass, and length can count as the right type of maximally specific properties which give rise to vague objects if vaguely instantiated.[89] So, applied to a case of PM, a vague object may arise in such a case that 'there is a spacetime point p such that it is neither determinately the case nor determinately not the case that it falls within [an object's] boundaries' (2004, p. 187), or where an object is 'indeterminate in mass by being indeterminate in composition' (2004, p. 105).[90] Adopting R&S' broadly fuzzy view of vagueness, a point property may be vaguely instantiated insofar as it is not a degree-1 nor degree-0 instance of that property. So, there may be a series of spacetime points that are determinately part of MK's boundary (degree-1), some that are determinately not part of MK's boundary (degree-0), and other spacetime points that, to a non-determinate degree (between 0 and 1), fall within MK's boundaries.

Per R&S, then, there is a distinct vague object (MK) with a series of vaguely instantiated point properties (spacetime points falling on MK's boundary). While these vague objects involve the vague instantiation of point properties, where each point property would be instantiated to a particular degree between 0 and 1, for R&S this does not result in vague objects having "particular" properties that the object would possess. In their words: 'if we told you that we had an object that was genuinely indeterminate in size or mass, you should take us to mean that there was no particular or determinate mass property or size property that the thing determinately possessed' (2004, p. 188). So, the vague instantiation of some property, even perhaps to a precise degree between 0 and 1, does not result in the "determinate" instantiation of some property. On their view, then, while a vague object may indeterminately instantiate a point property such as mass, the object does not determinately possess any such property. A consequence of this is that R&S are susceptible to the criticism we levelled to Akiba, namely, their view does not provide a single boundary for vague objects; rather, a vague object would indeterminately instantiate such a property (such as possession of a precise location or set of spacetime points) without determinately possessing any such property. As with Akiba, this is an issue with some significance

[89] Rosen and Smith further develop their account of point properties (2004, pp. 188-195); however, for our purposes the gloss presented here should suffice.

[90] Here Rosen and Smith consider their view of metaphysical vagueness (and vague objects) as one that comports with a world 'composed by a perfectly determinate array of democritean atoms' which are themselves well-defined with respect to various 'fundamental physical parameters' (2004, p. 195). In short, their account runs on the idea that indeterminacy 'emerges, as it were, from a more fundamental stratum of utterly determinate fact' (2004, p. 196).

when considering issues arising from PM and the need to respect our counting intuitions with respect to ordinary objects.

Here one could perhaps envisage R&S responding that this would be an expected result from a view espousing a form of metaphysical vagueness grounded in fuzzy logic. Indeed, the absence of such a boundary would seem to comport nicely with the vague instantiation of some point property of mass or location. Further, this is a response that aligns with the argumentative scope of R&S' paper, for they merely wish to demonstrate that 'if the fuzzy view is correct', then they can 'make sense' of the intelligibility of one type of metaphysical vagueness, rather than 'argue that [metaphysical vagueness] exists' (2004, pp. 186-187). While the absence of clearly delineated boundaries for ordinary objects is a drawback, perhaps it can be readily accepted as part of R&S' fuzzy view of the world; perhaps we should expect such results for our metaphysics when wielding a fuzzy logic.

So, we can readily agree that R&S have succeeded in developing a view of metaphysical vagueness that, at the least, "makes sense". Further, R&S did not require the use of precisificational machinery to do so – there is good work to be done in considering alternatives to precisificational machinery. Perhaps unsurprisingly however, we demur with respect to their use of fuzzy logic, for our thesis has explicitly defended commitment to the use of classical logic (per **1.3**). Further, the logical apparatus one uses has clear implications for our metaphysics – accepting R&S' use of fuzzy logic gives rise to a plenitude of other fuzzy notions, such as vaguely instantiated properties. If we do not wish to construe our metaphysics as such, then an alternative view will be preferable. Indeed, if we can retain classical logic, then we can minimise the need to revise our metaphysics as a result of our chosen logic. There is thus motivation to explore other views which more closely comport with classical logic and non-fuzzy metaphysics. We turn to such an account in **3.4.2** by looking at Jessica Wilson's "determinable-based" account. Rather than apply fuzzy logic to the instantiation of properties and other worldly items, Wilson's account instead makes use of a particular metaphysical apparatus to present her own intelligible and explanatorily useful account of vagueness. Let us evaluate the prospects of her view.

3.4.2 Wilson and Indeterminate States of Affairs

Jessica Wilson presents a detailed non-precisificational account of metaphysical vagueness.[91] Wilson's account is, specifically, a "determinable-based" account of metaphysical vagueness; a view focusing on the idea that certain states of affairs can be underdetermined or overdetermined, leading to there being no unique determinate of a determinable at some level of determination. In Wilson's words:

> What it is for an SOA [state of affairs] to be MI [metaphysically vague] in a given respect *R* at a time *t* is for the SOA to constitutively involve an object (more generally, entity) *O* such that (i) *O* has a determinable property *P* at t, and (ii) for some level *L* of determination of *P*, *O* does not have a *unique* level-*L* determinate of *P* at *t* (2014, p. 366, emphasis added).

