VAGUENESS AND THE LOGIC OF THE WORLD

by

Zack Garrett

A DISSERTATION

Presented to the Faculty of

The Graduate College at the University of Nebraska

In Partial Fulfillment of Requirements

For the Degree of Doctor of Philosophy

Major: Philosophy

Under the Supervision of Professor Reina Hayaki

Lincoln, Nebraska

May, 2020

VAGUENESS AND THE LOGIC OF THE WORLD

Zack Garrett, Ph.D.

University of Nebraska, 2020

Advisor: Reina Hayaki

In this dissertation, I argue that vagueness is a metaphysical phenomenon—that properties and objects can be vague—and propose a trivalent theory of vagueness meant to account for the vagueness in the world. In the first half, I argue against the theories that preserve classical logic. These theories include epistemicism, contextualism, and semantic nihilism. My objections to these theories are independent of considerations of the possibility that vagueness is a metaphysical phenomenon. However, I also argue that these theories are not capable of accommodating metaphysical vagueness.

As I move into my positive theory, I first argue for the possibility of metaphysical vagueness and respond to objections that charge that the world cannot be vague. One of these objections is Gareth Evans'[1] much-disputed argument that vague identities are impossible. I then describe what I call the logic of states of affairs. The logic of states of affairs has as its atomic elements states of affairs that can obtain, unobtain, or be indeterminate. Finally, I argue that the logic of states of affairs is a better choice for a theory of vagueness than other logics that could accommodate metaphysical vagueness such as supervaluationism and degree theories. Preference should be given to the logic of states of affairs because it provides a better explanation of higher-order vagueness and does a better job of matching our ordinary understandings of logical operators than supervaluationism does and because it provides a more general account of indeterminacy than the account given by degree theories.

[1]Evans, 1978.

Table of Contents

Introduction **i**

0.1 What is Vagueness? i

0.2 Previews of the Theories iii

0.2.1 Epistemicism iii

0.2.2 Contextualism iv

0.2.3 Supervaluationism v

0.2.4 Semantic Nihilism vi

0.2.5 Trivalent Theories vi

0.2.6 Degree Theories vii

0.3 Goal viii

1 The Abductive Argument for Epistemicism **1**

1.1 Introduction 1

1.2 The Arguments for Epistemicism 2

1.2.1 Sorensen's Argument 2

1.2.2 Williamson's Argument 3

1.2.3 The Epistemic Thesis 8

1.3 The Structure of Williamson's Argument 9

1.4 How Virtuous is Epistemicism? 10

1.4.1 Intentional Vagueness 14

1.5 Vagueness as a Metaphysical Phenomenon 16

1.6 The Weight of Simplicity 17

1.6.1 Semantic Nihilism 18

2 Context and Vagueness **20**

2.1 What is Contextualism? 21

2.1.1 Three Questions 23

2.2 Ellipses 26

2.2.1 Is There a Sorites Argument Immune to Context Shifts? 28

2.3 A Complete Theory 29

2.3.1 Individual Cases 29

2.3.2 Psychological and Epistemological Explanations of Individual Cases 29

2.3.3 The Need for an Answer to the Semantic Question 30

2.3.4 Consistency with other Theories 31

2.4 Contextualism as a Springboard 32

2.4.1 Framework 32

2.4.2 The Multi-Range Theory 33

2.4.3 Relativized Truth 35

2.4.4 Higher-Order Vagueness 37

2.4.5 Vagueness All the Way Down 38

2.4.6 Ordinary Truth 39

2.4.7 Vagueness in the World 40

3 Precisifications **41**

3.1 Supervaluationism 41

3.1.1 Objection: Disjunctions and Existential Generalizations 44

3.1.2 Objection: Permissible Precisifications 46

3.1.3 Objection: Incomplete Precisifications 47

3.1.4 Quick Summary 48

3.2 Semantic Nihilism 49

3.2.1 Is Semantic Nihilism Self-Defeating? 51

3.2.2 Braun and Sider's Nihilism 53

3.2.3 Objection: The Vagueness of 'Truth-Apt' 55

3.2.4 Objection: Metaphysical Vagueness 56

3.2.5 Objection: Borderline Approximate Truth 57
3.2.6 Simplicity 57
3.2.7 Quick Summary 58
3.3 Expressivism 58
3.3.1 Objection: Plurivaluationism Revisited 60
3.3.2 Objection: Higher-Order Vagueness 62
3.3.3 Quick Summary 63

4 The Possibility of Vagueness in the World 64

4.1 What Does Vagueness in the World Look Like? 65
4.1.1 Vagueness in Property Instantiation 66
4.1.2 Vagueness in Existence 66
4.1.3 That These are Metaphysical 67
4.2 Argument for the Possibility of Vagueness in the World 68
4.2.1 An Argument from Ordinary Objects 68
4.2.2 The Gunky Vague World 70
4.2.3 Indeterminacy at the Bottom 72
4.3 The Argument against Vague Identities 74
4.3.1 The Argument 75
4.3.2 Do We Need to Respond to Evans? 76
4.4 Problems with the Evans Argument 78
4.4.1 Indeterminate Existence 78
4.4.2 Some Remarks on Leibniz's Law 79
4.4.3 The Failure of Lambda Abstraction 80
4.4.4 Higher-Order Objection 82
4.5 Other Responses 83
4.5.1 Bruce Johnsen 83
4.5.2 E.J. Lowe 85
4.5.3 Peter van Inwagen 86
4.6 Concluding Remarks 89

5 The Logic of States of Affairs **91**

5.1 States of Affairs and Indeterminacy . . . 92

5.1.1 Obtaining-Values . . . 93

5.1.2 Consequence . . . 94

5.2 Sentential Logic . . . 95

5.2.1 The Atomic Elements . . . 95

5.2.2 On Choosing Definitions . . . 96

5.2.3 Tautologies and Contradictions . . . 97

5.2.4 Negation . . . 98

5.2.5 Conjunction . . . 102

5.2.6 Disjunction . . . 105

5.2.7 Quick Summary . . . 108

5.2.8 Conditionals . . . 108

5.2.9 Biconditionals . . . 111

5.2.10 The Operators in the Metalanguage . . . 112

5.3 Quantificational Logic . . . 113

5.3.1 The Domain . . . 113

5.3.2 Universal Quantification . . . 113

5.3.3 Existential Quantification . . . 116

5.3.4 Note on Indeterminate Existence . . . 118

5.4 Modal Logic . . . 118

5.4.1 Necessity . . . 118

5.4.2 Possibility . . . 119

5.4.3 Equivalences . . . 119

5.4.4 Metaphysical Possibility . . . 119

5.5 Higher-Order Vagueness . . . 120

5.5.1 The Operators . . . 121

5.5.2 Tripartite Division . . . 122

5.5.3 Indeterminately Indeterminate . . . 123

5.5.4 Is this Really Trivalent? . . . 125

5.6 Advantages . . . 125

6 Degree Theories **128**
6.1 An Argument for Degrees 129
6.2 Degree-Functional Semantics 133
6.2.1 Some Approaches to Validity 134
6.2.2 Solving the Sorites Paradox 136
6.2.3 Smith's Approach to Validity 138
A Problem 140
6.3 Non-degree-functional Approach 142
6.3.1 Edgington's Semantics 142
6.3.2 Objections to Degree-Functionality 143
6.3.3 Smith's Response 144
6.3.4 Problems with the Non-Degree-Functional Approach 145
6.4 General Problems with Degree Theories 146
6.4.1 Higher Order Vagueness 146
6.4.2 Plurivaluationism 147
6.4.3 Non-Gradable Indeterminacy 148
6.5 The Logic of States of Affairs and Degree Theories 148
6.5.1 What is a Truth Value? 149
6.5.2 States of Affairs and Degrees of Obtaining 150

A Higher-Order Vagueness **152**
A.1 Vagueness of Vague 152
A.2 Indeterminate Indeterminacy 153
A.3 The Structure of Higher-Order Vagueness 154

B Important Results from the Logic of States of Affairs **156**
B.1 Semantics Review 156
B.2 Excluded Middle 158
B.3 Non-Contradiction 158
B.4 $\phi \to \phi$ 158
B.5 De Morgan's 158
B.6 Contraposition 159

B.7 Biconditional Exchange 160
B.8 Double Negation 160
B.9 Instantiation 161
B.10 Existential and Universal Quantifiers (Incomplete) 161
B.11 *Modus Ponens* 162
B.12 General Points 163

Bibliography **165**

Introduction

0.1 What is Vagueness?

On October 31st, 2006, the Superior Court of the state of Massachussetts ruled that a burrito is not a sandwich. The legal case in question arose from a conflict between White City Shopping Center and PR Restaurants, who owned a Panera Bread restaurant located in the shopping center. The lease that PR Restaurants signed with the shopping center stated that Panera Bread would be the only sandwich restaurant in the mall. PR Restaurants claimed that the shopping center's plan to add a Qdoba restaurant was in breach of their contract. Associate Justice Jeffery A. J. Locke wrote the following in his opinion on the case.

> Given that the term "sandwiches" is not ambiguous and the Lease does not provide a definition of it, this court applies the ordinary meaning of the word. The New Webster Third International Dictionary describes a "sandwich" as "two thin pieces of bread, usually buttered, with a thin layer (as of meat, cheese, or savory mixture) spread between them." Merriam-Webster, 2002. Under this definition and as dictated by common sense, this court finds that the term "sandwich" is not commonly understood to include burritos, tacos, and quesadillas, which are typically made with a single tortilla and stuffed with a choice filling of meat, rice, and beans. As such, there is no viable legal basis for barring White City from leasing to Chair 5.[2]

There are a couple of points we need to correct in Justice Locke's opinion. First, "sandwiches," though not ambiguous, is vague. The boundaries of what counts as a sandwich are blurred. Consider the definition Justice Locke uses for "sandwich." Consider a child with a sweet tooth who wants to swap the peanut butter in their PB&J for Nutella (a hazelnut chocolate mixture). Presumably, the result is still a

[2] *White City Shopping Ctr., LP v. PR Rests., LLC* 2006.

sandwich, even though all of the interior components are sweet, but it feels less sandwich-like than an actual PB&J. Now, suppose also that the bread that is used is cut very thick; imagine a whole loaf of sandwich bread cut lengthwise with the same amount of Nutella and jelly in the middle. The resulting dish seems even less sandwich-like, perhaps to the point where we are unwilling to call it a sandwich. What counts as a sandwich is controversial. Hotdogs, hamburgers, calzones, open-faced sandwiches, Pop-tarts, etc. all have sandwich-like features, but not enough to be clear examples of sandwiches. The quesadillas served by Qdoba fall into this category of sandwich-like foods that are not clearly sandwiches nor clearly non-sandwiches.

Of course, this example of vagueness is a somewhat trivial one, but it shows just how pervasive vagueness and its blurry boundaries are in our everyday lives. Almost every word we use is vague. When someone counts as bald is vague. When someone has consented to a contract is vague. When someone has died is vague. At first, vagueness may seem like an annoyance of language, but there are deeper problems lying in wait. Consider the following argument.

P1. A person who is 240 cm tall is tall.

P2. A person who is 0.1 cm shorter than a tall person is also tall.

C. A person who is 50 cm tall is tall.

Here we have a problem. The first premise is clearly true, and the second premise seems true as well. The conclusion is clearly false. However, the argument seems valid. Something must be wrong here. Either the argument isn't really valid, our intuition that both of the premises are true is wrong, or our intuition that the conclusion is false is wrong. This paradox is known as the sorites paradox. There is nothing special about the choice of 'tall' in the argument above. Sorites arguments like the one above can be constructed with any vague word, though doing so will sometimes require some ingenuity when a predicate is not easily connected with something quantifiable.

The task for those who wish to resolve the sorites paradox is primarily to determine what is wrong with the argument above and then secondarily to explain why we find the argument so paradoxical. A satisfactory answer to the sorites paradox, however, will give us a theory of vagueness in general, not just how it appears in arguments like the one above.

It is difficult to give a definition of 'vagueness' here since there is no way of doing so without first delving into the theories of vagueness. The question of whether vagueness is an epistemic phenomenon, a linguistic one, or, as I will argue in chapter 4, a metaphysical one needs to be answered to give a definition of 'vagueness.' According to the position I argue for in chapters 4 and 5, vagueness is the phenomenon whereby it is indeterminate whether or not a property is instantiated as a result of there being no sharp boundary delineating when something instantiates that property and when it doesn't.

0.2 Previews of the Theories

This dissertation is roughly divided into two parts. The first half focuses on theories of vagueness that attempt to preserve classical logic. The second half discusses theories that reject classical logic in some way. The preview of the theories here is in the order in which they appear in the dissertation.

0.2.1 Epistemicism

Epistemicism is the position that vagueness is a merely epistemic phenomenon. Epistemicists believe that, contrary to our intuitions, the boundaries of vague words are precise. They also claim that we cannot know where these boundaries are located. The semantic part of their position would entail that there is a precise moment in the day when it is no longer 'noonish'—a last noonish nanosecond. There is a precise amount of money down to a fraction of a cent such that someone with a net worth greater than that value is rich and everyone with a net worth of that value or lower is not rich.

Epistemicism treats the sorites argument above as unsound. P2 is just false because a person who is 0.1 cm shorter than a tall person could be just below the precise requirement for being tall. However, we cannot know where that point is. Timothy Williamson[3] argues that our epistemic failures come from the fact that our beliefs about cases that are near the border of a vague word are unsafe—we would believe the same things in nearby possible worlds where our belief is false. For example, if the last noonish nanosecond is 12:15:06 and I believe (truly) that 12:15:05 is noonish, that belief would not count as knowledge because in nearby possible worlds 12:15:05 is not noonish but I still believe that it is. Roy Sorensen[4] claims that not all true sentences have truthmakers and that true sentences without truthmakers are unknowable. According to Sorensen, true sentences that apply predicates to borderline cases lack truthmakers, and so they are unknowable. '12:15:05 is noonish' may be true, but it is not made true by anything.

Epistemicism preserves classical logic by keeping the sharp divide between when a predicate applies and when it doesn't. This means that for any atomic sentence there are only two possibilities; the predicate applies to the object or it doesn't. So, for any atomic sentence, it is either true or false.

0.2.2 Contextualism

Contextualism attempts to solve the paradox by appealing to context shifts that occur as we attempt to evaluate the truth of premises like P2 above. Suppose we have a line of people arranged in such a way that the first person in the series is 240cm tall and each successive person is 0.1cm shorter than the person in front of them. As we move along the series of people, the context we are in shifts in such a way that the person we are evaluating will get the same evaluation as the ones on either side of them. P2 isn't true, but we have an intuition that it is because every time we attempt to evaluate it the context shifts in such a way that each instance seems true.

[3] Williamson, 2002b.

[4] Sorensen, 2001.

Contextualism is not committed to preserving classical logic, but it appears to give an explanation of the sorites paradox that does not require a move away from bivalence.

In recent years Diana Raffman[5] has moved away from contextualism, introducing a theory she calls the multi-range theory. Her new theory treats truth as relativized to precise contexts. The context shifts from contextualism still play a role, but they are not the whole story in the multi-range theory. Raffman's new theory preserves classical logic since the contexts she relativizes truth to are perfectly precise.

0.2.3 Supervaluationism

Supervaluationism does not preserve classical logic, but it does preserve classical tautologies when we avoid concerns about higher-order vagueness (the vagueness of vagueness). There are many ways of making our language more precise. We could treat 'rich' as meaning 'has more than $300,000 in net worth' and 'tall' as meaning 'is taller than 185cm.' Each of these ways of precisifying language are called precisifications. Among the precisifications there are acceptable ones and unacceptable ones. For example, a precisification that treated the cutoff for 'rich' as $1 would not be acceptable. To determine the truth-value of a sentence, evaluate it on every acceptable precisification. If it is true on every acceptable precisification, then the sentence is supertrue. If it is false on every acceptable precisification, then it is superfalse. Finally, if it is true on some and false on others, then it is indeterminate in truth-value.

Supervaluationism preserves classical tautologies because those tautologies would be such that they are always true on all acceptable precisifications. Consider the law of excluded middle. On every precisification, every object falls within the extension or anti-extension of every predicate. As such, any sentence of the form $\phi \vee \sim \phi$ where ϕ is an atomic sentence will be true on every precisification.

[5] Raffman, 2013.

0.2.4 Semantic Nihilism

Semantic nihilism is the view that sentences containing vague words are not truth-apt. On David Braun and Ted Sider's[6] version of the theory, this is because propositions are perfectly precise and sentences get their truth-evaluability by picking out propositions. Since sentences containing vague words fail to pick out propositions, they are not truth-evaluable.

Since the word 'vague' is vague, Braun and Sider's theory appears to entail that the sentences composing their theory are not truth-apt. However, they claim that truth on all acceptable precisifications amounts to acceptable assertibility. Since their own theory is, they argue, true on all admissible precisifications, it is acceptable for them to assert it—though it is not actually truth-apt.

Semantic nihilism preserves classical logic by restricting truth-evaluability to truth-bearers that can only have classical values.

Another view that is discussed alongside semantic nihilism is Peter Unger's[7] theory on which all ascriptions of vague predicates to objects are false. He uses sorites arguments like the one earlier in this introduction to argue that ordinary objects do not exist. At the end of chapter 3, I will also discuss John MacFarlane's[8] expressivism, which shares features with Braun and Sider's semantic nihilism and Raffman's multi-range theory.

0.2.5 Trivalent Theories

Some theories add a third truth-value in order to capture those cases that fall in the border areas of vague words. For example, suppose that American adult men who are 200cm tall or taller are tall and people who are 175cm tall or shorter are not. An American adult man ,*s*, who falls in between 175cm and 200cm is not clearly tall nor are they clearly not tall. Trivalent theories treat a truth-bearer like '*s* is tall' as indeterminate—a third truth-value in between true and false.

[6]Braun and Sider, 2007.
[7]Unger, 1979.
[8]MacFarlane, 2016.

Trivalent logics are incompatible with classical logic. The introduction of the new truth-value is often seen to invalidate the law of excluded middle. The middle is no longer excluded, so it need not be the case that 'either *S* is tall or she isn't' is true—it could be indeterminate.

The logic I will propose in chapter 5 will be, nominally, a trivalent logic. Complications emerge with higher-order vagueness that will turn most trivalent logics into logics with not just three but effectively infinitely many truth-values. I call the logic I describe in chapter 5 'the logic of states of affairs' because it treats states of affairs as truth-bearers. For example, the state of affairs of *s*'s being tall is indeterminate in value. The state of affairs of Jeff Bezos' being rich obtains, where obtaining is the equivalent of truth in the logic of states of affairs.

0.2.6 Degree Theories

Degree theories treat truth as coming in degrees, usually represented by the real numbers from 0 to 1. A degree of truth of 0 is complete falsehood and a degree of truth of 1 is complete truth. Part of the justification for such a theory is that vague adjectives are gradable. Someone can be very bald, somewhat bald, not bald at all, and everything in between. For all of the vague words that are not gradable, constructing a sorites argument will give us a way of quantifying them such that we can have a sense for what it would mean for that word to apply to degree .5.

Degree theories treat the change from tall to not tall in the sorites argument above as gradual rather than sudden. Doing so violates classical logic in ways similar to those seen with trivalent logics.

In chapter 6, I will describe two degree theories. Nicholas J.J. Smith[9] provides a degree-functional approach to the logical operators and introduces elements that are similar to Raffman's multi-range theory. Dorothy Edgington's[10] degree theory treats the logical operators as non-degree-functional.

[9] Smith, 2008.
[10] Edgington, 1992.

0.3 Goal

The goal of this dissertation is to show that (i) vagueness gives us a strong reason to reject classical logic, (ii) vagueness is a metaphysical phenomenon, and (iii) a trivalent-esque theory like the one I present in chapter 5 is the best way to accommodate the vagueness in the world. Chapters 1, 2, and 3 focus on the theories that preserve classical logic or, in the case of supervaluationism, come close to preserving classical logic. In chapter 4, I provide an argument that vagueness is a metaphysical phenomenon and defend that position against Gareth Evans' famous argument that metaphysical vagueness is impossible. In chapter 5, I provide a logic that captures the vagueness in the world. Finally, in chapter 6, I argue that the logic provided in chapter 5 is a better option for dealing with vagueness in the world than a degree theory.

Chapter 1

The Abductive Argument for Epistemicism

1.1 Introduction

Epistemicism is the view that (i) all of our vague predicates have sharp borders and (ii) we cannot know where these borderlines are. (i), the **semantic thesis**, allows us to keep the theoretical virtues of classical logic—its simplicity and usefulness—while (ii), the **epistemic thesis**, explains our intuition that vague predicates lack sharp borders as a case of ignorance. According to the epistemicist, vagueness is a purely epistemic phenomenon.

This stands in contrast to a number of theories that treat vagueness either as a semantic phenomenon or as a metaphysical one. Many of these other theories end up rejecting classical logic to varying degrees. Rejections of classical logic often result in significantly more complicated logics. Timothy Williamson claims that the theoretical virtues of classical logic, primarily its simplicity, stand as good reasons to reject such non-classical theories of vagueness.

In this chapter I will present reasons why epistemicism is not as simple as it first appears as well as an argument that it does not fit some of the evidence about language use. Finally, I will argue that the possibility of metaphysical vagueness shows that epistemicism loses out to other theories when it comes to explanatory power. Though epistemicism is a simpler theory than the one I will sketch in chapter 5, its advantage in the area of simplicity is not large enough to overrule its problems in others.

1.2 The Arguments for Epistemicism

1.2.1 Sorensen's Argument

Roy Sorensen provides an admittedly simple argument for epistemicism. Consider the following sorites argument:

> **Base Premise:** 11:55:00 is noonish.
>
> **Inductive Premise:** For all times t, if t is noonish, then t plus one millisecond is also noonish.
>
> **Conclusion:** 23:59:00 is noonish.

One millisecond will not be enough to make the difference between a time that is noonish and a time that is not noonish, and so the inductive premise appears to be true. 11:55:00 is clearly noonish and 23:59:00 is clearly not noonish. In addition, the argument is classically valid; after all, it is just a number of applications of universal instantiation and *modus ponens*. So, we have a valid argument with true premises and a false conclusion—something must have gone wrong.