Wilson defends her account with reference to an example of 'relativised' determination. For instance, the 'highly iridescent feathers of a hummingbird gorget' appear to genuinely have different colours at the same time, depending upon an observer's angle of perception (2013, pp. 367-368). Though Wilson initially motivates her account with this account of colour, we will not rely on referring to this view; rather, what we are interested in is the idea that some level of determination may produce the absence of a unique determinate of some determinable.

Consider Wilson's view as applied to vagueness with respect to macro-object boundaries of the kind associated with PM. Here Wilson's determinate-determinable picture can be applied (see 2013, p. 377; 2016, p. 108):

> What it is for a macro-object O to have an indeterminate boundary is for it to be determinately the case (or just plain true) that (i) O has a determinable boundary property P but (ii) O does not have a unique determinate of P at t.

This may initially seem puzzling; how does an object have a determinable boundary property with no determinate of that determinable? To understand this initially peculiar idea, it is important to note that Wilson's account rests on the idea that "macro-objects" depend on "micro-aggregates". To provide a gloss of the view, features of these macro-objects are "realised" by (possibly many) of these micro-aggregates. So, object boundaries for things like MK are 'intimately dependent upon – realised by multiple

[91] Wilson refers to the existence of "indeterminate states of affairs" rather than "vague objects"; however, we will treat the two as relatively analogous for our purposes (leaving aside the particularities of the ontology of states of affairs, objects, etc. aside).

lower-level aggregates and their properties' (2016, p. 108). On this view it is not the case that there is one micro-aggregate for an object like MK; rather, MK's boundary is "multiply realised" by a number of micro-aggregates.[92]

This account is supported by Wilson's unique reading of the determinate/determinable relation. As Wilson quite candidly states, she takes there to be cases where determinables may not have unique determinates for some level of determination. So, while objects may have determinable properties (such as having a boundary, or being the colour blue), they may not have a unique determinate of that property (having one specific or singularly delineated boundary, having a particular shade of blue, etc.). For example, a mountain may have a determinable boundary property, but no unique determinate for that determinable (such as the determinate property of possessing a precise number of electrons). In Wilson's words: 'it is reasonable to assume that Mount Everest has a single determinable boundary property, which is determined, at any given time, by *multiple more specific boundary properties* (2016, p. 109, emphasis added). Here, then, Mt Everest has a determinable boundary that is "multiply determined" or "multiply realised" by there being 'many such realising micro-boundary properties associated with a given determinable macro-boundary at a time' (2013, p. 374).

Applied to PM, Wilson states that her account offers a 'straightforward" answer to PM: 'there is just one cat there: the one with the determinable boundary' (2016, p. 378). Here Wilson assumes that the various cat-candidates, or cat-constituters, have precise boundaries, and so do not have the same claim to cathood as the vague object with the determinable boundary that lacks a unique determinate.[93] Here our view has a notable advantage. Consider an example of PM where all but one candidate was to suddenly disappear; we are left with just one candidate. In this situation, Wilson's view would still need to distinguish the determinate cat-constituter from the "mere" determinable property – there would be apparent coincidence between a determinate and determinable. On the other hand, our view can treat this example by saying that the one remaining constituter is just identical to the object being constituted, for the set of "common" parts of the constituters would just be identical to those possessed by the

[92] See Wilson (2009, 2012) for a fuller explication of Wilson's view of the determinate/determinable relation.

[93] Wilson also presents an 'optional deeper answer' by adding that 'a cat-constituter cannot survive certain changes to its boundary that the cat it constituters can survive … the cat's existence and features are compatible with the determinable boundary's being determined by the determinate boundary of any appropriate candidate cat-constituter' (2016, p. 379).

one remaining constituter. So, the constituted object would only have precise parts, given that there would be no disagreement between multiple candidate and so no vague parts. By extensionality, the constituted object would be identical to the one remaining constituter. This result seems more plausible then Wilson's, which requires the peculiar posit of a determinate coinciding with a determinable in such a situation.

So, Wilson's account relies on the vague cat being determinately distinct from various precise candidates – we can readily agree on this. However, Wilson holds that all cat-constituters (or cat-candidates) are precise – she eschews the idea that there may be such things as vague candidates (indeed, this distinction is used to distinguish Tibbles the vague cat from the many cat-candidates/cat-constituters). As a result, Wilson's view is also susceptible to worries arising from the problem of 1001 vague cats, for in the absence of further principles, one could envisage many suitably construed vague cats with highly similar "multiply determined" determinable boundaries.

Though Wilson does not explicitly engage with this derivative problem, we can envisage a possible reply based on her account. Specifically, Wilson's view may just reject the possibility of there being many candidate determinable boundaries for some object – while it perhaps makes sense to speak of many candidate determinates (precise sets of micro-aggregates), Wilson's view may hold that there can only be one such determinable property. If so, then this would suffice for there being only one vague object where there was previously thought to be grounds for many, namely, that there are not many closely related determinable boundary properties in the way that there can be many closely related determinate properties. The grounds for rejecting there being many determinable boundary properties are unique to Wilson's view. Indeed, the grounds for this reply seem strong compared to the idea available to the precisificational view, namely, of rejecting there being many candidate vague objects in the precisificational framework.

Issues arising from the paradox of the 1001 cats notwithstanding, Wilson's view has a number of virtues,[94] and also some clear structural similarities to our novel account. First, her account does not rely on using an entirely different set of logical or metaphysical tools in embracing non-classical logic or (for instance), or the adoption of an additional precisificational dimension to be added to our regular understanding of spatio-temporal dimensions. Instead, her view relies only on revising one particular

[94] For instance, see Barnes and Cameron (2016) and Wilson (2016) for discussion.

notion from our metaphysical toolkit – the determinate-determinable relation. Wilson's view remains reasonably non-revisionary with respect to much of our logic and metaphysics.

Further, Wilson's account presents a systematic and *reductive* account of vagueness, in the sense that it offers an explanation for how and why things may be vague with the use of the familiar conceptual notions of determinate and determinable properties. This enables a clear sense in which the account offers an explanation of metaphysical vagueness in terms of other familiar concepts; Wilson offers a sense of "what it is" for something to be metaphysically indeterminate. There is no fundamental primitive or conceptual primitive doing the explanatory work. Wilson's view also respects our counting intuitions with respect to PM, for there is just one object present – the one with the determinable boundary property. Further, Wilson's view also seems to respect our desire to retain strict identity conditions for vague objects – her response to the Evans proof is to deny that her view gives rise to vague identity conditions, as we also defended earlier in the thesis. Wilson's view defends the idea that one can have vague objects (or metaphysically indeterminate states of affairs) without inducing vague identity.

Wilson (2016) also defends her view against the charge that it is a "Third-Way View" (Barnes and Cameron, 2016) by noting that her view does not include any propositional indeterminacy. For Wilson, it is not propositions that are indeterminate; rather, it is state of affairs that are indeterminate. She defends this by noting that for any state of affairs which her view renders indeterminate, this will nonetheless produce determinate propositions. For instance, she holds that 'it is true that Mount Everest has the determinable boundary property, and for any unrelativised determinate boundary property, it is false that Mount Everest has that property' (2016, p. 110). As with our view, Wilson does not introduce any third truth-values or categories which would 'suggest that propositions are ever anything other than true or false' (2016, p. 110); metaphysical vagueness need not affect our semantics by inducing indeterminacy in our propositions.

However, Wilson's view is susceptible to criticisms of its own. Specifically, Wilson's view requires revision with respect to the intuitive thought that determinables have unique or precise determinates, requiring revision to the idea that an object may definitely have a boundary, but that there is no unique fact of the matter (it is underdetermined or overdetermined) what precisely instantiates this boundary. Wilson

explains such a picture in other work (2009, 2012), but here, then, we are required to revise our initial understanding of the determinate/determinable relationship.

As mentioned, a virtue of Wilson's view is its systematicity, in that her view seeks to provide a clear explanation of *what it is* for something to be metaphysically vague for all such phenomena. Yet, this very systematicity leaves Wilson open to a pointed objection, namely, that if there is one instance of metaphysical vagueness that does not fit within her determinable-based view, then the claim to systematicity fails. [95] In short, her view needs to accommodate all instances of metaphysical vagueness to secure the advantage of systematicity. Wilson accepts this point (2016, pp. 113-114), replying that her view can suitably account for all relevant phenomena. [96] By comparison, our view is more modest, seeking only to account for vague objects and indeterminate spatial boundaries while remaining silent with respect to other phenomena such as the open future and quantum indeterminacy. As a result, while our view does not possess the virtue of offering a systematic account of metaphysical vagueness, we avoid the worry of over-extending with respect to our explanatory remit.

Let us summarise the points of comparison between Wilson's view and the proposed novel solution defended by this thesis. Both views respect the counting intuition of PM (there is one object present), and both similarly present a means of explaining the apparent metaphysical vagueness as a product of something attributable to the object (its determinate property for Wilson, its set of parts on our account). However, our view's "revisionary" notions of vague parthood and vague coincidence are derived from the more readily accepted machinery of precisifications, rather than Wilson's view relying on a particular and specialised reading of determinates and determinables. Our novel account also presents a clearer response to the paradox of 1001 vague cats, and do so without the need to reconsider intuitive thoughts regarding the determinate-determinable properties – points in our view's favour. In saying this, both views, on balance, share a number of structural similarities in locating vagueness in the actual world (rather than in a precisificational dimension, or in a representation of the actual world).