Sorensen's simple argument for the semantic thesis goes as follows:

> **P1:** The base premise of the sorites argument is true.
>
> **P2:** The conclusion of the sorites argument is false.
>
> **P3:** The sorites argument is valid.
>
> **C1:** Therefore, the inductive premise of the sorites argument is false.
>
> **C2:** Therefore, there is some time t, such that t is noonish but t plus one millisecond is not noonish.[1]

If an argument is valid, then it's impossible for the conclusion to be false when the premises are true. Since the conclusion of the sorites argument is clearly false and the argument is valid, it must be the case that one of the premises is false. The base premise is clearly true, and so the false one must be the inductive premise. Sorensen gets from C1 to C2 in the above argument by accepting the negation of the inductive premise, which in classical logic is equivalent to the existential generalization in C2.

The choice of 11:55 and the choice of 'noonish' are not important for Sorensen's argument. In order for a sorites argument to be paradoxical it must appear valid,

[1] See Sorensen, 2005, p. 679 and Sorensen, 2001, p. 1

have a clearly true base premise, and have a clearly false conclusion. These are just the features needed to construct Sorensen's argument. However, there is a clear flaw in the move from C1 to C2. After all, Sorensen accepts from the outset that classical logic is applicable in this case.[2] Those who are skeptical of classical logic will not necessarily allow the step from C1 to C2. In some non-classical logics that attempt to solve the sorites paradox, the falsity of the inductive premise does not entail the existence of a sharp cutoff for the relevant predicate. The trivalent logic I will defend in chapter 5 treats the inductive premise as indeterminate, rather than false. In such a trivalent logic, we cannot conclude from the indeterminacy of the inductive premise that there is a sharp cutoff, as Sorensen does in C2. Sorensen appears to be begging the question against the proponent of a non-classical solution to the paradox.

For the remainder of this chapter, I will focus on Timothy Williamson's epistemicism. Williamson approaches vagueness in a similar but more detailed way. He first attempts to establish classical logic through an inference to the best explanation. The semantic thesis then follows from classical logic in the way Sorensen describes. The theories differ the most in their explanations of our ignorance with regard to borderline cases. However, Williamson gives a more plausible explanation than Sorensen, and explaining Sorensen's position will take us too far afield into the metaphysics of truthmakers for the purposes of this chapter.[3]

1.2.2 Williamson's Argument

Williamson advises caution when contemplating revisions of classical logic.

[2]Sorensen finds it intriguing that so many philosophers miss such an easy solution to the sorites paradox. "I blink. I marvel at how difficult it is to believe in sharp thresholds. The guiding question of *Vagueness and Contradiction* is 'Why do wise philosophers have trouble accepting the simple (decisive) argument for sharp thresholds?'." (Sorensen, 2005, p. 683)

[3]Sorensen claims that truths about borderline cases of a predicate lack truthmakers—there is nothing about the world that makes them true. Since there is nothing in the world that makes them true, there is no way for us to know that they are true. This differs from Williamson's theory in that truths about borderline cases are knowable by an omniscient being for Williamson and unknowable by any being for Sorensen.

> Humans are better at logic than at philosophy. When philosophical considerations lead someone to propose a revision of basic logic, the philosophy is more likely to be at fault than the logic.[4]

While this could stand as an argument against revising classical logic in light of the problems posed by vagueness, we need some reason for thinking that our intuition that vague predicates lack sharp borders and our arguments against epistemicism are more likely to be at fault than classical logic. Williamson argues that any theory that revises classical logic to account for vagueness is giving up the virtues of classical logic. This, he says, may even be a sufficient argument for rejecting the countervailing evidence presented by vagueness.

> If one abandons bivalence for vague utterances, one pays a high price. One can no longer apply classical truth-conditional semantics to them, and probably not even classical logic. Yet classical semantics and logic are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains. It would not be wholly unreasonable to insist on these grounds alone that bivalence *must* somehow apply to vague utterances, attributing any contrary appearances to our lack of insight. Not every anomaly falsifies a theory. That attitude might eventually cease to be tenable, if some non-classical treatment of vagueness were genuinely illuminating. No such treatment has been found.[5]

Classical logic is simple and powerful, has experienced past success, and is integrated with theories in other domains. If we move to one of the non-classical theories of vagueness, we will have to complicate the logic. This would be a lot to give up to accommodate our intuition that vague predicates lack sharp borders, and one could argue that the loss of simplicity is enough to warrant skepticism of the objections to epistemicism.

Williamson does not rely solely on an inference to the best explanation—he offers a positive argument for the principle of bivalence. However, if one is willing to take a step towards vague metalanguages, the argument returns to the abductive argument for classical logic. Supposing that u is an utterance, the principle that Williamson sets out to prove is:

[4]Williamson, 1997, p. 215.
[5]Williamson, 2002b, p. 186.

(B): If u says that P, then either u is true or u is false.

He makes use of the following two principles about truth and falsity that he thinks are fundamental to the notion of truth.

(T): If u says that P, then u is true if and only if P.

(F): If u says that P, then u is false if and only if not P.

The proof proceeds by *reductio ad absurdum*.

(0) u says that P. (Assp. for *reductio*)

(1) Not: either u is true or u is false. (Assp. for *reductio*)

(2a) u is true if and only if P. (MP (0) and (T))

(2b) u is false if and only if not P. (MP (0) and (F))

(3) Not: either P or not P. (Substitution of equivalents (1), (2a), and (2b))

(4) Not P and not not P. (DeM (3))[6]

From the assumption that there is a counterexample to (B), Williamson derives a contradiction. The counterexample to (B) is an utterance u such that u says P and it is not the case that either u is true or u is false. Through some simple logic this leads to the conclusion that not P and not not P.

The problem with this argument lies in the meaning of 'not.' Williamson uses it as classical negation, allowing him to use De Morgan's Law in the final step. The proponent of non-classical logic would argue that assuming (1) does not entail the falsity of "either u is true or u is false." When rejecting the law of excluded middle as it appears in (1) and (3), the proponent of non-classical logic will argue that just because excluded middle is not true does not mean that it is false. The counterexample to bivalence, according to the proponent of non-classical logic, need not be a sentence, P, such that it is false that $P \vee \sim P$. All that is needed for a counterexample is a sentence, Q, such that it is indeterminate, or some other non-classical value, that $Q \vee \sim Q$. To demonstrate further how this argument fails, consider the difference between strong and weak negation. The truth table for the former is on the left,

[6] Williamson, 2002b, pp. 188-189.

and the truth table for the latter is on the right. '$\neg$' will be used to represent weak negation.[7]

TABLE 1.1: Truth Tables for Negation

p	$\sim$	p
T	F	T
#	**#**	#
F	T	F

p	$\neg$	p
T	F	T
#	**T**	#
F	T	F

Whereas classical negation says something like "it is false that," weak negation says "it is not true that." For the former, if the input is indeterminate, then the output is indeterminate. For the latter, if the input is indeterminate, then the output is true. Consider the many-valued theorist's rejection of the law of excluded middle. The claim is not that sentences of the form $P \vee \sim P$ are sometimes false. Rather, the claim is that such sentences are sometimes indeterminate or some other non-classical value other than true. With this understanding of negation, the final use of De Morgan's Law will fail. The proponent of non-classical logic will read the 'not' in (1) as weak negation. So, after the substitution, (3) reads $\neg(P \vee \sim P)$. This is not equivalent to, nor does it entail, (4).

Williamson considers this objection. In his response he uses "Ne" to represent weak negation.

> Appeals to weak negation neglect higher-order vagueness. If first-order vagueness sometimes makes it incorrect to assert that an utterance is true or false, then second-order vagueness sometimes makes it incorrect to assert that an utterance is true or weakly false. If u is such an utterance, then ne: either u is true or u is weakly false. But this assumption regenerates the argument from (1) to (4) with 'ne' in place of 'not.' The conclusion 'Ne P and ne ne P', is a contradiction if weak negation is any kind of negation at all.[8]

Williamson claims that the only way for the proponent of non-classical logic to get out of the new contradiction is to propose yet another form of negation, weakly weak negation. This, he claims, will lead to a regress that forces the proponent of

[7] I use '#' here to stand for any non-classical truth values. It could be used to represent any indeterminate truth value, a value of 'neither true nor false,' or a value less than 1 and greater than 0.

[8] Williamson, 2002b, p. 193.

non-classical logic to accept an infinite number of primitive forms of negation. As such, the complexity of the logic increases.

To get a better look at this argument let's use a trivalent logic. $\neg Fo$ is true iff Fo is either indeterminate or false. However, when something is indeterminate or false is vague. This means that it is vague what counts as true and what counts as weakly false. For more an introduction to higher-order vagueness, see appendix A. If we assume that there is something that is neither true nor weakly false, namely a borderline case between the two, then we appear to reach a similar contradiction to the one that Williamson did in his original argument.

(0) *u* says that *P*. (Assp. for reductio)

(1) Ne: either *u* is true or *u* is weakly false. (Assp. for reductio)

(2a) *u* is true if and only if *P*. (MP (0) and (T))

(2b) *u* is weakly false if and only if ne *P*. (MP (0) and (F))

(3) Ne: either *P* or ne *P*. (Substitution of equivalents (1), (2a), and (2b))

(4) Ne *P* and ne ne *P*. (DeM (3))

The proponent of a trivalent logic is forced to introduce another form of negation, what Williamson called "weakly weak negation." *Fo* would be weakly weakly false if it is indeterminately indeterminate, determinately indeterminate, or false. Again, we can construct the same argument, because the line between indeterminately indeterminate cases and true cases is vague. The border between the weakly weakly false and the true will be blurred, and so we can come up with another form of negation, "weakly weakly weak negation." This will continue *ad infinitum*. Of course, this makes trivalent approaches to vagueness look significantly more complex than classical logic, but since it avoids the *reductio*, Williamson must lean back on his abductive argument to respond.

Williamson's argument for bivalence relies on the assumption that the negation in (1) is classical negation. However, the proponent of a non-classical logic can reject the claim that it is classical negation, accepting vague metalanguages. One can then argue that accepting vague metalanguages is not problematic, as Rosanna Keefe does.[9] However, accepting the infinite forms of negation and the infinite vague metalanguages is an intuitive cost that may hurt non-classical theories in calculations for determining the most virtuous theory.

Non-classical logics throw a curveball when it comes to Williamson's argument. Williamson's two principles of truth and falsity, (T) and (F), are not true according to various non-classical logics. One consequence of introducing indeterminacy into a logic is that some instances of the T-Schema are not true. Suppose that u says P and that it is indeterminate whether or not P.

$$\text{If } u \text{ says that } P\text{, then } u \text{ is true if and only if } P \tag{1.1}$$

(1.1) is indeterminate in such a situation. 'u is true' is false because P is indeterminate. As such, the biconditional in the consequent of (1.1) is either indeterminate or false—it certainly isn't true. Given that the antecedent of (1.1) is true, the whole sentence is either indeterminate or false, depending on the truth-value of the consequent. If (T) is not true, then it cannot be used to establish (2a) in Williamson's argument. It should be granted that there is a slight intuitive cost in invalidating the T-Schema, but it would beg the question against non-classical logics to take the T-Schema as a premise in one's argument against the non-classical logics.

1.2.3 The Epistemic Thesis

We will not spend too much time on the epistemic thesis. Sorensen and Williamson disagree about the cause of our epistemic failures regarding borderline cases. Sorensen argues that true ascriptions of predicates to the borderline cases of those predicates do not have truthmakers. True propositions that do not have truthmakers are unknowable since there is nothing that makes them true for us to use as a justification for believing that they're true.

Williamson has a less controversial argument for the epistemic thesis. His argument relies on a principle of safety. In order for us to know something, it must be the case that in close possible worlds where our belief is false we don't believe it. For

[9] Keefe, 2000, pp. 202-208.

example, if I believe that there are at least 90,000 people in a packed stadium on the basis of what I see, then my belief is not safe when there are, in fact, 90,000 people in the stadium. In a close possible world, there are exactly 89,999 people and I still believe there are at least 90,000 since my eyes cannot tell the difference between the two. Williamson claims that we need a margin for error such that our beliefs are safe. When we apply a predicate to something close to the cutoff for that predicate, our beliefs are not safe, and so they do not count as knowledge.[10]

1.3 The Structure of Williamson's Argument

Williamson's abductive argument first appears to be an argument from the simplicity of classical logic to its preferability over non-classical logics. Once the epistemicist has established classical logic, then she can deal with the sorites paradox easily with Sorensen's simple argument. The epistemic thesis of epistemicism is an explanation of our reactions to the sharp borders posited by the semantic thesis. There is no doubt that in some sense classical logic is simpler than its non-classical opponents. However, abductive reasoning is not based solely on simplicity and the way the simplicity of a theory is determined is complex. We need to consider more closely how this abductive argument for classical logic is supposed to work.

To begin, the consequences of our theories must be included in the calculations for the theoretical virtues of the theory. For example, suppose I have two competing theories for how a book fell from my bookshelf. The first assumes that two squirrels entered through the window and ran along the shelf, knocking the book down. The second assumes that a mischievous spirit knocked it down to play a prank on me. It would be too hasty to count just the number of entities that play a role in the explanation. The explanation that invokes a spirit entails the existence of supernatural creatures. It should be clear which theory is the more virtuous explanation. The consequences of an explanation, relevant or otherwise, play a role in how theoretically virtuous the theory is.

[10] For more on the topic, see Williamson, 2002a

The semantic thesis of epistemicism is a consequence of a straightforward[11] classical logic. This result is shown by Sorensen's simple argument. The epistemic thesis follows from facts about our epistemic limitations and the complexity of the function that determines such precise boundaries. When we evaluate the theoretical virtues of classical logic, we cannot stop at its simplicity; we need to include the counterintuitiveness of every word's having a perfectly precise meaning—one determined, as Williamson claims, by a function that is unknowable to humans. In the next section, I will give a more thorough accounting of the virtues of classical logic and epistemicism.

1.4 How Virtuous is Epistemicism?

Williamson has conveniently given us a list of the virtues of classical logic—we need only add to his list the virtues and vices of epistemicism. However, there are some problems with Williamson's list. On his list he includes simplicity, strength, integration with other theories, and past success. Let's consider each of these in turn.

First, Williamson claims that classical logic has experienced past success. What past successes can we attribute to classical logic? Presumably, past successes of classical logic must be instances of correct argumentation using principles of classical logic. The actual principles themselves cannot count as the past successes, since they are the very thing in contention now. If the instances of correct argumentation count as past successes to classical logic, then they count equally well for most of classical logic's competitors. Though some trivalent theories remove some inferences and some logical truths, one can still be confident in one's use of these forms when they do not contain any vagueness.[12]

Second, Williamson claims that classical logic is integrated with our other theories. It is, however, not clear how it is integrated with these theories. Quine makes a similar claim in a paper where he intends to neither defend nor attack bivalence, but only to discuss the costs of accepting bivalence and of rejecting it.

[11] For my purposes here, supervaluationism is non-classical even though it preserves classical logical truths.

[12] For other examples see Williamson, 2002b, pp. 151-152

> Bivalence is a basic trait of our classical theories of nature. It has us positing a true-false dichotomy across all the statements that we can express in our theoretical vocabulary, irrespective of our knowing how to decide them.[13]

It is not clear that bivalence or other general principles of classical logic play this role in our other theories. For example, the belief that the law of excluded middle is valid is important for many of our mathematical beliefs. However, these beliefs do not require that excluded middle is valid *simpliciter*; they only require that it is truth-conducive when the atomic propositions in the inference can only receive classical truth-values. Presumably, the kinds of propositions that play a role in the proofs of mathematics do not involve vagueness, and will therefore have classical values. If it turns out that excluded middle is not always true when we are talking about bald people, it will not require much of a revision in mathematics. Andrew Wiles' proof of Fermat's Last Theorem will be safe from revision—as will pretty much any other results in mathematics. This holds true of pretty much every part of classical logic. Most accepted theories in other domains do not require that classical logic hold as is, only that certain parts of it hold—parts that are unconcerned with vagueness. So, rejecting bivalence does not require a significant revision of our system of beliefs.

Third, classical logic is stronger than trivalent logics. That is to say, it can prove more things because it has more valid argument forms and more logical truths. It is not clear that strength should be so easily added to the list of virtues of a logic. A logic that is too strong is just as bad as a logic that is too weak. We want a logic that gets close to the right level of strength, but we won't know how strong of a logic we want until we know what kinds of argument forms and sentence forms should count as valid. That is, we won't know until we have already made decisions about vagueness.

Finally, classical logic is simple. Simplicity is the only virtue that classical logic uncontroversially seems to hold over its competitors. Trivalent logics often turn into infinite-valued logics when we introduce higher-order vagueness. They contradict

[13]Quine, 1981, p. 94.

some of our intuitions about what kinds of sentences should be tautologies. The undermining of these intuitions should not be counted as a point against non-classical logics with regard to past success, but they do count against it in the area of simplicity.

Epistemicism has three main components: classical logic, the epistemic thesis, and a theory of meaning. The semantic thesis is a consequence of classical logic as applied to the sorites arguments. The solution of the sorites paradox using classical logic tells us that every vague word has a sharp border—there is a last noonish nanosecond, a centimeter that makes the difference between being tall and being not tall, and a cent that makes the difference between a rich person and a non-rich person. The epistemic thesis is an explanation for why we are perplexed by the sorites paradox. We cannot know where the border is for a vague predicate—we can't even imagine one existing—and so we don't think there is one.

There is a worry that epistemicism involves a certain arbitrariness. That is, if there are sharp borders for our vague predicates, then they must be set in some way akin to us arbitrarily determining a cutoff point. We do this on some occasions when we need to resolve an issue with vagueness to take some action, such as in a court case concerning a vaguely written law. However, when someone normally attempts to avoid the problem of vagueness by stipulating a cutoff, it feels that they are not genuinely attempting to address the problems with the meaning of the word. In response to this kind of objection, Williamson gives an explanation of the mechanism by which the borders for words are set.

Williamson claims that meaning supervenes on use and that our uses determine specific meanings for our words. As our use of the word "bald" changes, so does the cutoff for its application. We can come up with some toy examples that will make this sound very plausible. For example, suppose that a community of English speakers creates a new word "schmall," which applies to natural numbers. Their dispositions[14] to use or not use "schmall" to describe different numbers are represented in Table (1.2). The percentages in the table represent the percent of the

[14]We do not include just actual uses, but also the dispositional uses of members of the linguistic community. Person *A* may never evaluate the number 1 with regard to schmallness, but were they to do so, they would evaluate it as schmall.

language community that has a disposition to use or not use the word "schmall" when evaluating the number in that row.

TABLE 1.2: Schmall Simple

#	schmall	not schmall
1	100%	0%
2	100%	0%
3	100%	0%
4	100%	0%
5	0%	100%
6	0%	100%
...	...	...

In Table (1.2) we can see that the uses of "schmall" and "not schmall" mirror each other in such a way that it is clear that the border is between 4 and 5. If the border was instead between 3 and 4, then all of the uses of "schmall" for 4 would be incorrect. Presumably the people in this community have some knowledge of the use of their new word and some control over its meaning, so it is implausible that they would all use it incorrectly.

While it may seem intuitive in simple toy cases, there is a worry about how it works in more complicated cases. Williamson claims that the function by which uses determine meanings is chaotic. Minor changes in use can drastically change the output of the function, and this occurs in such a complicated manner that we have no hope of understanding it. The chaotic nature of the function plays into the epistemic thesis since our inability to learn the function gives us reason to think that we can never know the cutoffs of our vague predicates.

In (1.3) there is disagreement about whether or not 4 is schmall. Half of the members of the linguistic community would use the word "schmall" for 4 and half would not. Supposing that everyone in the community is an equally competent speaker, there seems to be no reason to draw the line between 3 and 4 as opposed to between 4 and 5 or vice versa.

TABLE 1.3: Schmall Complex

#	schmall	not schmall
1	100%	0%
2	100%	0%
3	100%	0%
4	50%	50%
5	0%	100%
6	0%	100%
...	...	...

There are even more complicated examples than the one I just described. Scenarios where the distributions are not symmetrical like the ones in

(1.3). The border area in (1.3) is very small, but we may have an example with a large border area and disagreement, not just between members of the language community, but also inconsistencies in a single speaker's usage of 'schmall.'

Rosanna Keefe responds to Williamson's claim about the supervenience of meaning on use in the following way. Suppose the cutoff for 'noonish' were 12:05:01 according to one function from use to meaning. A very slightly different function could put the border for 'noonish' at 12:05:02. There are many, sometimes infinite, equally good options for functions from use to meaning, and so the mere existence of a supervenience relation between use and meaning does not resolve the arbitrariness of the epistemicist position. In addition to supervenience and a function we need a reason for why one particular function wins out over the others.[15]

It is likely unfair to require Williamson to describe a law of nature that explains the reason why the function between use and meaning is the function that it is. After all, by the lights of epistemicists, this task is likely impossible. However, Keefe is right that epistemicists need something like a law to push the bruteness down to a lower level of explanation. It is better to have a function that determines the meanings of words than for those meanings to be brute. The law that determines the function from use to meaning needs to be included in our calculations of the simplicity of the theory. Unlike other linguistic theories that tie our uses of words to their meanings, the law that epistemicists require picks out one incredibly precise meaning for each word and none of the infinite other slightly different precise meanings.

1.4.1 Intentional Vagueness

It is ambiguous to say that uses determine meanings. Exactly what aspects of a use should be included as an input to the function? Probably the competence of the speaker matters and should add more weight to the uses of that speaker. The intentions of users should also affect the function. If I am coining a new word and I intend it to have some meaning, those intentions should have some effect on the meaning of the word. Now, what happens when our intentions in the use of a word

[15] Keefe, 2000, pp. 79-84.

specifically call for there to be no sharp borderline? For a recent example, consider the handball rule in soccer[16], which changed as of the 2019-2020 season, causing an enormous amount of discontent among fans and players alike. The previous rule was concerned primarily with whether or not the ball's being touched by a player's hand was a deliberate act on the part of the player. This was meant to rule out cases where the player's hand is hit by accident by an errant ball. The current rule hinges on whether or not the player's hand was in a position that made their body "unnaturally bigger"—that is to say, their hand was outside of their natural standing silhouette.[17] The current version is meant to be more precise than the previous version and remove some of the guess work of refereeing. Though "unnaturally bigger" is still vague, using the abundant on-field cameras it is often clear whether or not a player's hand is stretched beyond their natural silhouette. Players' intentions are much harder to judge and a camera will not help much in doing so.