[95] See Barnes and Cameron (2016, pp. 129-131).

[96] See Wilson and Calosi (2019) for recent work attempting to apply Wilson's account towards issues surrounding quantum phenomena.

3.5 Lessons for a Defence of Vague Objects

Our analysis of precisificational and non-precisificational accounts of vague objects has helped with identifying a number of virtues, and also some explanatory gaps, in existing accounts of metaphysical vagueness. The virtue of using precisifications, as seen in B&W and Akiba's view, is that they provide conceptual machinery with which to respect classical logic and provide a means of capturing the sense in which vagueness gives rise to the possibility of things being made more precise. However, as our analysis demonstrated, precisificational views often struggled with providing satisfactory responses to PM which respected our counting intuitions, or in responding to the paradox of 1001 cats. The very precisificational machinery that provided these views with their explanatory power also limited their ability to respond to PM. So, the lesson to draw from these accounts is to take note of the virtues of using precisifications, but to consider how their use can be extended to present a cogent response to PM with a more robust metaphysics supporting our conception of vague objects. This discussion motivated an exploration of non-precisificational accounts of metaphysical vagueness.

The non-precisificational accounts discussed were able to provide a clear connection between their view and the attempt to clarify how vague objects exist in the world. While R&S were able to bridge the gap in offering a clear sense of how a property may be vaguely instantiated, this came at the expense of relying on non-classical logic and a broader "fuzzy" metaphysics. On the other hand, while Wilson's view avoided such issues, and enabled us to avoid issues pertaining to propositional indeterminacy, her view nonetheless produced the view that there are determinable properties with no unique determinates. Even granting this to be a viable picture, it requires granting the use of multiply relativised determination to let such a view get off the ground. From discussion of non-precisificational accounts, then, the lesson to draw is while furnishing our metaphysics can provide additional explanatory power, it should be mediated where applicable to avoid additional commitment to revisionary notions.

From our discussion we can better see the virtues of our novel view. We adopt a broadly Wilson-esque approach to providing criteria for what it is for an object to be vague in the actual world (possessing a set of definite and vague parts on our view); crucially however, we also built our account using the machinery of precisifications while adopting this approach. We aimed to provide a clearer explanatory link between vague objects and their closely related candidates and precise counterparts, using

precisifications to build notions of vague parthood and vague coincidence (rather than the more revisionary notion of a multiply determined determinable). The notions we use play a clear role in enabling the selection of one vague object from many candidates, enabling the provision of a “many-one” response that considers the relationship between a vague object and its many “d-candidates”. Doing so also enables our novel view to respect our counting intuition and avoid problematic instances of selecting one among many vague objects – our view selects a single vague object based on the unique maximal property it possesses. Our view presents a satisfactory response to the paradox of 1001 cats – an issue to which other accounts had difficulty responding.

So, ultimately our view can be taken as a broad synthesis of existing views; possible only as a result of the good work that has recently been done with respect to responding to PM (including newly created “many-one” solutions) and producing accounts of metaphysical vagueness (adopting a similar approach to metaphysical vagueness as Wilson, but using similar precisificational machinery to B&W in the model). Issues arising from discussion of this view merit an extended discussion (such as a fuller account of the semantics of our view, and of the metaphysical status and ground of precisifications); here we have only touched on these issues briefly. Nonetheless, our hope is that by beginning with vindication, and presenting a view that builds on the good work and progress made in recent decades, we can then motivate further work extending this thesis to provide a more developed treatment of these issues in constructing a systematic account of vague objects.

3.6 Taking Stock – Looking Forward

This chapter has considered existing accounts of metaphysical vagueness, and in particular, how they present accounts of vague objects that comport with our metaphysics and logic. We first considered the distinction between vague objects and metaphysical vagueness in **3.1**, before evaluating precisificational and non-precisificational accounts of metaphysical vagueness in **3.2** and **3.4** respectively. We showed that each view discussed offer intelligible and cogent accounts of metaphysical vagueness, but that each required revisionary metaphysics or logic, or presented unsatisfying answers with respect to PM or derivative paradoxes. **3.3** and **3.5** articulated the comparative virtues of our novel view, demonstrated how it built from the precisificational machinery used by views in **3.2**, and how it built from Wilson’s broad approach to metaphysical vagueness in **3.4.2**. In short, we have situated our

view in relation to the contemporary literature, showing that our construction of vague objects is intelligible and, we argue, favourably positioned to provide a cogent account of vague objects, and to provide a solution to PM.

Yet, there are still a number of live issues in the literature on which we have remained silent due to reasons of space. For instance, if we endorse the cogency of metaphysical vagueness, then it may be important to determine whether this vagueness features at the “fundamental” bedrock of reality, or whether it is merely a derivative phenomenon from more precise components (see Barnes 2014). Another fruitful line of enquiry engages with quantum mechanics, in which various interpretations have been raised to attempt to capture a sense of “metaphysical indeterminacy” in line with the possession of properties such as spin and location (see Bokulich 2014; Wolff 2015; Wilson and Calosi 2019).