Presumably a more precise rule would be better, allowing for clarity and fairness in the application of the rules, but the new rule is hated by both players and spectators, many of whom would like to go back to the vaguer rule. The reason for this is that sports are primarily about entertainment and only secondarily about completely fair competition. The vague rule allows for more entertaining play since players can more easily skirt the line of what counts as a handball. Since it is easier to skirt the line of what counts as a handball, players do not have to worry as much when they are accidentally near the line. For the most part players and spectators want the word 'handball' to be vaguer than it is defined in the current rules. For similar reasons to the ones given against the current rule, soccer fans would probably disagree with a rule that chose a cutoff point for when a player's touching of the ball with their hand was deliberate. Creating such a cutoff would harm the entertainment value of soccer.

Fans would like it to be the case that "handball" does not have a precise cutoff. If the official soccer organizations return the rule to its previous version with the intention of removing precision, then presumably the word "handball" would not

[16]"Football" for those who would prefer a more descriptive name.
[17]See IFAB, 2019, pp. 104-105

have a sharp cutoff. In the case of 'handball,' we have a linguistic authority that sets the meaning of the word. So, if the International Football Association Board (IFAB) decides that they are going with a purposefully vague meaning for 'handball,' then due to the effect that the IFAB's authority would have on the meaning of the word, their uses of 'handball' would result in a word without a sharp border. So, there would not be a function from use to a precise meaning for 'handball.' Even if we discount the authority of the official soccer organizations, similar decisions on the uses of words may occur in everyday pickup games of soccer. Perhaps the boundary of the playing area is left purposefully blurred—more so than it might usually be—in order to have a more organic playing experience where pauses in play can be left up to the discretion of the players.

This phenomenon of intentional vagueness is not rare. The stipulation of precision in the meanings of our words creates arbitrary boundaries that can make games less fun, or in more serious situations, lead to injustices. Having laws that are vague allows room for discretion on the part of a judge. Vagueness is useful, and so it plays a role in our linguistic toolbox. Epistemicism claims that even when we intend our words to be vague they have precise meanings. This does not match our linguistic activities. Posed as an error theory, epistemicism may claim that we are just wrong when we think we've intentionally left the meaning of a word vague. However, epistemicism would need to be simpler than it actually is to make a strong abductive case that we are in error here.

1.5 Vagueness as a Metaphysical Phenomenon

In chapter 4, I will argue that vagueness is not merely an epistemic or a linguistic phenomenon—it is a metaphysical one. I will give a quick preview of the argument from that chapter here.

There are a number of reasons to think that there is genuine vagueness in the world. First, we already have indeterminacy in the world in the form of quantum mechanical indeterminacy. Sometimes it is indeterminate whether or not electron *a* is identical to electron *b*. Vagueness is just a kind of a broader category of metaphysical indeterminacy. Second, it is possible that the world is vague all the way down. One attempt to avoid metaphysical vagueness is to treat putatively vague properties as illegitimate and foregoing them for more fundamental and more precise properties. However, if the more fundamental properties are also vague, then we have not avoided metaphysical vagueness. Finally, those who accept ordinary objects into their ontology should already be predisposed to accept metaphysical vagueness because the best way to cash out what a baseball is involves metaphysical vagueness.

Since vagueness is at least possibly a metaphysical phenomenon, our logic should have something to say about situations where someone indeterminately instantiates the property <being bald>[18]—or any other indeterminate property instantiation. Here is where epistemicism will fail to give us a full explanation of the phenomenon of vagueness. Epistemicism claims that vagueness is a merely epistemic phenomenon. This may be true of the actual world—though I argue that it is probably not true. However, if there are possible vague property instantiations, then classical logic will need to be amended to handle this possibility. In addition, our lack of knowledge about modal claims involving metaphysically borderline cases will not be caused by margins for error. Instead, our lack of knowledge comes from the lack of a fact of the matter in the given case.

1.6 The Weight of Simplicity

The goal of this chapter was to show that epistemicism is not as simple of a solution to the sorites paradox as it might first appear. The epistemicist cannot rely wholly on the simplicity of classical logic to justify their answer to the paradox. The simplicity of epistemicism as a whole is not enough to overcome issues with evidentiary fit and explanatory power. The other theoretical virtues Williamson claims in favor of classical logic cannot be granted to classical logic until we know what the correct theory of vagueness is. Its past successes could equally be successes of non-classical

[18]Throughout this dissertation I use $<$ and $>$ to indicate properties.

logics. Rejecting classical logic does not require much revision in other domains. Though it is more powerful than most non-classical logics, it may very well be too powerful.

The biggest point in favor of epistemicism is the simplicity of classical logic. However, classical logic applied to the sorites paradox entails that every word has a precise meaning. 'Noonish' means, for example, 12:05:01 and not 12:05:02, or any of the infinite other options. The law that governs how these meanings are set is a large theoretical burden on epistemicism, given how big of a role it must play in the theory.

Though epistemicism is simpler than many of its competitors, it has problems that block the use of an abductive argument. First, Williamson's linguistic claims do not match the way that we use language. There are times when we want our words to be imprecise. The meanings of words can be heavily influenced by those with authority in the language community. If those in authority desire the word to be imprecise, then it very likely is imprecise.

In addition, vagueness is not merely an epistemic phenomenon. As I will argue in chapter 4, it is possible to have indeterminate instantiations of properties. This possibility cannot be handled by epistemicism.

A theory may be virtuous with regard to some theoretical virtues and not with regard to others. However, for a theory to overlook its vices, it needs to be very virtuous in other respects. Epistemicism wins out on simplicity, but not by enough to make a compelling case that we are mistaken about intentional vagueness or the possibility of metaphysical vagueness.

1.6.1 Semantic Nihilism

Epistemicism is not the only theory that preserves classical logic. Semantic nihilism, the theory that sentences containing vague words are not truth-evaluable, preserves classical logic without the added complexity in the semantic theory that epistemicism needs.[19] In chapter 3, I will argue against semantic nihilism. An attempt to

[19] See Gómez-Torrente, 1997, pp. 240-243 for a discussion of nihilism and the abductive argument for epistemicism.

use abductive reasoning to establish semantic nihilism over its non-classical opponents will fail for similar reasons to the ones cited in this chapter with regard to epistemicism.

Chapter 2

Context and Vagueness

Introduction

Contextualism provides a intuitive place for proponents of classical logic to settle. It is clear that vague terms are context-sensitive. Ambiguous words are also context-sensitive and we can often remove some ambiguity by fixing a context. For example, the ambiguity in the sentence “Winston went to the bank” is often resolved by specifying the context as one where Winston is planning to skip rocks. In the case of vagueness, the evaluation of “Joanne is tall” may be unclear when a context has not been set.

However, when we specify the context as “among basketball players,” and supposing that Joanne is 170 cm tall, it is clear that the sentence is false. If we can solve the sorites paradox in a more intuitive way than Williamson’s epistemicism and preserve classical logic, that would be fantastic. However, contextualism alone is not a full theory of vagueness—it is only an explanation of some of the phenomena surrounding the sorites paradox. To become a complete theory it will need to be supplemented with one of the other theories of vagueness. Doing so will commit the contextualist to the intuitive costs of the theory to which it is attached. The first half of this chapter will deal with the concerns just described. The second half of the chapter will discuss Diana Raffman’s multi-range theory—a theory that relativizes truth to contexts. Raffman’s theory is a natural jumping-off point for the contextualist. Their solution to the sorites paradox involves shifting contexts that shift the borders of vague predicates. If we wish to preserve classical logic, then relativizing truth to the precise extensions generated by contexts would be a natural step from contextualism. However, the multi-range theory suffers from serious problems dealing with higher-order vagueness. There can be contexts such that it is indeterminate whether or not we can relativize truth to them. If preserving classical logic is the goal, then the multi-range theory is not the way to go.

2.1 What is Contextualism?

Suppose we have a line of piles of sand where each pile in the line has one more grain than the previous pile. Contextualism is the view that all vague words are context-sensitive and that the context shifts as we move down the line of piles of sand. It shifts in such a way that any two adjacent piles will either both be included or both be excluded from the extension of "heap." The reason that we fail to find a precise border between the non-heaps and the heaps is because as we search for it the context shifts.

Suppose we have a series of patches of color starting with shades that are clearly red and ending with shades that are clearly orange. The patches are labeled #1-#1000 and are arranged in such a way that any two adjacent patches are indistinguishable to normal humans. Consider the following sorites argument.

Base Premise: #1 looks red.

Inductive Premise: If #n looks red, then #(n + 1) looks red.

Conclusion: #1000 looks red.

By stipulation of the case, the base premise is true and the conclusion is false. The inductive premise seems like it should be true since any adjacent pairs will be indistinguishable. Diana Raffman asks her reader to envision themselves moving from patch to patch in the above sorites series. Each patch starts off looking red, but as we get further along, our evaluations become more and more tentative. Eventually, we decide that one of the patches looks orange. How can it be the case that each patch is indistinguishable from its adjacent patches and yet one of them is placed in the category 'orange' and the other is in the category 'red?'

> Ah, but there's the catch: they are not in different categories—or, rather, they are not in different categories when judged *pairwise*.[1]

Raffman contends that when judged pairwise each pair of patches will be such that the members both look red or both look orange. She claims that the inductive premise as it is written in the argument above is false. However, a variant on that inductive premise is true:

Inductive Premise Variant: If #n looks red, then #(n + 1) looks red, *insofar as #n and #(n + 1) are judged pairwise*.[2]

When we find the sorites argument paradoxical it is because we are mistaking the actual inductive premise for this variant. This mistake is understandable because evaluating the inductive premise requires us to judge the patches in a pairwise way. Each replacement of the variable n in the original inductive premise above creates a pairwise assessment in order to determine the truth-value of the resulting conditional.

Once more imagine moving along the series of patches of color. Eventually, you decide that one of the patches looks orange, let's say it is patch #658. How is this category shift consistent with the inductive premise variant's being true? Consider moving backwards from the point you are at, returning to patch #1. Your evaluation of #657 will have changed—you'll think that you were incorrect in your evaluating #657 as 'red.' Continuing towards #1 you'll become more and more tentative in saying that the patches look orange until you eventually decide that one of them looks red. Suddenly some of the patches you thought looked orange now look red. There are limits to how far you can stretch the extensions of red and orange, but they can stretch within certain ranges.

Contextualism claims that the context shifts as we move along a sorites series. Raffman gives a thorough look at the ways that context shifts for vague predicates.

> What is distinctive about the sorites is that it exposes the variability of the correct application of a vague predicate even *independently of any*

[1] Raffman, 1994, p. 46.
[2] Raffman, 1994, p. 47.

> *variations in "external" context* (including variations in the identity of the judging subject). Hold fast the subject and the entire external context, ..., and still the extension of 'red' will vary.[3]

Even if the external context—things like the lighting environment—and the subject[4] are held fixed, the extension of 'red' can vary. Raffman attributes this variability to three different contextual elements. First, vague words can vary with presentational context—whether we are judging cases singly, pairwise, threewise, etc. They can also vary with historical context. Patch #501 may look red when we've evaluated #1-#500 as red, but not if we were coming from #1000 down to #501. Finally, vague words can vary with what Raffman calls *categorial context*. The categorial context is one where, even when the other two contexts are held fixed, the "brute mechanical variations in the operation of our categorizing homunculus" causes us to change the way we categorize things, thereby shifting the extension of the vague word. All of these, plus the external context and other features of the judging subject, form the total context that a subject may be in when evaluating the application of a predicate. Because of these context shifts, the inductive premise variation comes out true and we are thereby fooled into believing the inductive premise comes out true.

2.1.1 Three Questions

For this subsection, I will make use of the following sorites argument about heaps, wherein I have formalized the inductive premise as a univerally quantified conditional. Let Hx mean 'x is a heap' and Rxy mean 'y has one fewer grain of sand than x.' Finally, imagine that we have a series of piles of sand starting with a pile with 1 grain and ending with a pile of 1000 grains. p_1 is the pile with 1 grain, p_2 is the pile with 2 grains, and so on.

Base Premise: Hp_{1000}

Inductive Premise: $\forall x \forall y((Hx \wedge Rxy) \rightarrow Hy)$

[3] Raffman, 1994, p. 64.

[4] It is not wholly clear what parts of the judging subject we are supposed to keep fixed. As she points out shortly after this passage, it is possible that the same subject may give different evaluations to the same patch under the same viewing and historical conditions. The cause of this difference in evaluation is something internal to the subject's categorizing homunculus—the internal component of the subject that handles categorizing. So, not every feature of the subject is held constant.

Conclusion: Hp_1

Delia Graff Fara[5] claims that not all solutions to the sorites paradox focus on the same question. Fara sets out three questions that she thinks are the targets of various theories of vagueness.

1. *The Semantic Question*: If the universally generalized inductive premise of the sorites argument is not true, then must the classical equivalent of its negation be true? That is, if $\sim\forall x \forall y((Hx \land Rxy) \rightarrow Hy)$, is it the case that $\exists x \exists y((Hx \land Rxy) \land \sim Hy)$?

2. *The Epistemological Question*: If the universally generalized inductive premise of the sorites argument is not true, then why are we unable to say which of its instances is not true?

3. *The Psychological Question*: If the universally generalized inductive premise of the sorites argument is not true, why are we so inclined to accept it?[6]

For an example of how a theory might answer these questions, Williamson's epistemicism answers the semantic question by saying that the classically equivalent negation is true—there is an x and a y such that $Hx \land Rxy \land \sim Hy$. To the epistemological question, he claims that we cannot know because when we are near the borderline, we do not have a large enough margin for error. Finally, to the psychological question, we are inclined to accept the inductive premise because we cannot imagine the transition.

Fara's main targets are the epistemological and psychological questions. As for an account of the borderline, "I am happy to wait and see what story about borderline cases, and the characteristic hedging responses they provoke, might naturally flow from such an account."[7] Her answer to the epistemological and psychological questions, which she calls a *Bare Bones Solution*, begins with a discussion of four constraints on the variation of standards for the application of a vague predicate.

[5] Fara, 2000.
[6] Fara, 2000, p. 50.
[7] Fara, 2000, p. 54.

1. *Clear Case Constraint*: There are certain things to which the predicate must apply and certain things to which it must not apply. "Small" must apply to quarks. We cannot vary the context so much that quarks are not "small."

2. *Relational Constraint*: Some vague predicates like "tall" have relational constraints such that for any x and y if x is tall and y is taller than x, then it must be the case that y is tall.

3. *Coordination Constraint*: There are constraints on the coordination of relational predicates. We cannot have a context where something counts as both rich and poor.

4. *Similarity Constraint*: If x is relevantly and saliently similar to y and y meets the standards for a vague expression, then so does x.

It is this last constraint that is of most importance for Fara's contextualist answer to the epistemological and psychological questions. The reason that we do not know where the borderline is for a vague term is because when we pick out two cases that might form the border we force their similarity into salience. Since the similarity between the cases is salient, by the similarity constraint, we cannot push ourselves into a context in which one but not the other of the cases is outside of the extension of the vague term. This has the effect of rendering each instantiation of the inductive premise true when we observe it in isolation. As for the psychological question, Fara answers that our inability to find a borderline—a result of the similarity constraint—motivates our intuition that the inductive premise is true.

An answer to the semantic question requires a discussion of the changes or lack thereof we are proposing for classical logic. Though Fara is an adherent of bivalence, her *Bare Bones Solution* is not committed to bivalence. For example, supervaluationists might take all of the constraints mentioned above as constraints on the range of admissible precisifications for a vague predicate. If a is one millimeter taller than b and this similarity is salient, then there will be no admissible precisification of "tall" such that a is tall and b is not. The applicability of the *Bare Bones Solution* to theories that do not preserve classical logic leaves open what the contextualist response to the semantic question would be.

2.2 Ellipses

Jason Stanley[8] argues that contextualists fail to solve some versions of the sorites paradox. Stanley attempts to create the force of the sorites argument while keeping the context fixed throughout. He proposes to do this by making use of verb-phrase ellipses. A verb-phrase ellipsis is a linguistic structure where a verb-phrase is left out.

Example Sentence: Mary will go see the new Marvel movie, and Denise will go see the new Marvel movie too.

Verb-Phrase Ellipsis: Mary will go see the new Marvel movie, and Denise will too.

With Ellipsis Present: Mary will go see the new Marvel movie, and Denise will ... too.

He argues that context remains fixed through verb-phrase ellipses, and so it would remain fixed throughout the argument. Below is an example of how such a sorites argument would go.

Base Premise: p_{1000} is a heap.

Inductive Premise: For all x, if p_x is a heap, then p_{x-1} is too.

Conclusion: p_1 is a heap.

The verb-phrase ellipsis appears in the consequent of the conditional in the inductive premise, otherwise this argument is the same as the sorites arguments we've seen so far.

Stanley claims that indexicals are context-invariant over verb-phrase ellipses.

> Here is a fact about indexical expressions. Indexicals have invariant interpretations in Verb Phrase ellipsis.[9]

[8] Stanley, 2003.
[9] Stanley, 2003, p. 271.

In support of this point he offers a number of intuitive examples.

1. John likes me, and Bill *does too*.
2. Hannah lives here, and Bill *does too*.
3. Hannah is supposed to be in Syracuse now, and Mary *is too*.
4. John saw Hannah's film, and Bill *did too*.
5. John read that, and Bill *did too*[10].

For each of these sentences there is no interpretation such that the indexicals have different meanings. There is no available interpretation under which Hannah and Bill live in different locations. Similarly, there is no available interpretation under which John and Bill read different things.

Returning to the universally quantified inductive premise of the sorites argument above, for each of the conditionals we get from instantiating the universal quantifier, there is no available interpretation on which the "heap" in the antecedent has a different meaning than the implicit "heap" in the consequent.

The problem with Stanley's version of the sorites paradox is that though the context does not shift within the scope of each conditional or within the universally quantified sentence, it can move as we examine each instantiation of the inductive premise on its own. The contextualist response to Stanley's sorites argument can be the same as their response to the sorites argument I provided at the beginning of this chapter. We think that the universally quantified inductive premise is true because when we examine each instantiation the context shifts in such a way that that instantiation comes out true. When we examine each of the instantiations in Stanley's sorites argument, the context does not shift in the course of that particular sentence, but the contextualist only needs the context to shift as we examine each instantiation. Unless the verb phrase ellipsis is spread across the entire sorites series, we cannot keep the context fixed.

$$\text{If } p_{1000} \text{ is a heap, then } p_n \text{ is too.} \tag{2.1}$$

[10]The words in italics represent the verb phrase ellipses in these sentences. The emphasis is mine.

Assuming p_n is clearly not a heap, (2.1) is false and does not generate a sorites paradox.

Another response to Stanley's attempt at a context-invariant sorites argument is provided by Diana Raffman. Raffman's theory of vagueness treats vague words similarly to indexicals, but with a few important differences. In particular, she notes that context is not invariant over verb phrase ellipses with regard to vague words and she gives a few examples.

Shaquille is tall, and so is the Empire State Building. (2.2)

The comparison class in the first clause is different than the comparison class in the second clause. This indicates that context can shift over verb phrase ellipses with vague words. If we mean by contexts the same thing as Raffman, where miniscule differences in our psychology can result in context shifts that affect the extension of a predicate, then even the examples Stanley gives could involve context shifts through the verb phrase ellipsis. The context will not shift dramatically, so it will not be noticeable.

2.2.1 Is There a Sorites Argument Immune to Context Shifts?

Stanley's argument represents the obvious way to attack contextualism. If we can find a sorites paradox that keeps the context fixed, then contextualism fails to answer the epistemological and psychological questions with regard to that version of the paradox. Is it possible to come up with a version of the paradox that is not susceptible to context shifts?

I believe the answer is 'no.' To generate the paradox we need small changes between any two items in our sorites series. These small changes will always be small enough to trigger the similarity constraint. By virtue of the fact that the steps of a sorites series have to be small we are guaranteed to be in a context in which both or neither of two adjacent members of the series fall within the extension of the predicate we are constructing the sorites argument about.

So contextualism seems safe from attempts to create sorites paradoxes that are immune to context shifts. However, sorites paradoxes are not the only phenomena for which we need psychological and epistemological explanations. In the next section, I will argue that contextualism fails to provide psychological and epistemological explanations for some vagueness phenomena. In addition, the success of contextualism in providing answers to Fara's specific psychological and epistemological questions should not overshadow the fact that it does not give us an answer to Fara's semantic question or similar semantic questions.

2.3 A Complete Theory

2.3.1 Individual Cases

Though contextualism provides us with a good answer to Fara's psychological and epistemological questions when we are moving along a sorites series, it does not give us an adequate explanation of our reactions when we observe individual instances of borderline cases and do not consider their similarity with nearby objects on a sorites series. If, instead of moving along a series of piles of sand, one is presented with a single pile, p, in the border area, then what should be our evaluation of 'p is a heap'? Let's assume that the historical context will not play a decisive role because it has been too long since you last evaluated a pile of sand for its heapness. The presentational context is one where we evaluate the object singly.

2.3.2 Psychological and Epistemological Explanations of Individual Cases

Individual Semantic Question: If x is a borderline case of some predicate F, then what is the truth-value of Fx?

Individual Epistemological Question: If x is a borderline case of some predicate F, then why are we unable to say whether x is F or not?

Individual Psychological Question: If x is a borderline F, then why do we struggle to evaluate whether or not x is F?

It is clear that something special is going on in the case of an individual borderline case. For example, when trying to evaluate whether or not a borderline bald person is bald, we feel like neither answer is completely correct. On the epistemological side, we clearly do not know whether the person is bald or not. We cannot appeal to the similarity constraint to explain our feelings here, since we have been presented with a single borderline case out of the blue. Whatever context we are in is not one that pushes the person into or out of the extension of 'bald.' One might respond that categorial context will adjudicate the case. However, there seems to be no reason to believe that the brute mechanical operations in our brains would always resolve borderline cases. As such, context will not supply us with an answer to the questions presented in this subsection.

What kind of explanation would work here? A theory that accepts indeterminacy, like the one I will describe in chapters 4 and 5, will hold that it is genuinely indeterminate whether the person is bald, and so we could not know that he is bald, nor could we know that he is not bald. Similarly, the psychological unease from attempting to classify borderline cases is explained by the fact that the person is in an indeterminate position between baldness and non-baldness.

2.3.3 The Need for an Answer to the Semantic Question

To give a complete theory of vagueness the contextualist needs to answer the individual semantic question. The reason for this requirement is that we cannot explain them in terms of context shifts. The fact that we have an individual case that is borderline according to the current context indicates that the phenomenon is not caused by shifting contexts.

Fara's *Bare Bones Solution* is very bare bones since it only gives an explanation of one particular phenomenon related to vagueness—the sorites paradox, and even then it only gives answers to Fara's epistemological and psychological questions. A full theory of vagueness will provide explanations of all of the phenomena of vagueness, including both sorites series and individual borderline cases.

2.3.4 Consistency with other Theories

One reason a contextualist might avoid answering the semantic question is that the contextualist answers to Fara's epistemological and psychological questions appear to be consistent with multiple competing theories of vagueness that offer answers to the semantic question.

Contextualism is pretty clearly compatible with epistemicism. When we start moving along the sorites series we are approaching a sharp border, but when we get close we change contexts such that the border in the new context is in a different location in the series.

Fara describes a supervaluationist theory that is compatible with contextualism. The normal supervaluationist response to the inductive premise of the sorites argument is to claim that it is superfalse. This is because on every precisification of "heap," there will be some a and some adjacent b such that a is a heap and b is not. So, there will be at least one instantiation of the universal in the inductive premise that is false on every precisification. This is their answer to the semantic question. The similarity constraint can give them answers to the other two questions. If a and b are adjacent cases in a sorites series, then what counts as an admissible precisification of "heap" will either include both or exclude both. As such, each individual instantiation of the universal will be supertrue. As we move along the sorites series, the context shifts such that what counts as an admissible precisification shifts as well.

One way the contextualist could give a complete theory of vagueness is to choose one of the other theories of vagueness and combine contextualism with that theory. It is clear that vague words are context-sensitive. It also seems likely that we are subject to a shifting context as we move along the sorites series. As such, contextualism can offer an interesting insight into the psychological and epistemological phenomena going on when we move along a sorites series. However, if one intends to preserve classical logic, then contextualism is not clearly a route for doing so.

2.4 Contextualism as a Springboard

As we saw in section (2.1), in 1994[11] and 1996[12], Diana Raffman argued for contextualism. She claims, however, that she is no longer a contextualist.

> I have since come to believe that vagueness has nothing essentially to do with context-sensitivity (though the two often come together). The competent application of a vague word is variable, both inter- and intra-subjectively, in roughly the way I initially had in mind, but not because of any contextual sensitivity. In fact, I now think it is of the essence of vagueness that the competent application of a vague word varies even when the operative context is fixed.[13]

Unlike the proposal made at the end of the previous section, Raffman provides a new theory of vagueness with a different answer to the semantic question that combines nicely with the rest of the contextualist framework. As this theory claims to preserve classical logic it will be the target of the remainder of this chapter.

2.4.1 Framework

The framework for Raffman's theory begins with Kaplan's[14] analysis of demonstratives. Kaplan takes Frege's distinction between the sense and reference of an expression and divides sense into character and content.[15]

The character of a demonstrative is similar to the dictionary meaning of the word. For example, 'she' might be defined in the following way: 'she' refers to some contextually relevant female-gendered object. The character of a demonstrative is set by linguistic conventions. The content of the demonstrative is fixed by the context. So, when Winston utters "she is drinking gin" while pointing at a woman at the bar, the content of the word 'she' is set by the character and the context—the context being that Winston is making the utterance while pointing at a woman at the bar and intending to talk about her. The content, and therefore, the referent of the instance of 'she' in Winston's utterance is the woman sitting at the bar.

[11] Raffman, 1994.
[12] Raffman, 1996.
[13] Raffman, 2013, *xi*.
[14] Kaplan, 1979.
[15] Kaplan, 1979, p. 83.

Raffman claims that vague words, like demonstratives, have some context-invariant meaning. Consider a case where I point to something and say 'That's a tall one.' You do not know what I am pointing at, and so you do not have the full context of the utterance.

> Assuming you are a competent speaker of English, you understand at least part of what I have said—roughly, *that that* [to which I am pointing] *is large in spatial height*, or, perhaps better, *relatively large in spatial height.* What you understand is something like a character of the predicate 'tall.'[16]

Vague words have some stable content that is not context-variant, and this stable content resembles the characters of demonstratives in Kaplan's theory. Vague words are not the same as demonstratives. We have already seen a difference in the discussion of Stanley's sorites paradox above. However, there does seem to be some similarity with respect to the character/content framework.

The next step in setting up the framework for Raffman's theory is to spell out the material that forms the context. The entire context, whatever that may require, is captured by what Raffman calls a V-index. A V-index includes elements like a contrastive category, a comparison class, etc.

$$\text{Shaquille is tall.} \tag{2.3}$$

A V-index that can give a content for 'tall' in (2.3) might look like $\langle$height, average height, basketball players, 2000, $w_{@}\rangle$. The character of 'tall' involves being large in spatial height. The context tells us that we are evaluating with regard to height, in contrast to average height, relative to basketball players in the year 2000 in the actual world.

2.4.2 The Multi-Range Theory

A character and a context are not sufficient to pick out a unique extension for vague words. Raffman argues that the character and v-index provide a number of ranges

[16] Raffman, 2013, p. 76.

of application for the word. These ranges of application provide acceptable stopping points in the application of the predicate. For example, consider a sorites series for tallness starting at 230 cm and going down to 160 cm. We will use a fixed context, described by the V-index, $\langle$height, average height; NBA players, 1980-2019, $w_@\rangle$. For the purposes of the example we will assume that a person with a height of 210 cm counts as a clear borderline case.[17] One person may stop applying 'tall' when they get to 208 cm in our sorites series. Another might stop at 212 cm. Each of these points is an acceptable stopping point and so the ranges (230 cm, 208 cm) and (230 cm, 212 cm) are ranges of application picked out by the character of 'tall' and the provided V-index.[18]

Raffman argues that the decision of where to stop applying 'rich' in our sorites series is arbitrary, and so as long as one stops in the border area they are not wrong. There is no principled reason to stop in one spot over another in the border area. Since we have no nontrivial reasons for stopping in one place over another, there is no nontrivial reason for choosing one range over another.

Raffman is attempting to give a theory of vagueness that fits as closely as possible our competent uses of vague words.

> All else being equal, a semantics should square with competent use. Hence at this juncture I propose to take the character of competent use as evidence of the semantic structure of vague words. Specifically, I propose that the multiple competent ways of applying a vague predicate relative to a given V-index reflect multiple *ranges of application* in the semantics of the term.[19]

In a given circumstance, much of which will come from the V-index, we can get an extension from a range of application. So, once we have the character and a V-index, we get a set of ranges of applications corresponding to competent ways of using the word. From the ranges of applications, we get a set of extensions.

[17] The average height of an NBA player is 200 cm.

[18] These ranges have upper limits because they are referring to stopping places on a particular sorites series where the tallest person in the series is 230 cm in height. This is not to say that only heights within the range (230 cm, 208 cm) are tall. Rather, anyone taller than 208 cm will count as tall according to this range of application.

[19] Raffman, 2013, p. 96.

2.4.3 Relativized Truth

Raffman's theory may sound a lot like supervaluationism so far. Supervaluationists claim that truth just is truth on all precisifications where a precisification sets precise borders for all of our words. The extensions we get from the ranges of applications seem like precisifications. The big difference between Raffman's theory and supervaluationism is that she does not identify truth with truth on all ranges of applications.

> Lastly, whereas the supervaluationist identifies ordinary everyday truth with truth on all admissible precisifications or "super-truth," the present view identifies ordinary everyday truth with truth relative to a single range of application.[20]

The truth or falsity of a sentence containing vague words is relativized to the ranges of applications of the vague words given the context.

> According to the multiple range theory, the truth or falsity of a sentence containing a vague word is relative to a V-index and a range of application.[21]

Raffman's theory is an example of what are called plurivaluationist theories. Instead of treating a sentence like "Shaq is tall" as true simpliciter, plurivaluationists relativize truth to either more precise versions of the sentence or contexts that give us precision. Since plurivaluationists relativize truth to more precise circumstances, they are better able to preserve classical logic.

Though Raffman tries to make the case that ranges of applications combined with V-indexes are not precisifications, the differences she puts forth in support of her claim fail to provide a significant difference. She provides four points of difference between supervaluationist precisifications and ranges of application.

1. A complete precisification of a predicate F contains both F and not-F items with a sharp boundary between them. A range of application of F contains only F items and no boundary.
2. A precisification can contain "gappy" items with respect to which "x is F" receives no truth-value. This does not happen with ranges of application.

[20] Raffman, 2013, p. 103.
[21] Raffman, 2013, p. 109.

3. The predicate 'borderline *F*' has multiple ranges of application, just like any other vague predicate.
4. Supervaluationism identifies everyday truth in terms of truth on all precisifications. The multi-range theory identifies everyday truth with truth on a single range of application.[22]

Regarding number 1, this is an unimportant difference between ranges of application and precisifications. After all, a range of application does create a sharp boundary between all of the values within the range and all of the values outside of the range. Making this sharp boundary is important for getting an extension relative to a range of application and a V-index. If a range of application did not create a sharp boundary between the objects to which the predicate applies and the objects to which it doesn't, then the multi-range theory would not preserve classical logic.

Moving on to number 2, it is important to note that supervaluationists will not use 'precisification of *F*' that divide the objects in the universe into ones that are *F*, ones that are not *F*, and ones that are neither. One could precisify a predicate in this way, but this possibility does not act as a wedge between Raffman's use of ranges of application and a supervaluationist use of precisifications.

As for number 3, Raffman is arguing that some ranges of application of 'borderline *F*' contain all and only borderline *F* cases. She claims this is not possible with precisifications because they must classify some items as *F* and some as not-*F*. Number 3 is no different than number 1; it is just one level higher in terms of higher-order vagueness. Raffman seems to have mistaken what a precisification of the predicate 'borderline *F*' would involve. Just like the multi-range theory, supervaluationism must contend with higher order vagueness. 'Borderline *F*' can be vague since it can be unclear which precisifications are acceptable for determining supertruth. A precisification of 'acceptable precisification' will allow us to precisify 'bordeline *F*' and will separate cases into borderline-*F* cases and not-borderline-*F* cases. This higher-order precisification will not resolve the first-order issue by separating objects into *F* and not-*F* cases.

Finally, number 4 is the difference with which I started this subsection. This is definitely a difference between supervaluationism and the multi-range theory, but

[22]Raffman, 2013, pp. 102-103.

it is not a difference between ranges of application and precisifications. Raffman's multi-range theory amounts to a relativization of truth and falsity to the kinds of precisifications used by supervaluationism.

2.4.4 Higher-Order Vagueness

The first response that might come to mind when looking at the multi-range theory is that it is subject to higher-order vagueness. The character and context of our vague words are supposed to give us a set of ranges of applications that correspond to competent stopping points for the application of vague words. However, it is vague what counts as a competent stopping point. So, it is vague which ranges of application are picked out by a character and context. Raffman claims that "range of application" has its own ranges of application. In turn, "range of application of 'range of application'" can have ranges of application, and so on *ad infinitum*.

Let's see how higher-order ranges of application work. Many heights (in cm) are acceptable cut-off points for the application of the predicate "tall", but many are not. Suppose h is a height that is not a clearly acceptable cut-off point—it might be too high, or too low, but it's neither clearly too high nor clearly too low. Assuming that context is fixed, let a be the borderline range of application that sets h as the minimum for the application of "tall". Then a is not clearly a range of application, because it might be unacceptable, but it is a borderline range of application. This means that some higher-order ranges of application include a as a range of application, and others exclude it. Now, suppose that Winston's height exceeds h. Then "'Winston is tall' is true relative to a" is true on any range of application that includes a as a range of application, but false on any range of application that excludes a as a range of application.

Raffman notes that at every level of the higher-order vagueness structure, "range of application" will be vague and will therefore have ranges of application. Let's consider a range of application, r, of the predicate 'tall' such that some ranges of application of 'range of application of tall' include r and others do not. Still further, all of the ones that include r are themselves borderline acceptable ranges of application of 'range of application of tall.' There is nothing blocking the possibility of a first-order range of application being borderline borderline ... borderline acceptable. When this happens, it is indeterminate whether or not truth can be relativized to r.

We will delve deeper into the phenomenon of higher-order vagueness as this dissertation proceeds. Here is a brief note to clarify the argument in this subsection. Take Winston as our example. Suppose that he is borderline borderline tall. That is to say, it is unclear whether or not Winston is borderline tall. He's in the border area between the borderline tall cases and the clearly tall cases. We can go further with borderlineness. Winston's height may be such that it is in the border area between the borderline borderline cases of tall and the borderline clear cases of tall, in which case, he would be borderline borderline borderline tall. This process can be repeated *ad infinitum*, and so Winston may be infinitely borderline tall, i.e. borderline borderline ... borderline borderline tall.

The argument against the multi-range theory that I give in this subsection claims that these infinitely borderline cases make it such that we cannot get an acceptable range of application to which we can relativize truth. A range of application may be such that it is infinitely borderline acceptable. Can truth be relativized to this range of application? We shouldn't answer 'yes;' truth is relativized to acceptable ranges of application. For the same reason, we shouldn't answer 'no;' this putative range of application isn't clearly unacceptable. Note that the example with baldness is not an isolated case. Every vague word is open to the possibility of this kind of infinite borderlineness. As such, we can end up with cases with these ranges of application to which it is unclear whether or not we can relativize truth.

2.4.5 Vagueness All the Way Down

The following is a problem for any theory that attempts to get full precision, including both supervaluationism and the multi-range theory. A range of application needs some dimension on which to specify an acceptable stopping point. For example, a range of application for 'tall' needs to be specified on a dimension like 'height.' However, the predicates we use for setting acceptable stopping points are vague as well. Whether or not a person is 180 cm tall can be unclear—it relies on questions about which atoms belong to the person, which is vague. Merely having a set of ranges of application for 'tall' without also having sets of ranges of application for when an atom belongs to a person would fail to give us precise evaluations of some sentences. If a person is borderline 180 cm and we have a range of application that sets 180cm as the cutoff for 'tall' within some context, then that person will be indeterminately 'tall' relative to that range of application.

The obvious response is to combine the ranges of application for 'being an atom belonging to x' and 'tall' and relativizing truth to this combination. However, a problem emerges if we never reach a precise predicate that does not need multiple ranges of application.

If this process bottoms out at some precise dimension, then everything works out for the multi-range theory. Truth will be relativized to these complexes of ranges of applications. However, if there is no bottom to this process, if we encounter vagueness all the way down, then there will be no point at which we complete the process. Ranges of application are supposed to give us precise extensions, but they can only do so if the dimension on which the range exists is precise. If there never is such a precise dimension we may not be able to get a precise extension from a particular range of application.

2.4.6 Ordinary Truth

Supervaluationism identifies ordinary truth with supertruth, but Raffman's theory identifies ordinary everyday truth with truth relative to a range of application. Raffman considers this an advantage of her view. However, she states that she is avoiding giving an error theory and is instead preserving the ordinary sense of 'truth.'

The preceding subsections should show that the multi-range theory does not give us an ordinary sense of truth. Certainly, when the average person speaks of truth, they do not specify that they are speaking relative to a particular range of application. This alone is not a strong point against the multi-range theory because Raffman can give a story about how we do not need to be so precise in everyday language. Something that is clearly true will be true on all relevant ranges of application. However, this kind of story would support supervaluationism's identification of truth with supertruth more than the multi-range theories relativization of truth.

The real problems for the multi-range theory become apparent when we consider the complexity of the apparatus that emerges when we introduce higher-order vagueness. Even when we exclude the infinitely borderline acceptable ranges of application, we have cases like clearly borderline clearly acceptable ranges of application of 'bald.' In such a situation, the range of application of 'bald' is acceptable according the ranges of application of 'range of application of 'bald',' but some of those ranges of application are only borderline acceptable. As such, when I say "'Winston is bald' is true," I am relativizing the truth of 'Winston is bald' not just to ranges of application, but to ranges of application of 'range of application of 'bald."
This certainly does not seem like what we are doing when we assert the truth of a sentence.

2.4.7 Vagueness in the World

In chapter 4, I will argue that it is possible for there to be vagueness in the world. If vagueness is a metaphysical phenomenon, in addition to being a linguistic phenomenon, then the multi-range theory will, at the very least, not be a full theory of vagueness since it only describes linguistic vagueness. You could take the multi-range theory and adapt it to metaphysical vagueness. I will describe views like this when I discuss Elizabeth Barnes' metaphysical supervaluationist theory[23] and Nicholas J.J. Smith's plurivaluationist degree theory[24] in later chapters.

[23]Barnes, 2010.
[24]Smith, 2008.

Chapter 3

Precisifications

Introduction

In the previous chapter, we explored Diana Raffman's multi-range theory, which relativizes truth to ranges of application. I argued that ranges of application act a lot like the precisifications used by supervaluationists. This chapter is devoted to exploring other theories of vagueness that make use of precisifications. The views I discuss here include supervaluationism itself; semantic nihilism, the theory that sentences containing vague words are not truth-apt; and expressivism, the theory that utterances containing vague words are merely expressions of plans for the applications of predicates. Semantic nihilists like David Braun and Theodore Sider use precisifications to talk about assertability. Truth on all precisifications implies that it is acceptable to assert a sentence. John MacFarlane supplements his expressivist theory of vagueness by relativizing truth to precisifications—something we've seen already in Raffman's multi-range theory. For each of these theories, I will provide a number of objections. To supervaluationism and expressivism, the objections focus on the problems with using precisfications to evaluate truth-bearers. Against semantic nihilism, the objection focuses on the possibility of imprecise truth-bearers.

3.1 Supervaluationism

According to supervaluationism, the phenomenon of vagueness is one of semantic indeterminacy. Sentences containing vague words fail to pick out unique meanings. Rather, each such sentence picks out a number of similar meanings. Consider an obvious case of semantic indeterminacy—ambiguity. Suppose the following sentence is uttered in such a situation that it is unclear whether the sentence means that Winston is going to a financial institution or the side of a river.

Winston went to the bank. (3.1)

This sentence is ambiguous between two meanings:

1. Winston went to the money bank.
2. Winston went to the river bank.

The meaning of (3.1) could be either of these. Kit Fine argues that we can evaluate the truth-value of ambiguous sentences like (3.1).[1] Suppose that Winston went to both the money bank and the river bank. Fine claims that the sentence would be true. Suppose that, instead, he went to neither; then the sentence would be false. Finally, if he went to one and not the other, then the sentence would be indeterminate in value.

The same strategy can be applied to sentences containing vague words.

> Vague and ambiguous sentences are subject to similar truth conditions; a vague sentence is true if true for all complete precisifications; an ambiguous sentence is true if true for all disambiguations. Indeed, the only formal difference is that the precisifications may be infinite, even indefinite, and may be subject to penumbral connection. Vagueness is ambiguity on a grand and systematic scale.[2]

Take for example the following sentence contaning the vague word 'rich.'

Jeff Bezos is rich. (3.2)

(3.2) picks out a number of precise meanings corresponding to different cutoffs for the application of the predicate "rich."[3] It could mean 'Jeff Bezos has more than 1 billion dollars in net worth.' or it could mean 'Jeff Bezos has more than 1 billion and 1 dollars in net worth.'

[1] Fine, 1975.
[2] Fine, 1975, p. 282.
[3] 'Jeff Bezos' is also vague, so any precisifications of the sentence will also precisify things like the atoms in Jeff Bezos' body and exactly when he came into and will go out of existence, etc.

To evaluate sentences, we look at every precisification of the language. A complete precisification sets a precise cutoff for the application of every word. If (3.2) is true on each precisification, then it is called supertrue. If it is false on every precisification, then it is called superfalse. Finally, if it is true on some and false on others, then the sentence is indeterminate. Raffman's multi-range theory relativizes truth to precisifications; supervaluationism, on the other hand, identifies truth with supertruth. Truth isn't relativized to precisifications, but is defined using precisifications. For a sentence to be true is just for it to be supertrue.

In order to bring supertruth in accordance with the ordinary use of truth, we need to make a change to the definition of supertruth. Instead of evaluating a sentence as supertrue when it is true on all precisifications, supervaluationists usually treat sentences as supertrue when they are true on all admissible precisifications. This is also the move Raffman makes in the multi-range theory. Instead of being relativized to all possible ranges of application, truth is only relativized to acceptable ranges of application. If we use all precisifications in defining supertruth, then sentences like 'Jeff Bezos is rich' will not be supertrue. After all, Jeff Bezos is not rich according to a precisification that sets the cutoff for 'rich' at 1 trillion dollars.

One of the primary appeals of supervaluationism is that it preserves classical tautologies. Consider excluded middle.

Winston is bald or he is not bald. (3.3)

Regardless of how "bald" is precisified, Winston will fall on one side or the other of the cutoff. If according to a given precisification he is bald, then the left disjunct is true on that precisification. If according to a given precisification he is not bald, then the right disjunct is true. So, for any precisification, either the left disjunct is true or the right one is, and so, for any precisification, the sentence is true. Therefore, the sentence is supertrue. Neither "Winston" nor "bald" are important for this example to work. For any instance of excluded middle, the same reasoning will hold. Once precise cutoffs are set for every predicate, any object must fall on one side or the other of the cutoff.

3.1.1 Objection: Disjunctions and Existential Generalizations

In the next three subsections, I will explore objections to supervaluationism. The first of these objections turns the benefit of supervaluation into a flaw.

Braun and Sider as well as MacFarlane argue that supervaluationism gets the wrong semantics for disjunctions and existential generalizations. According to supervaluationism, a disjunction can be supertrue, even though neither of its disjuncts is supertrue. Suppose that Winston is a borderline bald man. "Winston is bald or he is not bald" is supertrue, since no matter how we precisify 'bald,' Winston will either fall within the extension of 'bald' or he'll fall outside of the extension. However, since Winston is in the border area for the predicate 'bald,' he will not be bald on every precisification, nor will he be not-bald on all precisifications. So, each disjunct is indeterminate. Supervaluationism preserves excluded middle, but at the cost of allowing a disjunction to be true when neither disjunct is true. The claim made by Braun, Sider, and MacFarlane is that any correct meaning of 'or' will be one where a disjunction is only true when at least one of its disjuncts is true. This intuition is compelling. If we know that a disjunction is true, it is reasonable to inquire about which disjunct is true. In cases like (3.3), the supervaluationist will have to say that the inquiry will fail because there isn't a true disjunct.