We have also left room for more systematic work to be completed. Specifically, there is a need to provide a more comprehensive mereological and logical apparatus in which to situate our account of vague objects. For example, there are salient issues arising from debates concerning mereological universalism and mereological moderatism; we have left these pertinent questions regarding the nature of composition to the side. As a result, we have not sought to resolve van Inwagen’s (1990) Special Composition Question – certainly a live issue when considering vagueness with respect to spatial boundaries. Even further still, there are pertinent issues with respect to vague objects and their temporal extensions, or of metaphysical vagueness and the (perhaps) presently unsettled and open future. There are also live questions with respect to how vague objects arise in virtue of their constituents – perhaps the popular notion of “grounding”, or some other dependence relation, could play some role in offering explanatory utility for those engaged with the project of defending metaphysical vagueness. We have also not engaged with the rich literature on non-classical logic or on non-metaphysical views on vagueness. Perhaps some combination of views and logics will enable a fully systematic account that can aptly explain the puzzling phenomenon of vagueness in its myriad guises.

Overall, we hope our discussion has provided reason for optimism. The field of metaphysics is well and truly alive, engaging with serious issues and offering more refined accounts of phenomena that call for explanation. We are optimistic that the serious consideration of vague objects as existent things may help in contributing to these lines of inquiry.

Conclusion

In summary, we see good reason to endorse the existence of vague objects and to see them as being vindicated from their historically maligned reputation. Let us briefly retrace why this is so. In Chapter 1 we decoupled vague objects from contradiction, defending their existence against fervent criticism and the charge that endorsing their existence required the use of revisionary logic or the acceptance of contradictory notions of identity. In Chapter 2, we demonstrated how vague objects can be used to develop a cogent solution to the Problem of the Many – a solution respecting classical logic, standard mereology, and our counting intuitions with respect to ordinary objects. In Chapter 3 we situated vague objects within the contemporary literature. We considered various accounts of metaphysical vagueness and demonstrated the utility of our novel vague-objects view as a cogent instance of metaphysical vagueness.

Ultimately, our primary aim has been to vindicate the idea that vague objects are genuinely existent things – to argue against their historically maligned reputation. Throughout the thesis we have hoped to show that their vindication is, at the least, *plausible*. Our more ambitious hope for the future is that such vindication will lead to fruitful applications in various areas of metaphysical inquiry. We hope these prospects are theoretically *promising*, and that our ability to present a cogent response to the Problem of the Many is a means of demonstrating this promise. Having laid the foundations for such a line of research, we hope it can bear fruit in its future application to a variety of problems in metaphysics.

Bibliography

Abasnezhad, A 2016, ‘How (Not) to Argue Against Vague Object’, *Metaphysica: International Journal for Ontology and Metaphysics*, vol. 17, no. 2, pp. 195-205.

Abasnezhad, A and Hosseini, D 2014, ‘Vagueness in the World: A Supervaluationist Approach’, in K Akiba and Abasnezhad A (eds), *Vague Objects and Vague Identity: New Essays on Ontic Vagueness*, Springer, Dordrecht, pp. 239-255.

Abasnezhad, A and Jenkins, C 2018, ‘Metaphysical Vagueness Without Vague Objects’, *Thought: A Journal of Philosophy*, vol. 7, no. 4, pp. 278-283.

Akiba, K 2000, ‘Vagueness as a Modality,’ *Philosophical Quarterly*, vol. 50, no. 200, pp. 359-70.

Akiba, K 2004, ‘Vagueness in the World’, *Noûs*, vol. 38, no. 3, pp. 407-429.

Akiba, K 2014, ‘A Defense of Indeterminate Distinctness’, *Synthese*, vol. 191, no. 15, pp. 3557-3573.

Akiba, K 2015a, ‘Referential Indeterminacy with an Ontic Source? – A Criticism of Williams’s Defense of Vague Objects’, *Metaphysica: International Journal for Ontology and Metaphysics*, vol. 16, no. 2, pp. 167-178.

Akiba, K 2015b, ‘How Barnes and Williams have Failed to Present an Intelligible Ontic Theory of Vagueness’, *Analysis*, vol. 75, no. 4, pp. 565-573.

Barcan, RC 1947, ‘The Identity of Individuals in a Strict Functional Calculus of Second Order’, *The Journal of Symbolic Logic*, vol. 12, no. 1, pp. 12-15.

Barnes, E 2009, ‘Indeterminacy, Identity, and Counterparts’, *Synthese*, vol. 168, no. 1, pp. 81-96.

Barnes, E 2010a, ‘Arguments Against Metaphysical Indeterminacy and Vagueness’, *Philosophy Compass*, vol. 5, no. 11, pp. 953-964.