Existential generalizations have a similar problem. For example, "there is a smallest number of grains that make a heap" is supertrue, but there is no number of grains such that that number of grains is the smallest number that makes a heap. On every precisification, there is a sharp borderline for 'heap,' and so there is a number of grains such that a pile of sand with one fewer grains would not be a heap. However, there is no number of grains of sand such that on every precisification, that number is the smallest number of grains of sand that makes a heap. So, no instantiation of the above existential generalization is super true, but the generalization itself is supertrue. The intuition driving this objection is that an existential generalization can only be true if there is at least one true instantiation. This intuition may even be stronger than the one in the case of disjunctions. If we know that something exists, then we should be able to find which object it is. If our epistemic abilities are lacking, then an omniscient being should be able to determine which object it is. However, in the case of supervaluationism, existential instantiation is not valid.

Something feels dishonest in the supervaluationist semantics. Supervaluationists claim that Winston is either bald or not bald, but when pressed to say which one, they cannot. The inability to say which one is not an epistemic limitation, but a semantic one. John MacFarlane claims that the supervaluationist is mistaking acceptability for truth here.[4] It is acceptable to assert the disjunction even though it is not acceptable to assert either disjunct. It can be acceptable to assert an existential generalization even though it is not acceptable to assert any instantiation. MacFarlane's explanation of the mistake made by supervaluationists is similar to David Braun and Theodore Sider's. They claim that truth on all precisifications leads to the acceptability of assertion, not to truth.

Both supervaluationism and the intuition espoused by Braun, Sider, and MacFarlane are onto something. We have already seen how the classical intuition about disjunction is compelling. However, the supervaluationist intuition is not wholly wrong. Imagine a uranium atom, u, going through $alpha$-decay. Suppose that u is somewhere in the border area between being a uranium atom and being a thorium atom. According to supervaluationism, 'u is a uranium atom or it's a thorium atom' is supertrue. No matter how we draw the borderline between 'uranium' and 'thorium' u will fall on one or the other of the two sides. It cannot fail to be a uranium atom and fail to be a thorium atom—it cannot be the case that both disjuncts are false. This doesn't match our normal understanding of disjunction, but it does contain elements of the meaning of a disjunction. Part of what it means for something to be a disjunction is that it presents two options. The supervaluationist's 'or' does at least that much, making it similar to our ordinary meaning of 'or.'

This may seem little compensation to those who find the supervaluationist semantics for disjunctions problematic. The logic I will offer in chapter 5 will be open to either of these understandings of disjunction; however, I take it that Braun, Sider, and MacFarlane are correct in thinking that our normal uses of 'or' require a true

[4]MacFarlane, 2016, p. 24.

disjunction to have a true disjunct. As for existential generalizations, there seems to be a more serious problem. It may be possible to say that the supervaluationist sense of 'or' matches something in our ordinary meaning of the word. However, the ordinary meaning of 'exists' requires that existential instantiation be valid.

3.1.2 Objection: Permissible Precisifications

As I discussed alongside the definition of "supertrue" above, we cannot include all possible precisifications in determining the truth value of a vague word. If we did, then almost every sentence would come out as indeterminate. Instead, we only include permissible precisfications. What counts as a permissible precisification is vague. We can construct a sorites argument to show this.

Base Premise: A precisification that sets the cutoff for 'tall' at 200cm is a permissible precisfication of 'tall.'

Inductive Premise: For all n, if a precisification that sets the cutoff for 'tall' at n cm is a permissible precisification of 'tall,' then so is one that sets the cutoff at $(n-1)$ cm.

Conclusion: A precisification that sets the cutoff for 'tall' at 50 cm is a permissible precisification of 'tall.'

The sorites-susceptibility of "permissible precisification" should be sufficient evidence of its vagueness. The vagueness of "supertrue" and "superfalse" are problematic for the supervaluationist because they sometimes hinder us from getting a truth-value for a sentence. A supervaluationist may attempt to apply the supervaluationist semantics in the case of higher-order vagueness. We need only look at the permissible precisifications of 'permissible precisification.' If for every permissible precisification of 'permissible precisification' a given sentence comes out supertrue, then it is supertrue that the sentence is supertrue. Similarly, it can be superfalse or indeterminate whether or not a sentence is supertrue. This process will continue *ad infinitum*. Since 'permissible precisification' is vague, which precisifications of 'permissible precisification' are permissible is vague.

The problem of higher-order vagueness hits supervaluationism particularly hard because the vagueness is located in the definition of truth. The supervaluationist must perform an infinite number of operations in order to get a truth-value for a single sentence. To say that a sentence is supertrue requires that on all permissible precisifications of "permissible precisifications of 'permissible precisifications of ...'" the sentence is supertrue.

Permissible precisifications present another issue for supervaluationism or any theory that makes use of them. Some predicates do not appear to have any permissible precisifications. That is to say, any potential borderline for the predicate seems incorrect. Consider the predicate 'rich.' Attempt to come up with a dollar amount such that anyone with a net worth greater than that amount is rich and anyone with a lower net worth is not rich. The task is incredibly difficult. Even if two people with the same net worth and in the same context produce numbers, they will likely disagree about the permissibility of the other person's answer. Raffman argues that any point in the border area of a predicate is a permissible stopping place in a sorites series. However, it may be the case that we, in fact, feel that none of them are permissible stopping places. Our tendency to overshoot the border area of a predicate in forced-march sorites may be an indication that we found none of the points throughout the border area permissible places to stop applying the predicate.

If there are no permissible precisifications of a predicate, then there is no way to use the supervaluationist semantics to evaluate that predicate. However, many of our vague words appear to work similarly to 'rich.' It would be a serious problem if we couldn't apply this theory of vagueness to many of our vague words. Ditching the use of permissible precisifications and just using all precisifications isn't a good option either. Doing so will render almost every sentence indeterminate.

3.1.3 Objection: Incomplete Precisifications

In chapter 2, I argued that Raffman's multi-range theory was unable to handle the possibility that vagueness does not bottom out at some set of precise predicates. One might attempt to precisify 'bald' by appealing to numbers of hairs, but since it is vague what counts as a hair, we need to precisify further. 'Hair' might be precisified by number of keratin molecules, but this runs into vagueness when we consider whether or not two keratin molecules are attached. If, as we move to more fundamental properties, we never find a perfectly precise layer, we can never get a full precisification. We will either have to continue searching infinitely for a precisification, or, if we stop somewhere along the way, there will be some borderline cases left over.

It is not the end of the world if there are borderline cases left in a precisification, but if one's goal is to preserve elements of classical logic in the way supervaluationism does, then these cases will be a problem. Let's look at how it would affect the supervaluationist's semantics. Supervaluationism preserves classical tautologies because classical logic holds in each precisification. If we have indeterminate values within the precisification, then either we don't have a precisification at all (because precisifications require complete precision) or we have a precisification that, though it makes things more precise, fails to provide classical values for every atomic sentence. One of the values of supervaluationism is its ability to preserve classical logic, at least to some degree. If one's goal is to preserve as much of classical logic as possible, supervaluationism is not a fruitful theory to adopt. If one is less concerned with preserving classical logic, then, I will argue, other non-classical theories are preferable to supervaluationism.

3.1.4 Quick Summary

I've presented three objections to supervaluationism. First, supervaluationists are charged with providing an incorrect semantics for disjunctions and existential generalizations. Though supervaluationism preserves classical tautologies, it does so by altering the meanings of 'or' and 'exists.' Second, the vagueness of the definition of supertruth creates a problematic infinite regress. For some sentences, it can be the case that no truth-value can be determined because doing so would require resolving an infinite number of borderline cases of permissible precisfications. Finally, the possibility that vagueness never bottoms out at some set of precise predicates brings into question the possibility of a complete precisification of our predicates that removes all indeterminacy.

In the next section, I will describe and object to semantic nihilism. Some of the proponents of semantic nihilism make use of precisifications, but, as we saw in subsection 3.1.1, they reject the identification of truth on all precisifications with truth. Instead, they identify truth on all precisifications with assertibility.

3.2 Semantic Nihilism

Semantic nihilism is the theory that sentences containing vague components are not truth-apt. Gottlob Frege is cited as one of the early proponents of semantic nihilism. For Frege, words and sentences have both a sense and a reference. The sense of a word is akin to its dictionary meaning. The referent of a word is the thing/s it picks out. For example, the sense of 'baseball' is something like 'ball made of leather stitched together that is approximately 230 mm in diameter.' The referent of 'baseball' is the set of all baseballs. Frege claims that vague words lack a referent and that a sentence where some of the components lack reference also lacks a referent. Since the referent of a sentence is its truth value, a sentence containing a vague word lacks a truth value.[5] Frege does not give a clear argument, but the position he is taking is clear. The only way that predicates, for example, can be fully meaningful—having both a sense and a referent—is for them to pick out precise extensions. Since vague words do not, they are not fully meaningful. They are defective and should be purged from the logic.

One argument that nihilists have used in various forms makes use of sorites arguments like the following two.

P1a A single grain of sand does not make a heap.

P2a If n grains of sand do not make a heap, then n + 1 grains do not.

Ca A billion grains of sand do not make a heap.

[5] While Frege's views on vague concepts seem to indicate a nihilist position, they are also consistent with epistemicism. The determining factor will be how frequent Frege thinks the phenomenon of vagueness is. For more on Frege's actual views on vagueness, see Puryear, 2012.

P1b A billion grains of sand do make a heap.

P2b If n grains of sand make a heap, then n - 1 grains do as well.

Cb 1 grain of sand makes a heap.

Both of these arguments appear to be sound, but if that is the case, then 1 grain of sand both does and does not make a heap. If we cannot change instantly from heap to non-heap or vice versa, then we will never change. The predicate "heap" appears to both allow any number of grains to count as a heap or none. One response is to take "heap" to be meaningless. This is the route that Frege seems to take.

Kirk Ludwig and Greg Ray also take this route. They present the argument as a proof of the following theorem[6]:

THEOREM: IF (i) for all n, 'heap' applies to n iff 'heap' applies to $n + 1$ and (ii) for all n, 'heap' fails-to-apply[7] to n iff 'heap' fails-to-apply to $n + 1$, THEN (iii) 'heap' applies to every n, (iv) 'heap' fails-to-apply to every n, or (v) 'heap' neither applies nor fails-to-apply to any n.[8]

Here the antecedent captures P2a and P2b of the above sorites arguments. It says that if 'heap' does not draw a sharp border between cases where it applies and cases where it fails-to-apply, then it applies to every n, fails-to-apply to every n, or neither applies nor fails-to-apply to every n. Since 'heap' is vague, its P2a and P2b are true, the antecedent is fulfilled and one of the disjuncts of the consequent must be the case. Ludwig and Ray choose (v), stating that (iii) and (iv) are unappealing. "(iii) and (iv) are options which we take it no one is willing to accept. Each is completely incompatible with our practices."[9] Since there is nothing in the extension of "heap" and nothing in the anti-extension of "heap," Ludwig and Ray claim that "heap" is defective and sentences containing it are not truth-evaluable.

Contrary to Ludwig and Ray, there is someone who is willing to accept (iv). Peter Unger privileges one direction of the sorites over the other. He argues that the first sorites argument above is a direct argument that there are no heaps and the second is an indirect argument for the same conclusion.[10] In the indirect argument, we assume

[6]Adapted for use with "heap"

[7]'Fails-to-apply' does not mean to not apply—for 'heap' to fail-to-apply to n means that n is in the anti-extension of 'heap.'

[8]Ludwig and Ray, 2002, p. 435.

[9]Ludwig and Ray, 2002.

[10]Unger, 1979, p. 118.

that there are heaps, but if there are heaps, then we run into the absurdity that 1 grain of sand makes a heap. Accepting the apparent soundness of the sorites arguments leads us to either reject the meaningfulness of vague predicates or to claim that they are extension-less. We will either end up with non-truth-apt vague sentences or any sentence that entails that a vague predicate has a non-empty extension will be false.

The nihilist cannot appeal to sorites arguments as evidence for their position without first independently undermining the non-classical theories of vagueness by establishing the truth of P2a and P2b. For example, I will argue in chapter 5 that some instantiations of P2a and P2b are indeterminate in truth, and that therefore the universal generalizations are indeterminate. Since P2a and P2b are both indeterminate, the arguments are not sound. If we cannot get to the conclusions, then we do not have an inconsistency such that all or no number of grains make a heap. Put in the terms of Ludwig and Ray's theorem, the antecedent needs to be established.

It is not clear what argument can be given for P2a, other than our intuition that it is true. This intuition can be handled by non-classical theories of vagueness in a variety of ways, so it does not provide strong evidence. For example, proponents of non-classical logics have contextualism as an available option for explaining the psychological pull of P2a. In addition to the explanations of the intuitions that P2a is true, if the only thing supporting P2a is our intuition that it's true, then that intuition must be weighed against our intuition that the conclusion of the sorites argument is false. When asked to choose between P2a and the negation of Ca, it seems likely that many will choose the negation of Ca. Without a strong argument for P2a, the nihilist will need to argue that we should accept P2a because none of the analyses of it that make it not true succeed, but that defeats the purpose of using the sorites argument as an attempt at providing a positive argument for nihilism.

3.2.1 Is Semantic Nihilism Self-Defeating?

Semantic nihilism faces a difficult problem. The words that are used to express Unger's and Ludwig and Ray's theories contain vague words. For example, the word 'vague' is vague. Here is Sorensen's argument that 'vague' is vague.[11] Consider the predicate '1-small' which applies to numbers iff they are either small or less than 1. Now, we will construct a series of predicates starting with '1-small' and going to '10000-small' where for any n, 'n-small' applies to numbers iff they are either small or less than n.

Base Premise: '1-small' is vague.

Inductive Premise: If 'n-small' is vague, then so is '$(n+1)$-small.'

Conclusion: '10000-small' is vague.

The base premise is clearly true. Take a number that is borderline small. Since it is greater than 1, whether or not it is 1-small will be indeterminate. The conclusion is clearly false. Since 10000 is a big number the borderline cases of 'small' will all be less than 10000, and so they won't be borderline cases for '10000-small.' There are borderline cases somewhere between 1-small and 10000-small, and so there are borderline vague predicates.

The vagueness of 'vague' implies that Unger's theory, by its own lights, is false. Unger's theory may be summarized as 'Sentences that entail that a vague predicate has a non-empty extension are false.' The predicate 'vague predicate' is itself vague. Unger treats this thesis as presupposing that there are some vague predicates, and so, if Unger's thesis is true, then it is false.[12] In the case of Ludwig and Ray, their theory would not be truth-apt. Since 'vague' is defective in the same way as 'heap'—it allows for borderline cases—sentences containing 'vague' are not truth-evaluable. The thesis of semantic nihilism of the kind provided by Ludwig and Ray is that sentences containing vague components are not truth-apt. This thesis, by its own lights, is not truth-apt.

[11] Sorensen, 1985.

[12] If there are no vague predicates, then Unger's thesis will be fine, but it will be vacuously true.

3.2.2 Braun and Sider's Nihilism

David Braun and Theodore Sider also argue for the position that sentences containing vague words are not truth-apt. To defend it from the objection in the previous subsection they make use of precisifications. Braun and Sider assume that only fully precise propositions can be true or false. Since vague sentences fail to express unique propositions, they cannot be true or false. So far, this is just a standard nihilist theory. However, they attempt to lessen the counterintuitiveness resulting from the revisionary and self-undermining nature of nihilism by appealing to precisifications. While vague sentences fail to pick out unique propositions, they do pick out multiple similar propositions. The propositions picked out by a vague sentence correspond to the different ways the sentence could be precisified to give sharp extensions for the predicates and clear referents for the subjects. Braun and Sider stop short of identifying truth with supertruth, but they make use of the property <being true on all precisifications>. They claim that, though our vague sentences are not actually truth-apt, we can often ignore their vagueness for the sake of communication—it can be permissible to assert a sentence that is not truth-apt. The condition on assertibility is that the sentence be approximately true, which Braun and Sider define as truth on all admissible precisifications.

> On our view, ordinary speakers typically and harmlessly ignore vagueness. And when doing so, it is reasonable to speak, in a sense to be defined, the approximate truth.[13]

Returning to the question of self-defeat, Braun and Sider's theory claims that sentences containing vague words are not truth-apt. Their own theory contains vague words, and so the sentences comprising it are not truth-apt. Again, this is because what counts as a vague sentence is itself vague. Braun and Sider argue that this does not undermine their theory, because, though it is not true, it is approximately true. Under all precisifications of the vague words used to express their theory it comes out true, and so it is approximately true. As such, it is acceptable for them to assert it. Similarly, they argue that the revisions that semantic nihilism

[13] Braun and Sider, 2007, p. 4.

requires are not as problematic as they might first appear. Of course, all vague sentences fail to be truth-apt, but some can be approximately true, and therefore assertable. The reason they are assertable is that when something is approximately true it is permissible to ignore the vagueness in the sentence. When we normally communicate we ignore the vagueness of our words, and so we have no problem asserting and accepting sentences. However, when we are presented with a borderline case, we can no longer ignore the vagueness, and so we have to accept the untruth of the sentence. Braun and Sider claim that their theory is true on all admissible precisifications, making it approximately true, and therefore assertable.

There is no question that this response alleviates some of the pressure faced by nihilism. In particular, the objection that nihilism is self-defeating seems much weaker. However, this defense does not remove all of the counterintuitiveness of semantic nihilism. Braun and Sider's theory is still an error theory that concludes that we are almost always wrong about our own sentences. Sentences like "Bill Gates is rich" and "Quarks are small relative to planets" are not truth-apt under their account. Perhaps more problematic, classical tautologies are not true according to nihilism. "If Winston is bald, then Winston is bald" and "it is raining or it is not raining" are both not truth-apt. Similarly, "Spot is a dog and Spot is not a dog" is not false. Though Braun and Sider admit that the approximately true sentence like "Bill Gates is rich" are perfectly acceptable to assert, that does not change the fact that their theory undermines most of our everyday uses of the truth predicate.

Even if Braun and Sider have succeeded at reducing the intuitive costs of semantic nihilism, they are far from having eliminated them. Nihilists will need to provide a good reason for us to accept these costs. Let's now consider the argument Braun and Sider offer in favor of nihilism.

> We assume that the properties, relations, and propositions that are candidates for being the meanings of linguistic expressions are precise: any n-tuple of objects either definitely instantiates or definitely fails to instantiate a given n-place relation, and any proposition is either definitely true or definitely false. But the facts that determine meaning (for instance, facts about use, naturalness of properties, and causal relations between speakers and properties) do not determine a unique property

> to be the meaning of 'red'. There is no property that 'red' uniquely expresses, and therefore no unique proposition that a sentence containing 'red' expresses. Vagueness is a type of semantic indeterminacy.[14]

> To be either true or false, a sentence must have a unique meaning. Ambiguous sentences do not have unique meanings. Therefore, they are neither true nor false. Similarly, sentences containing vague expressions do not have unique meanings; therefore, they too are neither true nor false.[15]

Braun and Sider make a number of assumptions in their argument. They assume that propositions are composed of perfectly precise properties and relations that an object or n-tuple of objects either definitely instantiates or definitely fails to instantiate. Even more controversially, they assume that propositions are only either true or false. They assume also that to count as a truth-bearer, a sentence must pick out a unique proposition. Granting these assumptions, the conclusion follows almost directly. Few would argue that vague sentences pick out unique precise propositions. However, these assumptions should not be granted easily.

3.2.3 Objection: The Vagueness of 'Truth-Apt'

Braun and Sider's appeal to approximate truth may relieve some of the tension caused by the vagueness of 'vague.' However, another problem for the theory comes from this same source. Propositions, according to Braun and Sider, are perfectly precise. However, as we have seen, when a predicate is precise is a vague matter. There can be indeterminately vague predicates and hence indeterminately precise predicates. Suppose that n-small is one of these indeterminately vague predicates. What do we then make of a sentence like 'x is n-small.' Suppose that $x < n$, then the sentence should be true. However, it appears to be indeterminately truth-apt according to Braun and Sider's view. Truth-aptness requires precision. It is indeterminate whether or not 'x is n-small' is precise. So, it is indeterminate whether or not the sentence is truth-apt. It may pick out a precise proposition, or it may pick out a number of precise propositions.

[14] Braun and Sider, 2007.

[15] Braun and Sider, 2007.

As a response, semantic nihilists may claim that n-small does not correspond to a property in the world, and so there are no propositions that 'x is n-small' picks out. Taking this route would require that the semantic nihilist also exclude precise predicates like 10000-small. If the precise predicates like 10000-small correspond to properties then we could construct a sorites argument that demonstrates the vagueness of 'precise property.' One may not find it problematic to reject putative properties like <being 10000-small>, but one cannot reject them on the basis that they are vague. Rather, they must be rejected on something like a naturalness criterion. Perhaps disjunctive predicates like 10000-small are not natural enough to pick out properties.

3.2.4 Objection: Metaphysical Vagueness

In chapter 1, I gave a preview of the argument that vagueness is metaphysical that will appear in the next chapter. Here I will draw another consequence of the possibility of metaphysical vagueness. If it is possible that there is vagueness in the world, then it is possible that the objects in the world or the properties instantiated by those objects are vague. In the former case, it may be that the very existence of an object can be indeterminate. In the latter, it may be indeterminate whether or not a given object instantiates a given property.

This possibility is enough to undermine Braun and Sider's semantic nihilism. It is not the case that "any n-tuple of objects either definitely instantiates or definitely fails to instantiate a given n-place relation." Even our sentences that pick out unique propositions may fail to pick out something that is definitely true or definitely false. The source of vagueness is not only a failure of our sentences to pick out unique precise propositions—vagueness also comes from the world not being precise.

I will stress in the next two chapters that I am concerned foremost with providing a logic that explains the metaphysical phenomenon of vagueness. Providing a full theory that includes the connection between language and the world is beyond the scope of this dissertation. However, a point can be made here about the truth-evaluability of sentences containing vague words. We do not need to jettison such sentences from our logic if their underlying truth-bearers, propositions (in the sense used by Braun and Sider), contain vagueness as well. Sentences can pick out unique vague propositions, and therefore have a 'unique meaning' as Braun and Sider require for truth-evaluability. However, I do not think we need to stop there in evaluating sentences containing vague components.