Barnes, E 2010b, ‘Ontic Vagueness: A Guide for the Perplexed’, *Noûs*, vol. 44, no. 4, pp. 601-627.

Barnes, E 2013, ‘Metaphysically Indeterminate Existence’, *Philosophical Studies*, vol. 166, no. 3, pp. 495-510.

Barnes, E 2014, ‘Fundamental Indeterminacy’, *Analytic Philosophy*, vol. 55, no. 4, pp. 339-362.

Barnes, E and Cameron, R 2016, ‘Are There Indeterminate States of Affairs? No’, in E Barnes (ed), *Current Controversies in Metaphysics*, Routledge, New York, pp. 120-132.

Barnes, E and Williams, JRG 2009, ‘Vague Parts and Vague Identity’, *Pacific Philosophical Quarterly*, vol. 90, no. 2, pp. 176-187.

Barnes, E and Williams, JRG 2011a, ‘A Theory of Metaphysical Indeterminacy’, in K Bennett and DW Zimmerman (eds), *Oxford Studies in Metaphysics Volume 6*, Oxford University Press, Oxford, pp. 103-148.

Barnes, E and Williams, JRG 2011b, ‘Response to Eklund’, in K Bennett and DW Zimmerman (eds), *Oxford Studies in Metaphysics: Volume 6*, Oxford University Press, Oxford, pp. 173-182.

Barnes, E and Williams, JRG 2009, ‘Vague Parts and Vague Identity’, *Pacific Philosophical Quarterly*, vol. 90, no. 2, pp. 176-187.

Bokulich, A 2014, ‘Metaphysical Indeterminacy, Properties, and Quantum Theory’, *Res Philosophica*, vol. 91, no. 3, pp. 449-475.

Broome, J 1984, ‘Indefiniteness in Identity’, *Analysis*, vol. 44, no. 1, pp. 6-12.

Burgess, JA 1989, ‘Vague Identity: Evans Misrepresented’, *Analysis*, vol. 49, no. 3, pp. 112-119.

Curtis, BL and Noonan, HW 2014, ‘Castles Built on Clouds: Vague Identity and Vague Objects’, in K Akiba and Abasnezhad A (eds), *Vague Objects and Vague Identity: New Essays on Ontic Vagueness*, Springer, Dordrecht, pp. 305-326.

Dummett, M 1975, ‘Wang’s Paradox’, *Synthese*, vol. 30, no. 3-4, pp. 301-324.

Eklund, M 2008, ‘Deconstructing Ontological Vagueness’, *Canadian Journal of Philosophy*, vol. 38, no. 1, pp. 117-140.

Eklund, M 2011, ‘Being Metaphysically Unsettled: Barnes and Williams on Indeterminacy and Vagueness’, in K Bennett and DW Zimmerman (eds), *Oxford Studies in Metaphysics Volume 6*, Oxford University Press, Oxford, pp. 149-172.

Eklund, M 2013, ‘Metaphysical Vagueness and Metaphysical Indeterminacy’, *Metaphysica: International Journal for Ontology and Metaphysics*, vol. 14, no. 2, pp. 165-179.

Evans, G 1978, ‘Can There Be Vague Objects?’, *Analysis*, vol. 38, no, 4, p. 208.

Fine, K 1975, ‘Vagueness, Truth and Logic’, *Synthese*, vol. 30, no. 3-4, pp. 265–300.

Garrett, BJ 1988, ‘Vagueness and Identity’, *Analysis*, vo. 48, no. 3, pp. 130-134.

Garrett, BJ 2014, ‘Some Comments on Evans’ Proof’, in K Akiba and Abasnezhad A (eds), *Vague Objects and Vague Identity: New Essays on Ontic Vagueness*, Springer, Dordrecht, pp. 275-282.

Geach, P 1980, *Reference and Generality*, 3rd edn, Cornell University Press, Ithaca, New York.

Hawley, K 2002, ‘Vagueness and Existence’, *Proceedings of the Aristotelian Society*, vol. 102, no. 1, pp. 125-140.

Heck, RG 1998, ‘That There Might Be Vague Objects (So Far As Concerns Logic)’, *The Monist*, vol. 81, no. 2, pp. 274-296.

Hyde, D 2018, ‘Sorites Paradox’, in Edward N. Zalta (ed), *The Stanford Encyclopedia of Philosophy*, Summer 2018 edn., accessed November 2019, https://plato.stanford.edu/entries/sorites-paradox/

Johnston, M 1992, ‘Constitution is Not Identity’, *Mind*, vol. 101, no. 401, pp. 89-105.

Jones, NK 2010, *Too Many Cats: The Problem of the Many and the Metaphysics of Vagueness*, PhD thesis, Birkbeck, University of London.

Jones, NK 2011, ‘Williams on Supervaluationism and Logical Revisionism’, *The Journal of Philosophy*, vol. 108, no. 11, pp. 633-641.