3.2.5 Objection: Borderline Approximate Truth

Truth on all precisifications entails that a sentence is approximately true, according to Braun and Sider's view. Of course, this can't be **all** precisifications; we only include those that are picked out by the sentence, the permissible precisifications. 'Permissible precisification' as we saw in chapter 2 is vague. So, when something counts as approximately true can be vague. Suppose it is indeterminate whether or not a given sentence is true on all of its precisifications. There may be one precisification on which the sentence is false but that precisification is only borderline permissible. The presence of borderline approximately true sentences is problematic for Braun and Sider's response to the problem of self-defeat. They claim that their theory is approximately true because it is true on all precisifications. Since 'approximately true' is vague, the claim that their theory is assertible is, according to semantic nihilism, not truth-apt.

3.2.6 Simplicity

At the end of chapter 1 I remarked that among the theories that preserve classical logic epistemicism is not the simplest. Semantic nihilism has the benefits that come from the simplicity of classical logic without the added complexity of epistemicism's linguistic theory. Semantic nihilism has its own theoretical vices that counterbalance its virtues. It trades one counterintuitive consequence for another. Instead of the counterintuitiveness being located in the extreme precision of our language, semantic nihilism's counterintuitiveness comes from its being an error theory. When we think we have asserted something that can be true or false we are almost always wrong. Since, as we saw in the previous section, it may be very hard to find the permissible precisifications, it may also be the case that what we assert isn't clearly approximately true. The counterintuitive elements of semantic nihilism mixed with the possibility that vagueness is metaphysical give us the result that semantic nihilism does not hold a strong enough advantage in theoretical virtues to give it preference over other theories.

3.2.7 Quick Summary

Semantic nihilism claims that sentences containing vague components are not truth-apt. This claim appears to be self-defeating because of the vagueness of 'vague.' Braun and Sider attempt to get out of this problem by using precisifications. They claim that their theory, though not truth-apt, is approximately true, and therefore assertible. This doesn't entirely resolve the counterintuitiveness of saying that 'the Eiffel Tower both exists and does not exist' is not false. However, it does resolve some of the problems the theory faces. Braun and Sider's argument for nihilism hinges on the primary truth-bearers in the world being perfectly precise. In chapter 4, I will argue that these primary truth-bearers need not be precise. So, once we have a complete theory of how language connects to these truth-bearers, it can be the case that a vague sentence is truth-apt by virtue of its connection to a vague proposition.

In the last section of this chapter, I will describe a theory similar to semantic nihilism in its rejection of the truth-aptness of everyday sentences and similar to the multi-range theory in its relativization of truth to precisifications.

3.3 Expressivism

John MacFarlane makes the claim that our assertions of vague sentences are not assertions of propositions that can be true or false, but rather expressions of plans for the uses of our words. He calls the theory expressivism because of its relation to Gibbard's account of normative language. Gibbard's expressivism says that a sentence like "I ought to pack now" does not express some fact, but rather just a plan. MacFarlane attempts a similar move with vague sentences like his (2) below:

(2) That is a large apple.

> To judge (2), in a context where it is taken for granted that the demonstrated apple is 84mm across, is just to plan to apply the concept large (and the predicate 'large') to apples 84mm in diameter and larger. To assert (2) is to express this plan, and to propose it for joint adoption. Neither judging (2) nor accepting an assertion of (2) requires recognizing the truth of a proposition that it expresses.[16]

For his semantics, MacFarlane makes use of Gibbard's hyperplans—fully specified plans. I may have a plan to take a train from Prague to Dresden, but this plan does not include any information about what foot I should step on the train with. A hyperplan specifies every result of every eventuality. In the context of plans for the uses of words, a hyperplan would specify cutoff points for the applications of vague predicates. According to MacFarlane, a sentence is true or false relative to a context, world, and hyperplan. Our normal vague sentences are not truth-apt, since even though the context and world of evaluation may be set, the plans that are expressed when they are uttered are not hyperplans.

> Since a speaker's state of mind is generally compatible with many hyperplans, this semantics allows that nothing about the speaker's state of mind settles which delineation to use. Accordingly, it does not issue in truth values for sentences at contexts.[17]

The hyperplans give us something like the precisifications of supervaluationism. Understanding MacFarlane's theory in this way lets us see the connection with Braun and Sider. Sentences can be evaluated for truth relative to a precisification, but they are not true simpliciter.

Like Braun and Sider, though MacFarlane's theory does make nihilism more appealing, it does not remove the concern that it is too revisionary. "Winston is bald if Winston is bald" still comes out as non-truth-apt according to expressivism. This means that expressivism also undermines our normal use of the truth predicate. Given the nature of this revision, we need strong reasons for accepting the view. However, MacFarlane's expressivism goes one step further in its revisionism. In addition to the counterintuitive consequences of nihilism, expressivism makes the claim that we are doing something different than asserting something that can be

[16] MacFarlane, 2016, p. 2.
[17] MacFarlane, 2016.

true or false when we assert a declarative statement—we are merely asserting a plan. This is contrary to our intuition about what we are doing when we assert, for example, 'Jeff Bezos is rich.' Braun and Sider's theory maintains that we are still doing the same activity as when we assert "2 + 2 = 4"—we are asserting that a sentence is true, not merely expressing a plan. In the case of 'Jeff Bezos is rich,' we are just wrong when we think that we have successfully asserted something that can be true. In this sense, Braun and Sider's theory is less revisionary than expressivism.

We can make a strong case that sentences with vague components are truth-evaluable. MacFarlane claims that to evaluate a sentence we need to set its context, world, and hyperplan, but we do not in fact need all of these to evaluate a sentence. Suppose that the sentence is a necessary truth. Then we do not need to set a world, since the sentence is true in every world. We also do not need to set a hyperplan in order to evaluate a sentence. Consider the following analogy to plans. I might have a plan to take a train from Prague to Dresden, but it is not a hyperplan because I have not created plans for what foot I will use to step onto the train. If I do, in fact, take the train from Prague to Dresden, then I have fulfilled my plan, regardless of what foot I used to step onto the train. Similarly, we can sometimes determine whether or not I have fulfilled a plan for the application of a predicate. When I assert "Jeff Bezos is rich," I have a plan for the application of "rich," though it is not a hyperplan. My plan at least counts all people with more than one hundred billion dollars in net worth as rich. As such, applying "rich" to Jeff Bezos fulfills my plan for the application of the predicate. We do not need a hyperplan to evaluate every sentence.

3.3.1 Objection: Plurivaluationism Revisited

MacFarlane's expressivism has one big difference from Braun and Sider's semantic nihilism. If there is no precise property in the world like <being a cat>, then there won't be a proposition ascribing that property to an object. The likely consequence of semantic nihilism is that the properties that exist in the world are not the ordinary properties like <being bald> or <being a baseball>. Instead, there are properties like <having fewer than 1028 hairs>. When we use words like 'bald' we fail to pick out one of these precise properties, sometimes picking out many of them. We do not need precise properties for truth-evaluability according to MacFarlane's expressivism. Instead, we need fully specified plans for the application of a predicate. However, this means that sentences containing vague words like 'cat' are truth-evaluable, but only relative to a hyperplan. We've already seen a view like this in Raffman's multi-range theory. Like Raffman's theory, the semantic aspects of MacFarlane's theory give us a kind of plurivaluationism—truth relativized to precisifications. Of course, there are differences between MacFarlane's expressivism and Raffman's multi-range theory. For example, MacFarlane gives a particular story about what we're doing when we make assertions containing vague words.

A point of interest here is that expressivism may not be subject to one of the objections I leveled against the multi-range theory in the previous chapter: namely, the vagueness of "acceptable" as it is applied to precisifications. MacFarlane could simply say that truth is relativized to hyperplans simpliciter, not just acceptable hyperplans. Our normal use of vague words only involves plans, and those plans will mostly be in the acceptable range for any given predicate. If a speaker expresses a plan that deviates too much from the standard plans used for the application of a given word, then that speaker's plan will not be adopted by others in the speaker's language community. Since we are not asserting truth-evaluable sentences and, instead, only expressing plans for the use of predicates, it doesn't matter that according to the hyperplan that sets a cutoff for the application of 'bald' at 10^{10} hairs everyone is bald. It is not clear that this move is available to other plurivaluationists. For Raffman, part of what defines the extension of available ranges of application is the context. Our context will not allow ranges of application that deviate too much from past usage.

An objection that does stick to MacFarlane's plurivaluationism is the possibility that complete precisifications may be impossible. To reiterate the point, precisifying a predicate requires an appeal to another predicate, e.g. 'bald' may be precisified by appeal to 'hair.' If this process of precisifying never reaches a precise predicate, then the precisification will never be completed. We could have an object that lies perfectly within the border area for each of the infinite predicates in the chain. This same problem will arise for hyperplans. Either we sometimes have predicates for which there are no hyperplans (assuming a hyperplan must eliminate indeterminate cases) or there are hyperplans that fail to eliminate all indeterminacy. If ones goal is to preserve classical logic, then relativizing truth to hyperplans that fail to eliminate all indeterminacy is not the way to go.

3.3.2 Objection: Higher-Order Vagueness

Since expressivism does not restrict truth-evaluability to acceptable hyperplans, it does not suffer from the same higher-order vagueness problems that the multi-range theory does. However, there is a serious problem for MacFarlane's view. Presumably some of our sentences are not merely expressions of plans for the use of predicates. When we assert that '5 is an odd number,' we are not merely stating a plan to use 'odd number' for numbers like 5. Instead, we are asserting that 5 has the property <not being divisible by 2>. A view that says that even a sentence like '5 is an odd number' is not truth-apt would have to hold that no atomic sentence is truth-apt. That would make for a much more extreme version of expressivism than the one that inspired MacFarlane's view. The difference between '5 is an odd number' and 'that is a big apple' comes from the fact that, by virtue of being precise, 'odd number' expresses a hyperplan, not just a plan. So, '5 is an odd number' is truth-evaluable when asserted because we can simply relativize truth to the hyperplan expressed by 'odd number.'

The act of predication is sometimes just an expression of a plan and other times it is an expression of a hyperplan. The difference between the two scenarios lies in the vagueness of the predicates. Vagueness is vague. Whether or not a predicate is vague can be indeterminate. Consider a sentence that predicates such an indeterminately vague predicate of an object. Is the assertion of this kind of sentence merely an expression of a plan or is it an expression of a hyperplan? It is indeterminate whether or not the sentence expresses a hyperplan.

3.3.3 Quick Summary

Expressivism is the view that when we assert sentences containing vague words we are actually merely expressing plans for the applications of our words. For example, when someone says 'that is a large apple' in reference to an apple that is 84mm in diameter, they are expressing a plan to apply 'large' to apples of 84 or more millimeters in diameter. MacFarlane claims that we need a hyperplan to evaluate sentences for truth-values. A hyperplan is a completely specified plan for the application of a predicate. Truth, then, is relativized to these hyperplans, along with contexts and worlds.

It is not clear why we need a hyperplan to evaluate sentences for truth. It can be clear that a plan has been followed through with even when that plan does not specify every detail. As for the relativization of truth to hyperplans, this suffers from some of the same problems that other plurivaluationist theories must contend with. If complete precisification is impossible, then some indeterminate cases will be left out. If one's goal is to preserve classical logic, then a plurivaluationist theory is not the one to choose.

Chapter 4

The Possibility of Vagueness in the World

Introduction

There are three viable options for what kind of phenomenon vagueness is. (i) It is a merely epistemic phenomenon: the world and our language are both precise, but we do not know the extensions of all of our words and concepts. Williamson's epistemicism[1] is a perfect example of this option. (ii) It is a semantic phenomenon: the world is precise, but our language can only paint a fuzzy picture of it. David Braun and Theodore Sider's form of semantic nihilism[2] is a view like this. Finally, (iii) it is a metaphysical phenomenon: the objects and properties in the world can be subject to vagueness. As for other popular theories of vagueness, contextualists[3] will most likely provide views that fall into (i) or (ii). Supervaluationists are most often in camp (ii) since the cleanest version of the view will treat the problem of vagueness as similar to ambiguity where our sentences fail to pick out unique propositions—assuming all propositions are perfectly precise.[4] Finally, proponents of a degree theory can fall into camps (ii) or (iii).[5]

My goal in this chapter is to argue that vagueness is not merely an epistemic or a linguistic phenomenon; it is also a metaphysical phenomenon. There are a number of objections to the view that there is vagueness in the world. One of the biggest problems with (iii) arises from an argument given by Gareth Evans[6] concluding

[1]See chapter 1 of this dissertation.
[2]Described in chapter 3.
[3]See chapter 2.
[4]However, Elizabeth Barnes' (Barnes, 2010) theory mimics supervaluationism and falls into camp (iii). I will describe Barnes' theory in more detail in the next chapter.
[5]These theories will be discussed in chapter 6.
[6]Evans, 1978.

that indeterminate identities, which are a consequence of metaphysical vagueness, entail a contradiction. This argument is the focus of much of this chapter. I will first give arguments for the possibility of vagueness in the world, and I will then defend metaphysical vagueness against Evans' argument.

4.1 What Does Vagueness in the World Look Like?

It is much easier to understand what is meant by "semantic indeterminacy" than to understand what is meant by "metaphysical indeterminacy." Semantic indeterminacy is present in examples like (4.1), which we've seen before.

$$\text{Winston went to the bank.} \tag{4.1}$$

The problems with evaluating a sentence like (4.1) do not come from some problem in the world. Instead, our difficulties come from a failure to specify a single meaning for "bank." Our sentence does not pick out a determinate meaning, and so we have a case of semantic indeterminacy.

What would it mean for there to be metaphysical indeterminacy, for there to be vagueness in the world? Terence Parsons and Peter Woodruff answer this question in the following way:

> The world consists of some objects, and some properties and relations, with the objects possessing (or not possessing) properties and standing in (or not standing in) relations. Call these possessings and standings-in states of affairs. Then the world determines that certain of them hold, and that certain of them do not hold, but leaves the rest undetermined.[7]

Parsons and Woodruff claim that objects and properties stand in states of affairs where an object possesses some property. There are some states of affairs that obtain, some that don't, and some that the world has left undetermined. Parsons and Woodruff describe one of the ways that there could be vagueness in the world—vagueness in property instantiation. I will describe one more that occurs as a result of vagueness in property instantiation—vagueness in existence.

[7] Parsons and Woodruff, 1995.

4.1.1 Vagueness in Property Instantiation

Objects may possess or fail to possess a property, but they may also indeterminately possess a property. That is, the state of affairs of their possessing the property is undetermined. For example, if Ivan is a borderline bald man, then his possessing the property <being bald> is indeterminate.

Related to vagueness in property instantiation, one other way that metaphysical indeterminacy is talked about is vagueness in boundary. Imagine Tibbles the cat sitting on a table. There is one hair that is barely attached to the cat. Is the hair part of Tibbles? Vagueness in boundary may be seen as a form of vagueness in property instantiation. The hair and Tibbles indeterminately stand in the part-of relation. In other words, Tibbles indeterminately instantiates the property <having that hair as a part>. R. M. Sainsbury calls this kind of vagueness "compositional vagueness."

CV x is compositionally vague $=_{df}$ x is such that, for some y, it is indeterminate whether or not y is a part of x.[8]

In addition to compositional vagueness, Sainsbury describes three other ways an object may be vague. These are modal vagueness, where it is indeterminate whether or not an object exists in a given world; temporal vagueness, where it is indeterminate whether or not an object exists at a time; and individuative vagueness, where it is indeterminate whether or not an object x is identical to some object y. I see these as species of vagueness in existence, and therefore also species of vagueness in property instantiation.

4.1.2 Vagueness in Existence

Another way that there could be vagueness in the world is for the existence or persistence conditions for an object to be vague. For example, if we remove one board at a time from the ship of Theseus, at some point the ship will cease existing. However, how many boards the ship can lose and still exist is vague, and so there can come a point where the existence of the ship is unclear. This kind of metaphysical indeterminacy goes a little further than just vagueness in property instantiation.

[8] Sainsbury, 1989, p. 101.

Vagueness in property instantiation is a matter of whether or not a given object has a property. Vagueness in existence is a matter of whether or not the object is even there to instantiate properties.

4.1.3 That These are Metaphysical

It is worth stressing that these kinds of indeterminacy are metaphysical and not just linguistic. In the case of property instantiation, the view is that there really are properties in the world that can be indeterminately instantiated by objects. The claim is not that there are a bunch of properties for baldness that have different cutoffs like <being bald$_{151}$> (baldness, but with a cutoff at 151 hairs), but rather that there is a property <being bald> that lacks a sharp cutoff.

In the case of vagueness in existence, the claim is that there can be indeterminacy in the list of existing objects. Whether or not the ship of Theseus should be included on the list of existing things after half of its boards have been removed is indeterminate.

Evidence that these kinds of vagueness are metaphysical comes from the kind of objection that would be strongest against the examples given above. If one is a sparse property theorist or an eliminativist about ordinary objects, then the natural response would be to deny that there is such a property as <being bald> or to deny that the ship of Theseus ever existed. The thought might be that only fundamental kinds of objects exist, and, since such entities are not sorites-susceptible, there is no vagueness in the world. One might think that fundamental properties are not sorites-susceptible because, by virtue of their being fundamental, there is no dimension on which to make a series of similar objects ranging from those that clearly have the property to those that clearly lack it. We need something like hairs, in the case of <being bald>, to make a sorites series. Since the response is to deny the existence of some properties or objects, it is clear that the kinds of metaphysical vagueness described above are understandable as metaphysical and not just as linguistic. The question, then, is just whether or not they are actually possible.

4.2 Argument for the Possibility of Vagueness in the World

In this section, I will give an argument that vagueness in the world, of the form just described, is possible. I restrict my argument to the mere possibility of metaphysical vagueness since this is both sufficient for the purposes of establishing the requirement that we accommodate such vagueness in our logic and easier to establish within the vagueness dialectic. If metaphysical vagueness is possible, then even if it is not actual, our logic will have to account for possible borderline cases of property instantiations. So, if the possibility of metaphysical vagueness can be established, it will undermine attempts at solving the sorites paradox that do not treat vagueness as metaphysical. Part of the task of addressing the sorites paradox is to determine whether or not classical logic can be maintained. Since modal claims can be analyzed by a logic and because some of these modal claims may involve indeterminacy, at the very least classical logic needs to be amended.

4.2.1 An Argument from Ordinary Objects

To begin, I will provide an argument aimed at those who already accept ordinary objects. Those who accept ordinary objects should accept metaphysical vagueness since the only alternative is epistemicism. Consider a baseball that has one atom at a time removed from it. After a sufficient number of atoms have been removed the baseball no longer exists. If we accept that there is a baseball at the beginning of the process and no baseball at the end, we must make a decision about the transition from baseball to non-baseball. Accepting metaphysical vagueness will allow us to treat the transition as either gradual, in the case of a continuum-valued theory, or as passing through borderline cases. If we reject metaphysical vagueness, we are left with a sharp transition from baseball to non-baseball. But, surely a single atom does not make the difference between a baseball and a non-baseball.[9]

[9]The same might be said about the transition from a determinate baseball to a borderline baseball. Even accepting metaphysical vagueness might not get us away from requiring sharp borders, just not the one between baseball and non-baseball. I will address this problem of higher-order vagueness in the next chapter.

If there are baseballs, then they are vague things. Note that the claim here is not that the word "baseball" is vague, but rather that the property <being a baseball> is vague. The sorites described is not a sorites on our descriptions of objects in the world, but a sorites on an object existing in the world. The most likely objection to this argument for the vagueness of the property <being a baseball> comes from supervaluationism. It is supertrue that there are baseballs, and similarly for other ordinary objects. However, the supervaluationist does not usually accept that baseballs are metaphysically vague.

It is not clear that the supervaluationist meaning of "x exists" entails that x is a member of the furniture of the world. That is to say, it is not clear that existence for the supervaluationist is a matter of ontological reality or merely the truth of the everyday sentence "baseballs exist." If the supervaluationist claims that baseballs are part of the furniture of the world and that the property <being a baseball> can be precisified[10] in multiple ways, then they will have accepted that baseballs are metaphysically vague. However, supervaluationists are apt to treat the actual properties in the world as precise and our language as failing to pick them out properly. For such a supervaluationist it is not clear that they accept ordinary objects like baseballs into their ontology.

The sorites on the instantiation of the property <being a baseball> assumes that there is some object that instantiates the property of being a baseball. Since the object instantiates a property and can be put into a position where it no longer determinately instantiates nor fails to instantiate that property, the property and the object are metaphysically vague. Or, if the object instantiating <being a baseball> is essentially a baseball, then its existence is vague. The kind of supervaluationist that would reject metaphysical vagueness is one who would reject ordinary objects in favor of more precise properties. The argument here is only that ordinary objects are vague objects. The next two arguments will be addressed to those who reject baseballs, and even to those who eliminate even more natural kinds like uranium

[10]Note that supervaluationists normally talk of precisifying words, not properties. What is meant here by precisifying a property is that the instantiation of the property depends on the instantiation of related precise properties. For example, Winston instantiates <being bald> iff he instantiates <baldness$_{2568}$> and <baldness$_{2569}$> and so on.

atoms.

4.2.2 The Gunky Vague World

Some philosophers hold the view that the world is gunky—that every object is complex and made up of other objects. We think that we are getting closer to the bottom level of reality, that everything is made up of quarks (let's say), and that quarks are not themselves made up of other things. However, every time we have thought we found the bottom level, we found out that our putative bottom level was itself composed of other smaller objects. After all, we use the word "atom" for something that is actually composed of objects that are composed of other objects.

Theodore Sider has argued for the possbility of gunk.

> At one point, at least, it was a legitimate scientific hypothesis that this process could go on forever, that there is no end to the world's complexity. Philosophical reflections on the nature of composition should not lead us to claim that a legitimate scientific hypothesis is metaphysically impossible. So we ought to accept the possibility of material objects made of gunk. But then, we ought to accept the possibility of the kinds of gunk worlds I've imagined.[11]

This may not be enough to overcome the strong intuition that there must be some fundamental atoms out of which everything else is built. An opponent of gunk might argue that gunk is impossible because if everything depends for its existence on something else, then nothing will exist since existence can never get off the ground. Ross Cameron argues quite convincingly that this is not a strong objection to gunk.

> Ontological priority is not temporal priority. It is not as if God has to have made a, b and c previously (literally) if he is to make the sum of a, b and c. Why can he not just make them all together? Similarly, why can he not just make all the infinitely many things that inhabit the gunky world at once, together with the relations of ontological priority that hold between them? My concern is that the intuition rests on our taking too seriously the temporal metaphor suggested by "priority."[12]

[11]Sider, 1993, p. 287.
[12]Cameron, 2007.

The existence of a complex object does not have to temporally follow the existence of its parts. It need only be the case that there are those parts for there to be the complex. If Cameron is right that the confusion of temporal priority and ontological priority is a major underlying motivation for the intuition that gunk is impossible, then we have a strong reason to reject the intuition. Without this intuition, it is hard to say that gunk is not at least possible.

Peter Unger[13] argues that every ordinary object is sorites-susceptible. In the process of removing its parts, as long as we do it slowly enough, we will end up with borderline cases.[14] While it seems possible that there could be a complex object that was not sorites-susceptible, it also seems possible that every kind of object in a gunky world could be sorites-susceptible. Imagine a world where every object is composed of other objects and that every kind of object is sorites-susceptible.

Let's use a world similar to the actual world as an example. What counts as an atom is sorites-susceptible. We just need to start breaking up the nucleus. What counts as a proton is sorites-susceptible as we slowly move the quarks apart; it is unclear when the strong force is no longer holding them together. Supposing that the example world is gunky and that quarks are composed of sub-quarks, we can do the same thing and push the sub-quarks apart and ask when the sub-strong force[15] is no longer holding them together.

The possibility of this gunky vague world would remove some of the objections to vagueness in the world. The process by which we remove vague properties and replace them with precise ones—exchanging <being bald> for <having fewer than x hairs>—fails when we never reach a bottom level. We cannot successfully precisify a property in the gunky vague world, since the properties we would use to precisify it are themselves vague and so on *ad infinitum*.

[13]Unger, 1979, pp. 120-121.

[14]There could be complex objects where their structure was such that it is impossible to remove one component without destroying the object and where removal of a component is itself not vague. Unger's point holds for every complex object I can think of in the actual world.

[15]Or whatever name best suits such a posit.

4.2.3 Indeterminacy at the Bottom

The worry that the proponent of metaphysical indeterminacy has is that vague properties do not exist and instead there are a number of precise properties like <having fewer than 10,000 hairs>. This may not be as appealing of a move with more natural properties, but that will not do much to convince those who are already willing to drop properties like <being a uranium atom>. Though the possibility of a vague gunky world would remove worries about eliminativism, a proponent of metaphysical vagueness may not want to accept such a possibility. In this subsection, I will discuss the possibility of indeterminacy at the fundamental level of reality. In this way, rejecting properties like <being a uranium atom> does not successfully avoid indeterminacy in the world. In addition, the possibility of indeterminacy at the fundamental level removes some of the complaints made against higher-level properties. When these are combined with defenses against the other objections made by eliminativists, we have more reason to accept ordinary objects and ordinary properties into our ontology.

E.J. Lowe argues that there are actual examples of indeterminacy of identity. While the example may not be one of vagueness as we generally understand it—having something to do with sorites-susceptibility—it does present the possibility of metaphysical indeterminacy.

> Suppose (to keep matters simple) that in an ionization chamber a free electron *a* is captured by a certain atom to form a negative ion which, a short time later, reverts to a neutral state by releasing an electron *b*. As I understand it, according to currently accepted quantum-mechanical principles there may simply be no objective fact of the matter as to whether or not *a* is identical with *b*.[16]

Lowe stresses that physicists do not normally interpret the indeterminacy here as merely epistemic. It is thought to be ontic indeterminacy.

> It is well known that the sort of indeterminacy presupposed by orthodox interpretations of quantum theory is more than merely epistemic - it is ontic.[17]

[16] Lowe, 1994, p. 110.
[17] Lowe, 1994, p. 110.

If the correct interpretation of quantum mechanics does involve indeterminacy, then we have a strong reason to also accept the possibility of indeterminacy caused by vagueness. Even if the correct interpretation of quantum mechanics does not involve indeterminacy, we still have some reason to accept the possibility of such indeterminacy. If the laws of quantum mechanics are not metaphysically necessary, then there could be a world that is not nomically possible but is metaphysically possible where the Copenhagen interpretation of quantum mechanics is correct and there is metaphysical indeterminacy.

This kind of metaphysical indeterminacy is not a case of vagueness in the world. Vagueness requires some dimension on which a property is sorites-susceptible. When dealing with fundamental entities there is no such dimension. Metaphysical indeterminacy is a broader genus of which metaphysical vagueness is a species.

The option to precisify a vague property will always be available (assuming we're not in the vague gunky world, we can ultimately succeed at precisifying). This is because, understood in the sorites-susceptibility sense, vague properties will always be vague on some dimension or other. As such, we could always get rid of the property and accept a number of more precise properties. However, the fact that metaphysical indeterminacy is possible supports the acceptance of properties like <being a uranium atom> or <being a cat>. That these properties allow for metaphysical indeterminacy is now a much less compelling reason to reject them, since we already have metaphysical indeterminacy. In other words, if one's goal is to avoid allowing any indeterminacy into our logic, then restricting our ontology to just fundamental objects will not help.

There may be other reasons to reject higher-level properties like <being a uranium atom> or <being a cat>. For example, there are arguments from causal redundancy that claim that ordinary objects like baseballs either lead to causal overdetermination or are causally inert—both of these options seem problematic. Amie L. Thomasson gives good defenses against claims of causal redundancy in her book *Ordinary Objects*[18]. She accepts a kind of causal overdetermination, but argues that this kind of overdetermination is not the problematic kind we are attempting to

[18]Thomasson, 2007.

avoid. She argues that we do not have problematic overdetermination as long as x's causing e analytically entails y's causing e. An example of analytic entailment would be that "Jones bought a house" analytically entails "Jones bought a building." Call the particles that compose the baseball b and the baseball B. b causes the window to break and so does B, but B's breaking the window is not independent from b's, since b's breaking the window analytically entails B's breaking the window.[19]

Adequately dealing with the eliminativist's arguments will take us too far afield. Even if they succeed, it will not get us out of the possibility of metaphysical indeterminacy. It is also not clear that they will get us out of the possibility of the gunky vague world. We will have to accept the possibility of vagueness in the world and incorporate it into our logic.

4.3 The Argument against Vague Identities

Now that we have seen reasons for accepting the possibility of metaphysical vagueness, it is important that we address the arguments that metaphysical vagueness is impossible. One consequence of metaphysical indeterminacy is the possibility for indeterminate identities. I will discuss two representative examples: the personal identity machine and quantum indeterminacy.

Suppose person P walks into a machine that will do something to them such that the person who walks out of the machine, call them S, is neither clearly a new person nor clearly the same person as P. For a simple example, if a bodily continuity theory of personal identity were correct, then the personal identity machine would change enough of P's atoms so that it is unclear whether or not S is P. It is indeterminate whether or not S is identical to P, since it is unclear whether or not P survives the personal identity machine. The vagueness of personhood entails the possibility of indeterminate identity.

As for vague identities arising out of quantum indeterminacy, consider again Lowe's example from the previous section. Electron a is taken up by an atom and electron b is released by that atom. It is indeterminate whether or not $a = b$.

[19]See Thomasson, 2007, pp. 15-17

4.3.1 The Argument

As the above examples show, for some a and some b it could be indeterminate whether or not $a = b$. Gareth Evans provides an argument that from the assumption that there is an indeterminate identity we reach a contradiction. Since the possibility of metaphysical indeterminacy appears to entail the possibility of indeterminate identities, if indeterminate identities lead to a contradiction, then so does metaphysical indeterminacy. The argument goes as follows:

(1) $\bigtriangledown(a = b)$ (Assumption for *Reductio*)[20]

(2) $\lambda x[\bigtriangledown(x = a)]b$ (From 1)[21]

(3) $\sim \bigtriangledown (a = a)$ (Premise)

(4) $\sim\lambda x[\bigtriangledown(x = a)]a$ (From 3)

(5) $\sim(a = b)$ (From 2 and 4 by Leibniz's Law)

(6) $\sim \bigtriangledown (a = b)$ (From 5)

First, we assume that "$a = b$" is indeterminate. From this it follows that b has the property <being indeterminately identical to a>. It is taken as a premise that "$a = a$" is not indeterminate, which entails that a does not have the property <being indeterminately identical to a>. But, since a and b have different properties, we can conclude via Leibniz's Law that a and b are not identical.

This does not clearly lead to a contradiction. After all, (5) simply says that a is not identical to b; it does not say that they are determinately non-identical. According to some understandings of indeterminacy, it may be the case that $\sim(a = b)$ is consistent with $\bigtriangledown(a = b)$, and so the argument would fail to establish the desired contradiction.

Brian Garrett[22] attempts to fix this problem in Evans' argument. He claims that (2) is true and that $\lambda x[\bigtriangledown(x = a)]a$ is "not just *not true* but *false*." That is, he is claiming that both (2) and (4) could have determinacy operators placed in front of them.

[20] $\bigtriangledown\phi$ means ϕ is indeterminate. $\bigtriangleup\phi$ means ϕ is determinate.
[21] b has the property <being indeterminately identical to a>.
[22] Garrett, 1991, p. 347.

Since there is a determinate difference in properties between a and b, they should be determinately distinct, and so we can derive (6). Evans' solution to the problem is to accept an axiom similar to the K axiom of modal logic, but applying to the determinacy operator, $\triangle$.

$$\phi \vdash \triangle\phi \tag{4.2}$$

If ϕ is derivable, then $\triangle\phi$ is derivable. So, once we get to (5), we can prefix a determinacy operator to it, which would then entail that it is not the case that a and b are indeterminately identical. This is because, in most logics of vagueness, $\triangle{\sim}\phi \vdash {\sim} \bigtriangledown \phi$.

4.3.2 Do We Need to Respond to Evans?

Unlike the responses to be covered in the remainder of this chapter, Michael Tye[23] accepts Evans' argument but claims that it does not disprove the possibility of metaphysical vagueness. His response goes as follows. Evans' argument demonstrates that indeterminate identities do not occur, but it does not show that vague property instantiations cannot occur. Tye denies that the possibility of vague objects entails the possibility of indeterminate identities. As I will argue for the remainder of this chapter, there is no need to take this conciliatory route since Evans' argument fails. However, Tye is incorrect in thinking that the proponent of metaphysical vagueness can simply sidestep Evans' argument.

Consider, again, the personal identity machine. Suppose that personhood is an essential property of P. P walks into the personal identity machine and S walks out of the machine. It is now indeterminate whether or not P is identical to S since, by description of the case, the personal identity machine does whatever has to be done to render the resulting person indeterminately the same as the one who enters. If the correct theory of personal identity is one of bodily identity—same body, then same person—then the personal identity machine will replace enough atoms such that it is indeterminate whether or not they have the same body after the process as before

[23] Tye, 1990.

the process. It is the vagueness of personhood that leads to the indeterminacy of the identity between P and S.

If we are accepting vague properties, then almost all properties will be vague. For any property that is a persistence condition for an object, we can create cases like the personal identity machine. As such, introducing vague properties also introduces countless possible vague identities.

R. M. Sainsbury argues that vague objects do not entail that the identity relation is vague.

> The slogan is: vague objects without vague identity.[24]

Simply because it is indeterminate whether or not an object instantiates a property, it does not follow that the property is vague. A property instantiation can be indeterminate when an object is vague and a property is indeterminate.

> For example, suppose that Snowdon is compositionally vague and that the sentence 'Snowdon has a surface area of just 1500 acres' is vague. Then the property of having a surface area of just 1500 acres neither definitely applies to Snowdon, nor definitely fails to apply to it; but, intuitively, this is a sharp property.[25]

The property in this example is supposed to be precise, but because the boundaries of the object are not precise, we can still run into indeterminacy. Sainsbury argues that the identity relation can work like this. It may be precise, even though there are vague objects. As a response to Evans' argument, this will do nothing. Responses to vague objects require only that the existence of vague objects entails the existence of indeterminate cases of identity, even if this indeterminacy is not due to vagueness in the identity relation. In other words, it could be the case that $\bigtriangledown(a = b)$ and $\bigtriangledown\lambda x(x = a)b$ even if the identity property is not the origin of any indeterminacy—the indeterminacy could originate from indeterminacy in the persistence conditions for persons.

The property <having a surface area of just 1500 acres> sets a sharp cutoff for when it applies. However, it can be vague whether or not a given object has this

[24]Sainsbury, 1989, p. 103.
[25]Sainsbury, 1989, p. 102.

property because the boundaries of objects can be vague. The property <being identical to a> also sets a sharp cutoff. Only objects that are identical to a will have this property. However, it may be indeterminate whether or not b is identical to a. Suppose a walks out of the personal identity machine after b walked in. It is indeterminate whether or not b still exists and so it is indeterminate whether or not a is identical to b. All that matters for Evans' argument is that the vagueness of objects and properties can lead to indeterminate identities.

4.4 Problems with the Evans Argument

4.4.1 Indeterminate Existence

It is worth noting that there may be cases of indeterminate identity for which Evans' argument fails. This is caused by the possibility of indeterminately existing objects. If we allow metaphysical indeterminacy in property instantiation, then, as a result of indeterminate instantiations, it is possible for objects to indeterminately exist. Consider the personal identity machine described earlier in the chapter. P walks into a machine that does whatever it has to do to make it indeterminate whether or not the person, S, who walks out of the machine is identical to P. It is therefore indeterminate whether or not P still exists after walking into the machine.

Lambda abstraction is not valid when applied to indeterminately existing objects. Only existing objects can instantiate properties. Suppose that P indeterminately exists; it can only indeterminately instantiate properties. As such, even though $\bigtriangledown(P = S)$, it does not follow that $\lambda x \bigtriangledown (x = S)P$.

This result may not be completely satisfying to the defender of metaphysical indeterminacy. It undermines Evans' argument in the case of indeterminately existing objects, but that is not enough to establish the possibility of general metaphysical indeterminacy. If the possibility of metaphysical indeterminacy implies the possibility of indeterminate identities that do not involve indeterminate existence, then Evans' argument still stands as a *reductio ad absurdum* of metaphysical indeterminacy.

For the rest of this chapter, I will assume that in Evans' argument both a and b completely determinately exist.

4.4.2 Some Remarks on Leibniz's Law

Some responses to Evans' argument claim that the use of Leibniz's Law in the argument is invalid because of the introduction of indeterminacy. I will use the name "Leibniz's Law" for the Principle of the Indiscernability of Identicals.

$$\forall x \forall y (x = y \rightarrow \forall F(Fx \rightarrow Fy)) \tag{4.3}$$

If a and b are indeterminately identical, then there must **not** be any properties that a and b completely determinately[26] disagree on. That is to say, there is no property F such that a completely determinately has F and b completely determinately lacks F, or *vice versa*.

Suppose that a and b are indeterminately identical but that they completely determinately disagree on some property, H, because a completely determinately instantiates H and b completely determinately lacks H. It follows that $\sim\forall F(Fx \rightarrow Fy)$ since $\sim(Ha \rightarrow Hb)$. By *modus tollens* on the conditional in Leibniz's Law, we get $\sim(a = b)$. There should be no problem affixing a determinacy operator in front of this negated identity. Since it is completely determinate that Ha and completely determinate that $\sim Hb$, it should be completely determinate that $\sim(Ha \rightarrow Hb)$. It follows that it is completely determinate that $\sim\forall F(Fx \rightarrow Fy)$, and therefore it is completely determinate that $\sim(a = b)$. This result, however, contradicts the assumption that a and b are indeterminately identical. To reiterate this point about Leibniz's Law, if a and b are indeterminately identical, then they do not completely determinately disagree on any property.

What happens when a and b completely determinately agree on all of their properties, including the properties <being identical to a> and <being identical to b>? We will start with the assumption that a determinately exists[27] and determinately

[26] For ϕ to be completely determinate, there is no question of its determinacy at any level of higher-order vagueness. Affixing any number of determinacy operators in front of ϕ will result in a truth.

[27] We will be assuming that a and b determinately exist for the remainder of the chapter.

has the property <being identical to a>.[28] Suppose that a and b completely determinately agree on all properties. It follows that b completely determinately has the property <being identical to a>. Well, if b completely determinately has the property <being identical to a>, then $a = b$ is completely determinately true. So, if a and b completely determinately agree on all properties, then they are not indeterminately identical.

4.4.3 The Failure of Lambda Abstraction

On the assumption that both a and b completely determinately exist, the target of my objection is the inference from (1) to (2).

We will assume that $\sim\lambda x(\triangledown(a = x))a$. As I will claim in the following sections, every object that completely determinately exists is completely determinately self-identical. Many of the other responses to Evans bear an intuitive cost by undermining this assumption. Combining $\sim\lambda x(\triangledown(a = x))a$ with the assumption that $\triangledown(a = b)$ allows us to conclude that b either lacks the property <being indeterminately identical to a> or indeterminately lacks this property, as represented in (4.4).

$$\triangle ... \triangle\sim\lambda x(\triangledown(a = x))b \vee ... \triangledown ...\sim\lambda x(\triangledown(a = x))b \tag{4.4}$$ [29]

One of the conclusions from the previous section was that indeterminately identical objects cannot completely determinately disagree with regard to some property. It is completely determinate that a is identical to a, and so it is completely determinate that a does not have the property <being indeterminately identical to a>. Since b cannot completely determinately disagree with a with regard to this property, either b completely determinately lacks the property <being indeterminately identical to a> or there is some indeterminacy regarding whether or not b has the property.

[28] We will see shortly that other responses to Evans do not accept that an object, x, must always completely determinately have the property <being identical to x>. I will argue that these responses suffer because they do not accept this point.

[29] I use '... $\triangledown$ ...' to indicate that there is some indeterminacy somewhere in the hierarchy of higher-order vagueness. For example, the sequence could be something as simple as $\triangle \triangledown \triangle$ or as complicated as $\triangle\triangle \triangledown \triangledown\triangle \triangledown \triangledown\triangledown$. The point is that whatever sequence of determinacy and indeterminacy operators is placed in front of $\sim\lambda x(\triangledown(a = x))$, it must include at least one indeterminacy operator.

If *b* completely determinately lacked the property <being indeterminately identical to *a*>, then *b* would be either completely determinately identical to *a* or completely determinately distinct from *a*, and both results are contrary to our stipulation. So, we must go with the other disjunct in (4.4)—there is some indeterminacy in whether or not *b* has the property <being indeterminately identical to *a*>. If this right disjunct is true, then there is some indeterminacy in line (2) of Evans' argument. This does not necessarily entail that (2) is first-order indeterminate, but it is not completely determinately true and this indeterminacy at some level does entail that we cannot conclude that (2) is true.

Consider an inference from a completely determinately true premise to an indeterminately determinate conclusion. To clarify, it is indeterminate whether or not the conclusion is determinately true or indeterminate. Does this particular assignment of truth-values demonstrate that the inference is invalid? If the conclusion were determinately true, then the assignment would not be a counterexample to the validity of the inference. If it were indeterminate, then the inference would not be determinately truth-preserving. Since the conclusion is indeterminately determinate, it is indeterminate whether or not the assignment presents a counterexample to the validity of the inference. At the very least, we cannot treat the conclusion as true, even when we know the premise is true. Similarly, in the case of Evans' argument, we cannot treat (2) as true, even assuming that (1) is true.

To put the argument of this section succintly, since *a* and *b* are indeterminately identical, there is no property that they completely determinately differ on. Since *a* completely determinately lacks the property <being indeterminately identical to *a*>, it follows that *b* must not completely determinately have this property. This leaves open the option that *b* indeterminately has the property <being indeterminately identical to *a*>. So, (1) is true and there is some indeterminacy in the value of (2).

4.4.4 Higher-Order Objection

Any putative difference between indeterminately identical objects a and b will be such that it is indeterminate whether or not it constitutes a difference. This applies to higher-order properties as well. For example, a completely determinately has the property <being determinately identical to a>. Does b have this property? Again, I reiterate, by Leibniz's Law and the assumption that $\bigtriangledown(a = b)$, either b has this property or there is some indeterminacy in whether or not b has it. The first disjunct is easy to rule out, and so we end up with the second. So, there is indeterminacy in whether or not b has the property <being determinately identical to a>. This is consistent with there being some indeterminacy regarding whether or not b has the property <being indeterminately identical to a>—the result from the last section. In particular, it would be indeterminate whether or not it is indeterminate whether or not b has the property <being indeterminately identical to a>.

We can go up another level. It is completely determinate that a has the property <being determinately determinately identical to a>. Suppose it were indeterminate whether or not b has this property. It is therefore indeterminately indeterminately indeterminate whether or not b has the property <being indeterminately identical to a>. Since we can continue this process infinitely, it will turn out that b indeterminately ... indeterminately has the property <being indeterminately identical to a>.

What is important here is that, for any property that is put forward as a difference between a and b, it will only ever be indeterminate whether or not it constitutes a difference. As such, we cannot get a completely determinate difference between them and so we cannot conclude that they are non-identical. Indeterminate differences do not lead to non-identity according to Leibniz's Law. If a has a property F and it is indeterminate whether or not b has F, and they are otherwise in agreement on all properties, then the consequent of Leibniz's Law is indeterminate. We cannot then use *Modus Tollens* to conclude that the antecedent is false—at most, we could conclude that the antecedent is indeterminate.

4.5 Other Responses

4.5.1 Bruce Johnsen

One of the common responses to Evans' argument is to question the use of Leibniz's Law in line (5) of the proof. Bruce Johnsen[30] argues that there are reasons for restricting Leibniz's Law in the context of arguments about vagueness. Johnsen's argument ends inconclusively, but he claims that at the very least Leibniz's Law must be altered to account for vagueness and that the altered versions that would work in Evans' argument are not clearly true. So, Johnsen concludes that vague objects appear to be logically consistent, but whether or not they are metaphysically consistent is still an open question—one that hinges on which versions of Leibniz's Law are true.

For an example of a possible breakdown of Leibniz's Law in a trivalent logic consider an object, a, which is indeterminately green. Taking one direction of Leibniz's Law we get the following formula.

$$(a = a) \rightarrow (Ga \rightarrow Ga)^{31} \tag{4.5}$$

Since a is indeterminately green, it follows that Ga is indeterminate. If we are using the strong Kleene truth-tables (see table (4.1)) for our connectives as Johnsen does, then the consequent of (4.5) is indeterminate and the antecedent is true.