Jones, NK 2015, ‘Multiple Constitution’, in K Bennett and DW Zimmerman (eds), *Oxford Studies in Metaphysics Volume 9*, Oxford University Press, Oxford, pp. 217-261.

Keefe, R 2000, *Theories of Vagueness*, Cambridge University Press, Cambridge.

Keil, G 2013, ‘Introduction: Vagueness and Ontology’, *Metaphysica: International Journal for Ontology and Metaphysics*, vol. 14, no. 2, pp. 149-164.

Kripke, S 1971, ‘Identity and Necessity’, in RJ Stainton (ed), *Perspectives in the Philosophy of Language: A Concise Anthology*, Broadview Press, Peterborough, pp. 93-122.

La Porte, J 2016, ‘Rigid Designators’, in Edward N. Zalta (ed), *The Stanford Encyclopedia of Philosophy*, Spring 2018 edn., accessed November 2019, <https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=rigid-designators>

Lewis, D 1983, ‘New Work for a Theory of Universals’, *Australasian Journal of Philosophy*, vol. 61, no. 4, pp. 343-377.

Lewis, D 1988, ‘Vague Identity: Evans Misunderstood’, *Analysis*, vol. 48, no. 3, pp. 128-130.

Lewis, D 1993, ‘Many, but Almost One’ in J Bacon, K Campbell, and L Reinhardt (eds), *Ontology, Causality, and Mind: Essays on the Philosophy of D.M. Armstrong*, Cambridge University Press, Cambridge, pp. 23-42.

López de Sa, D 2014, ‘Lewis vs Lewis on the Problem of the Many’, *Synthese*, vol. 191, no. 6, pp. 1105-1117.

Loux, MK and Crisp, TM 2017, *Metaphysics: A Contemporary Introduction*, Routledge, New York.

Lowe, EJ 1982, ‘On Being a Cat’, *Analysis*, vol. 42, no. 3, pp. 174-177.

Lowe, EJ 1995, ‘The Problem of the Many and the Vagueness of Constitution’, *Analysis*, vol. 55, no. 3, pp. 179-182.

Markosian, N 1998, ‘Brutal composition’, *Philosophical Studies*, vol. 92, no. 3, pp. 211-249.

McGee, V and McLaughlin, BP 2000, ‘The Lessons of the Many’, *Philosophical Topics*, vol. 28, no. 1, pp. 129-151.

McKinnon, N 2002, ‘Supervaluations and the Problem of the Many’, *The Philosophical Quarterly*, vol. 52, no. 208, pp. 320-339.

Merricks, T 2001, *Objects and Persons*, Oxford University Press, Oxford.

Merricks, T 2017, ‘Locating Vagueness’, The *Journal of Philosophy*, vol. 114, no. 5, pp. 221-250.

Montgomery, H and Routley, R 1966, ‘Contingency and Non-Contingency Bases for Normal Modal Logics’, *Logique et analyse*, vol. 9, no. 35-36, pp. 318-328.

Morreau, M 2002, ‘What Vague Objects are Like’, *The Journal of Philosophy*, vol. 99, no. 7, pp. 333-361.

Noonan, HW 1982, ‘Vague Objects”, *Analysis*, vol. 42, no. 1, pp. 3-6.

Noonan, HW 1990, ‘Vague Identity Yet Again’, *Analysis*, vol. 50, no. 3, pp. 157-162.

Paganini, E 2016, ‘Vague Objects within Classical Logic and Standard Mereology, and without Indeterminate Identity’, *Journal of Philosophical Logic*, vol. 46, no. 4, pp. 457-465.

Parsons, T and Woodruff, P 1995, ‘Worldly Indeterminacy of Identity’, *Proceedings of the Aristotelian Society*, vol. 95, pp. 171-191.

Pelletier, FJ 1989, ‘Another Argument Against Vague Objects’, *The Journal of Philosophy*, vol. 86, no. 9, pp. 481-492.

Priest, G 1998, ‘Fuzzy Identity and Local Validity’, *The Monist*, vol. 81, no. 2, pp. 331-342.

Rosen, G and Smith, NJJ 2004, ‘Worldly Indeterminacy: A Rough Guide’, *Australasian Journal of Philosophy*, vol. 82, no. 1, pp. 185-198.

Russell, B 1923, ‘Vagueness’, *The Australasian Journal of Psychology and Philosophy*, vol. 1, no. 2, pp. 84-92.

Salmon, N 1981, *Reference and Essence*, Princeton University Press, New Jersey.

Sattig, T 2010, ‘Many as One’, in DW Zimmerman (ed), *Oxford Studies in Metaphysics Volume 5*, Oxford University Press, Oxford, pp. 145-178.

Sattig, T 2013, ‘Vague Objects and the Problem of the Many’, *Metaphysica: International Journal for Ontology and Metaphysics*, vol. 14, no. 2, pp. 211-223.

Sattig, T 2014, ‘Mereological Indeterminacy: Metaphysical But Not Fundamental’, in K Akiba and Abasnezhad A (eds), *Vague Objects and Vague Identity: New Essays on Ontic Vagueness*, Springer, Dordrecht, pp. 25-42.

Sattig, T 2015, *The Double Lives of Objects: An Essay in the Metaphysics of the Ordinary World*, Oxford University Press, Oxford.

Sider, T 2001, “Maximality and Intrinsic Properties”. *Philosophy and Phenomenological Research*, vol. 63, no. 2, pp. 357-364.

Sider, T 2003a, ‘Maximality and Microphysical Supervenience’, *Philosophy and Phenomenological Research*, vol. 66, no. 1, pp. 139-149.

Sider, T 2003b, ‘Against Vague Existence’, *Philosophical Studies*, vol. 114, no. 1, pp. 135-146.

Sider, T 2009, ‘Against Vague and Unnatural Existence: Reply to Liebesman and Eklund’, *Noûs*, vol. 43, no. 3, pp. 557-567.

Simon, J 2014, ‘Indeterminate Comprehension’, *Thought: A Journal of Philosophy*, vol. 3, no. 1, pp. 39-48.

Smith, NJJ 2005, ‘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*, vol. 102, no. 8, pp. 381-421.

Smith, NJJ 2008a, ‘Why Sense Cannot Be Made of Vague Identity’, *Noûs*, vol. 42, no, 1, pp. 1-16.

Smith, NJJ 2008b, *Vagueness and Degrees of Truth*, Oxford University Press, Oxford.

Thomason, RH 1982, ‘Identity and Vagueness’, *Philosophical Studies*, vol. 42, no. 3, pp. 329-332.

Tye, M 2000, ‘Vagueness and Reality’, *Philosophical Topics*, vol. 28, no. 1, pp. 195-209.

Unger, P 1980, ‘The Problem of the Many’, *Midwest Studies in Philosophy*, vol. 5, no. 1, pp. 411–67.

Van Inwagen, P 1990, *Material Beings*, Cornell University Press, Ithaca and London.

Varzi, A 2016, ‘Mereology’ in Edward N. Zalta (ed), *The Stanford Encyclopedia of Philosophy*, Spring 2019 edn., accessed November 2019, <https://plato.stanford.edu/archives/spr2019/entries/mereology/>

Weatherson, B 2003, ‘Many Many Problems’, *Philosophical Quarterly*, vol. 52, no. 213, pp. 481-501.

Weatherson, B 2014, ‘The Problem of the Many’ in Edward N. Zalta (ed), *The Stanford Encyclopedia of Philosophy*, Winter 2016 edn., accessed November 2019, <http://plato.stanford.edu/archives/sum2009/entries/ problem-of-many/>

Wiggins, D 1965, ‘Identity-Statements’, in RJ Butler (ed.) *Analytical Philosophy*, 2nd edn, Blackwell, Oxford, pp. 40-71.

Williams, JRG 2008a, ‘Supervaluationism and Logical Revisionism’, *The Journal of Philosophy*, vol. 105, no. 4, pp. 192-212.

Williams, JRG 2008b, ‘Multiple Actualities and Ontically Vague Identity’, *The Philosophical Quarterly*, vol. 58, no. 230, pp. 134-154.

Williams, JRG 2008c, ‘Ontic Vagueness and Metaphysical Indeterminacy’, *Philosophy Compass*, vol. 3, no. 4, pp. 763-788.

Williamson, T 1994, *Vagueness*. Routledge, London, 1994.

Williamson, T 1999, ‘On the Structure of Higher-Order Vagueness’, *Mind*, vol. 108, no. 429, pp. 127-143.

Williamson, T 2003, ‘Vagueness in Reality’ in MK Loux and Zimmerman DW (eds), *The Oxford Handbook of Metaphysics*, Oxford University Press, Oxford, pp. 690-716.

Wilson, J 2009, ‘Determination, Realization, and Mental Causation’, *Philosophical Studies*, vol. 145, no. 1, pp. 149-169.

Wilson, J 2012, ‘Fundamental Determinables’, *Philosophers’ Imprint*, vol. 12, no. 4, pp. 1-17.

Wilson, J 2013, ‘A Determinable-based Account of Metaphysical Indeterminacy’, *Inquiry*, vol. 56, no. 4, pp. 359-385.

Wilson, J 2016, ‘Are There Indeterminate States of Affairs? Yes’ in E Barnes (ed), *Current Controversies in Metaphysics*, Routledge, New York, pp. 105-119.

Wilson, J and Calosi, C 2019, ‘Quantum Metaphysical Indeterminacy’, *Philosophical Studies*, vol, 176, no. 10, pp. 2599-2627.

Wolff, J 2015, ‘Spin as a Determinable’, *Topoi*, vol. 34, no. 2, pp. 379-386.

Woods, ET forthcoming, ‘Many, But One’*, Synthese*.