TABLE 4.1: Kleene Conditional

	T	i	F
T	T	i	F
i	T	i	i
F	T	T	T

It follows that (4.5) is indeterminate, and so Leibniz's Law is not a law of a trivalent logic that uses the Kleene tables.

The amended version of Leibniz's Law that Johnsen advises the opponent of vague identities to accept is (LL_{nv}), below. Note that Johnsen uses the determinacy operator to mean "not indeterminate." According to Johnsen's use of determinacy, 'determinately ϕ' and 'determinately $\sim\phi$' can be true at the same time, assuming that ϕ is not indeterminate. Instead

[30]Johnsen, 1989.

[31]In Johnsen's paper, he uses '$\lambda x(Gx)a$' instead of 'Ga.'

of using $\triangle$ for determinacy as I do throughout this dissertation, I have used 'D' to represent Johnsen's version of the determinacy operator.

$$\forall y \forall z((D\phi y \wedge D\phi z) \rightarrow ((y = z) \rightarrow (\phi y \leftrightarrow \phi z))^{32} \qquad (\mathrm{LL}_{nv})$$

This version accommodates the move to trivalent logic and it allows Evans' inference to go through. Johnsen, however, questions whether or not this version is a law of trivalent logic.

> Nonetheless, whether VI*[33] is sound must still depend upon whether or not LL_{nv} is in fact a law. After all, it is not at all clear that Evans's supporters are entitled to invoke anything stronger than LL_{v}, which would seem to be the only version strictly guaranteed by the supposed law of two-valued logic, LL, given above. Yet it is evident that no argument analogous to Evans's can succeed by invoking merely LL_{v}.[34]

Johnsen's approach to Evans' argument is to question the use of Leibniz's Law. He specifically avoids attacking the use of lambda abstraction. My argument against Evans, given above, questions the use of lambda abstraction. I think there are good reasons to hold onto Leibniz's Law when we move to a trivalent logic. To begin, the problems Johnsen points out with the original law could equally be seen as problems with the Kleene truth-tables. $\phi \leftrightarrow \phi$ seems like it should be tautologous, even if we allow for indeterminacy.

As for LL_{nv}, it seems very much justified by the following reasoning. If a and b determinately disagree on a property—that is, if a determinately has the property and b determinately lacks the property—then it is determinate that they are non-identical. Even if we allow for the possibility of objects' indeterminately instantiating properties, indeterminacy is not playing a role in this example. According to Evans' argument, a and b determinately differ with respect to the property $\lambda x(\triangledown(x = a))$ and so it should be determinate that a and b are not identical. If we are to defeat Evans' argument, then questioning the use of Leibniz's Law does not seem like a viable option unless one is willing to accept the wholly unintuitive

[32] As before, Johnsen uses '$\lambda x(\phi x)y$' instead of 'ϕy'.
[33] Johnsen's version of Evans' argument using LL_{nv}
[34] Johnsen, 1989, p. 109.

position that objects that determinately differ with regard to a property can be indeterminately identical.

4.5.2 E.J. Lowe

E.J. Lowe's preferred objection to the Evans argument is to reject the properties described on lines (2) and (4).

> But this 'fact' and this 'property' seem to be, to say the very least, of highly dubious status and completely without empirical significance. We might agree that *a* does not possess any such property, but hold this on the grounds that there is no such property and so *a fortiori* no such property to be possessed by *b* either.[35]

This response will not be available to every proponent of metaphysical vagueness, many of whom are already open to non-fundamental properties. Since proponents of metaphysical vagueness are disposed to accept more properties than those who reject metaphysical vagueness, it seems disingenuous to reject the property <being indeterminately identical>. Presumably the proponent of metaphysical vagueness accepts properties like the property <being borderline bald>. In addition, the common complaint that identity is not a property because everything bears the property to itself and nothing bears it to anything else, does not apply to the properties in Evans' argument. After all, it is not the case that everything is indeterminately identical to itself.

Lowe recognizes the difficulty in accepting this objection, so he provides one that will be more amenable to these philosophers. In his second objection he claims that by virtue of the indeterminate identity between *a* and *b*, it follows that *a* has the property <being indeterminately identical to *b*>. This property is symmetrical to the one that *b* has. Lowe argues that because *a* and *b* are indeterminately identical, the properties $\lambda x \bigtriangledown (x = a)$[36] and $\lambda x \bigtriangledown (x = b)$[37] are indeterminately distinct. The objection is aimed at the inference from lines 3 to 4. Since *a* has the property <being indeterminately identical to *b*>, and since that property is indeterminately

[35] Lowe, 1994, p. 113.
[36] The property of being indeterminately identical to *a*.
[37] The property of being indeterminately identical to *b*.

distinct from the property <being indeterminately identical to *a*>, it follows that *a* indeterminately has the property <being indeterminately identical to *a*>. As such, the inference from (3) to (4) is invalid.

This second objection requires the proponent of metaphysical indeterminacy to accept a very counterintuitive consequence. Assuming that *a* determinately exists, it should be determinate that *a* determinately lacks the property <being indeterminately identical to itself>. This intuitive principle acts to stop Lowe's argument against the inference from (3) to (4). Since *a* does not have the property <being indeterminately identical to *a*>, it follows that $\lambda x \bigtriangledown (x = a)$ and $\lambda x \bigtriangledown (x = b)$ are not indeterminately distinct, and therefore *a* is not indeterminately identical to *b*. Starting with $\bigtriangleup{\sim}\lambda x \bigtriangledown (x = a)a$[38] as a premise along with the assumption that $\bigtriangledown(a = b)$ we reach a contradiction just as in Evans' argument, but using Lowe's objection. The question, then, is whether or not the premise is plausible. It would be a large intuitive cost to say that *a* does not determinately lack the property <being indeterminately identical to itself>.

4.5.3 Peter van Inwagen

Peter van Inwagen argues that the validity of Evans' argument hinges on the semantics we choose, and that we can offer a plausible semantics for vagueness that invalidates the step from (3) to (4).

> The friends of vague identity are not trying to prove that vague identity *is* coherent but only to undermine the appeal of Evans's argument; to do that (I should think) they need only construct a reasonably plausible semantics according to which Evans's reasoning is formally invalid.[39]

In defense of his semantics van Inwagen claims that each of the elements added to the semantics are individually plausible. I will briefly describe van Inwagen's semantics before describing why I think it fails at its goal of plausibility.

To begin, van Inwagen's semantics has three truth values—1, $\frac{1}{2}$, and 0. The definition of validity makes use of models that determine an *extension* and a *frontier*

[38] *a* determinately lacks the property <being indeterminately identical to *a*.>
[39] Van Inwagen, 1988, p. 265.

for all abstracts—predicates involving the λ operator—and truth values for all sentences. Models also assign to each constant a referent and the objects to which its referent is indeterminately identical, called its fringe-referents. The extension of an abstract is the set of objects that definitely have the property picked out by the abstract. The frontier of an abstract is the set of objects that indeterminately have that property.

There are three steps by which the extensions and frontiers of abstracts are determined by a model:

1. The extension of an identity-predicate[40] contains just the referent of its term; the frontier of an identity-predicate contains just the fringe-referents of its term.
2. The result of prefixing '$\sim$' to a predicate having extensions e and frontier f is a predicate having extension $U \setminus (e \cup f)$—where U is the universe of the model—and frontier f.
3. The result of prefixing $\bigtriangledown$ to a predicate having a frontier f is a predicate having an extension f and an empty frontier.[41]

Van Inwagen's logic of vagueness invalidates the abstraction in line (4). He creates a model that gives lines (1)-(3) in Evans' proof a value of 1, but gives (4) a value of $\frac{1}{2}$. a and b are indeterminately identical, so a indeterminately has the properties that b has. Since $\lambda x \bigtriangledown (x = a)b$, it follows that $\bigtriangledown \lambda x[\bigtriangledown(x = a)]a$. Because b has the property $<$being indeterminately identical to $a>$ and a indeterminately has the properties that b has, it is indeterminate whether or not a has the property $<$being indeterminately identical to $a>$. The truth-value according to van Inwagen's semantics for $\lambda x[\bigtriangledown(x = a)]a$ is $\frac{1}{2}$. So, line (4) is also $\frac{1}{2}$. The inference from (3) to (4) is an inference from a true sentence to an indeterminate sentence, and it is therefore invalid.

The partial model that gives us this result is $\{\{A, B\}, \{\{A, B\}\}, f\}$. The first set is the universe, containing only two objects, A and B. The second set, which contains the set $\{A, B\}$, is the set of indeterminate identities. In this case, A and B are indeterminately identical to one another. This means that the set representing the

[40] $\lambda x(x = a)$ is an example of an identity-predicate.

[41] Some notation has been changed to match the notation used in this paper. Van Inwagen, 1988, p. 262

universe actually has an indeterminate number of members, since if A is identical to B, then the universe contains only one object. Any predicate logic that is meant to handle metaphysical vagueness is going to have issues with how the universes of its models are structured. Finally, the function, f, is a reference function that maps 'a' to A and 'b' to B.

Looking at the inference from (1) to (2) again, the model assigns the value 1 to $\bigtriangledown(a = b)$, since the referents of 'a' and 'b' are in one of the pairs of indeterminate identities. The extension of $\lambda x(x = a)$ is $\{A\}$ and its frontier is $\{B\}$. Since adding the indeterminacy operator takes the frontier of a predicate and makes it the extension of the resulting predicate, $\lambda x \bigtriangledown (x = a)$ has the extension $\{B\}$. So, $\lambda x \bigtriangledown (x = a)b$ is assigned 1. The problematic inference is the one from (3) to (4). The extension of $\lambda x \bigtriangledown (x = a)$ is $\{B\}$. Since A is a fringe referent of 'b,' it follows that $\lambda x \bigtriangledown (x = a)a$ is assigned $\frac{1}{2}$ by the model. By the meaning of $\sim$, $\sim\lambda x \bigtriangledown (x = a)a$ is also assigned $\frac{1}{2}$. But $\sim \bigtriangledown (a = a)$ is assigned 1. This is because $a = a$ is assigned 1 and so $\bigtriangledown(a = a)$ is assigned 0 and $\sim \bigtriangledown (a = a)$ is assigned 1.

A	$\sim A$
1	0
$\frac{1}{2}$	$\frac{1}{2}$
0	1

Though he gives a semantics where the inference from (3) to (4) is invalid, it is not clear that his semantics has the intuitive appeal that van Inwagen thinks it does. The objection I will give to van Inwagen mirrors the one given to Lowe's argument. While we will have to make some changes to our logic to handle vagueness, our semantics should preserve instances of self-identity in cases of determinate existence. That is, if a determinately exists, then $a = a$ is true—has value 1. It therefore seems incorrect for a to indeterminately have the property <being indeterminately identical to a>. Brian Garrett makes this claim about the plausibility of a determinately not having the property <being indeterminately identical to a>. Garrett claims that $\lambda x[\bigtriangledown(a = x)]a$ "is not just *not true* but *false*."[42] To put the objection another way,

[42] Garrett, 1991, p. 347.

whatever semantics we decide on for the logic of vagueness, it should be the case that if a picks some object out from the universe, then $\models \triangle(a = a)$ and, either derivable from this or as a logical truth all its own, $\models \lambda x \triangle(x = a)a$. If $\models \lambda x \triangle(x = a)a$, then $\models \sim\lambda x \bigtriangledown (x = a)a$. If a semantics for vagueness does not validate the reasoning here, then that is a serious intuitive problem—a problem that my own solution avoids.

4.6 Concluding Remarks

Though some have argued that vagueness is merely a linguistic phenomenon, there are strong reasons to think that it is primarily a metaphysical one. The metaphysical possibility of vagueness is the possibility of there being properties that are not always determinately held or lacked by objects, or objects that do not determinately exist or determinately fail to exist. Those who accept ordinary objects, such as baseballs, into their ontology are already in a good position to accept the possibility that those objects are vague. Those who would prefer to eliminate oridinary objects, keeping instead more fundamental and more precise objects, must show that the gunky vague world is impossible. If it is possible that every object is a complex object and that there is vagueness at every level of reality, then we cannot escape metaphysical vagueness by eliminating the vague objects from our ontology. Finally, vagueness is not the only kind of indeterminacy. The possibility of the kind of quantum mechanical indeterminacy that some physicists take to be present in the world forces us to accept the possibility of metaphysical indeterminacy. If we must already accept some kind of indeterminacy into our metaphysics, then the barrier to accepting vague indeterminacy is substantially lowered.

As for Evans' much-disputed argument that metaphysical vagueness is impossible, I have argued that the move from line (1) to line (2) in Evans' argument is not valid. By virtue of the assumption in (1) that a is indeterminately identical to b, it follows from Leibniz's Law that for any property that a completely determinately instantiates, b will not completely determinately lack that property. Since a

completely determinately has the property <being identical to *a*>, *b* does not completely determinately lack this property. All of this leads to the conclusion that there is some indeterminacy in whether or not *b* has the property <being indeterminately identical to *a*>. Therefore, line (2) of Evans' argument is not completely determinately true. Lambda abstraction is not valid when applied to propositions that are not completely determinately true. The strongest argument for the impossibility of metaphysical vagueness does not go through. Therefore, there are reasons to accept the possibility of metaphysical vagueness, and the strongest reason to reject that possibility can be refuted.

I have already discussed how this possibility affects some of the theories of vagueness that preserve classical logic. Epistemicism treats vagueness as a purely epistemic phenomenon, but properties in the world need not have sharp boundaries between the cases where they are instantiated and the ones where they're not. In addition, if the world is vague and our words pick out these properties, then there is reason to think that our uses of words are not setting precise boundaries for their application. The theories that make use of precisifications or precisification-like structures such as Raffman's multi-range theory, supervaluationism, and MacFarlane's expressivism will struggle to handle possibilities like the gunky vague world where precisifications cannot be completed. In addition, if those theories are treated as claims that vagueness is a failure of our language to pick out a precise world, then they are undermined by the metaphysical possibility that vagueness is not merely linguistic. Finally, semantic nihilists are forced into dealing with the possibility that the truth-bearers may contain vagueness. Presumably there are true propositions about objects in the gunky vague world. For example, 'object *a* is heavier than object *b*' could be true when it is determinate that *b* is a proper part of *a*. These true propositions will contain some imprecise elements.

In the next chapter, I will sketch out the logic of states of affairs, an attempt at capturing the vagueness in the world. Part of that chapter and most of chapter 6 will be devoted to showing that the logic I propose is preferable to the other options that can handle metaphysical vagueness, including a straightforward trivalent logic and various degree theories.

Chapter 5

The Logic of States of Affairs

Introduction

In the previous chapter I argued that it is possible for the world to be vague and that our logic needs to be able to capture this metaphysical vagueness. If there is indeterminacy in another possible world, then our modal semantics needs to accommodate that indeterminacy. After we remove the theories that do not handle metaphysical vagueness, we are left with three big contenders. As we saw in chapter 4, Michael Tye[1] accepts metaphysical vagueness and provides a trivalent logic to describe the phenomenon. In addition, some forms of supervaluationism can accommodate metaphysical vagueness, as can degree theories. In this chapter, I will put forth a semantics for a logic that is capable of accommodating metaphysical vagueness and that does so in a better way than its competitors. I call this logic the **logic of states of affairs** since the atomic elements and the truth-bearers of the logic are states of affairs.

In section 5.1, I will describe the truth-values[2] in the logic of states of affairs and how those truth-values interact with one another. Throughout most of the chapter I will treat the logic as trivalent, but when we factor in the phenomenon of higher-order vagueness an argument could be made that there are infinitely many truth-values in the logic of states of affairs. In section 5.2, I will give definitions for the

[1] Tye, 1990.

[2] I will call these obtaining-values rather than truth-values for reasons that will be explained in 5.1. Obtaining-values are the statuses states of affairs can be in. Truth-values, I treat as the values sentences can have.

standard set of logical operators and discuss how the logic of states of affairs preserves some classical tautologies like the principle of non-contradiction. In section 5.3, I will describe how quantification works in the logic of states of affairs and handle the problem of indeterminately existing objects that emerged in chapter 4. In section 5.4, I discuss the semantics for the modal operators. In section 5.5 I describe how higher-order vagueness works in the logic of states of affairs and discuss whether or not the logic is trivalent. Finally, in section 5.6, I discuss the advantages the logic of states of affairs has over other theories that attempt to deal with vagueness in the world.

5.1 States of Affairs and Indeterminacy

To begin, I will address my use of “states of affairs," in order to be clear on what kind of entity the logic of states of affairs takes to be the atomic bearers of truth-values. In the following passage, we see Plantinga’s exposition on states of affairs.

> There are such things as states of affairs; among them we find some that obtain, or are actual, and some that do not obtain. So, for example, Kareem Abdul-Jabbar’s being more than seven feet tall is a state of affairs, as is Spiro Agnew’s being President of Yale University. Although each of these is a state of affairs, the former but not the latter obtains, or is actual. And although the latter is not actual, it is a possible state of affairs; in this regard it differs from David’s having travelled faster than the speed of light and Paul’s having squared the circle.[3]

An atomic state of affairs contains an object and a property, or an ordered n-tuple of objects and an n-place relation. For example, Kareem Abdul-Jabbar and the property <being more than seven feet tall> are the object and the property of the first of Plantinga’s examples. Similarly, Spiro Agnew and the property <being the president of Yale> are the object and property of the second example. For an example of a state of affairs that contains a relation, Patrick Stewart’s being balder than Donald Trump is such a state of affairs.

A good example of the theoretical role of states of affairs comes from probability spaces. Suppose we are going to roll a six-sided die. It is probable that the die will

[3] Plantinga, 1978, p. 44.

land on a number greater than 1. It is improbable that the die will land on 6. There are things that are probable and things that are improbable. What are these things? We can express them with gerund clauses: the die's landing on a number greater than 1 and the die's landing on 6. It is these states of affairs that are probable or improbable. Note that in this example the states of affairs described by the gerund clauses have an object, the die, and a property, <landing on a number greater than 1>.

States of affairs are not facts. Facts must be the case. There is no fact that does not match the way the world is. States of affairs, on the other hand, may or may not coincide with the way the world is. For example, the state of affairs [Jeff Bezos, <being poor>][4] does not obtain. I use states of affairs for the logic of vagueness because they have this feature of obtaining or not obtaining. One could use facts as well, where instead of talking about a state of affairs that does not obtain, we make a negative existential claim about a fact. That is, if some state of affairs, S, does not obtain, then there does not exist a corresponding fact F_S. I opt for states of affairs to avoid these complications.

5.1.1 Obtaining-Values

States of affairs are the truth-bearers of the logic of states of affairs—they are the bearers of truth-values. However, henceforth, I will not call the values that states of affairs bear "truth-values." This is because truth is a property of sentences, not of states of affairs, and the logic of states of affairs is not concerned with sentences.[5] A better term than "truth-value" would be "obtaining-value." The logic of states of affairs is a logic of the obtaining of states of affairs.

There are three obtaining-values in the logic of states of affairs. A state of affairs can obtain. A state of affairs obtains iff the state's object or objects instantiate the state's property or stand in the state's relation. A state of affairs can unobtain, which happens iff the state's object or objects do not instantiate the state's property or do

[4] I will be using '[' and ']' to indicate states of affairs, and '<' and '>' to indicate properties.

[5] I hope to discuss the connection between sentences and states of affairs in future work. In that work, the truth-values of sentences will be in some way determined by the obtaining-values of states of affairs.

not stand in the state's relation. I use the word "unobtain" in the hopes of avoiding confusion with terms like "fails to obtain" or "does not obtain."[6] Finally, as we saw in the previous chapter, it is possible for an object to indeterminately instantiate a property. Suppose that Ivan indeterminately instantiates the property <being bald>. The state of affairs [Ivan, <being bald>] is indeterminate in value. It obtains iff Ivan has the property <being bald>, but it is indeterminate whether or not Ivan instantiates this property, so it is indeterminate whether or not the state of affairs obtains.[7] The third obtaining-value in the logic of states of affairs is indeterminacy.

Note that indeterminacy as it appears in the logic of states of affairs is metaphysical indeterminacy—it's not merely epistemic indeterminacy. When it is indeterminate whether or not an object instantiates a property, it is not the case that we merely cannot know whether the object instantiates the property. Rather, there is no fact of the matter about whether or not the property is instantiated by the object. Consider again Lowe's[8] example from the previous chapter. It is not the case that the electrons are determinately identical or determinately distinct and we just lack this knowledge. There is no fact of the matter about whether or not they're identical, and our lack of knowledge is a result of this metaphysical indeterminacy.

5.1.2 Consequence

The logical consequence relation in the logic of states of affairs differs from the consequence relations in other non-classical logics.

Consequence Obtaining Condition: A state of affairs ϕ is a consequence of a set of states of affairs Γ iff ϕ obtains whenever all of the members of Γ obtain.

Consequence Unobtaining Condition: A state of affairs ϕ is not a consequence of a set of states of affairs Γ iff it is possible for all of the members of Γ to obtain when ϕ unobtains.

One result of this definition of consequence is the possibility of vague cases of consequence. Suppose there is a possible combination of obtaining-values for the

[6] These terms are ambiguous between unobtaining and merely not obtaining.
[7] Symmetrically, it is indeterminate whether or not the state of affairs unobtains.
[8] Lowe, 1994.

atomic states of affairs according to which all of the members of Γ obtain while ϕ is indeterminate, call this combination $\mathscr{C}$, and there are no other combinations on which every member of Γ obtains and ϕ unobtains. Under this supposition, it is only indeterminate whether ϕ is a consequence of Γ. If ϕ obtained according to $\mathscr{C}$, then the consequence relation would hold. If ϕ unobtained according to $\mathscr{C}$, then the consequence relation would not hold. Since it is indeterminate whether or not ϕ obtains according to $\mathscr{C}$, neither the consequence obtaining condition nor the consequence unobtaining condition are met. So, it is indeterminate whether or not the consequence relation holds.

Many other non-classical logics treat consequence as precise—not allowing for indeterminate consequence relations. One way of giving such a consequence relation is to say that consequence relations only hold when it is impossible for ϕ to have a lower truth-value than the member of Γ with the lowest truth-value. For indeterminate states of affairs there is no fact of the matter about their obtaining or not. Using this sense of indeterminacy, treating consequence as imprecise works better. When the only putative counterexamples to a consequence relation are combinations where ϕ is indeterminate and all of the members of Γ obtain there is no fact of the matter about whether or not it is impossible for ϕ to unobtain when all of the members of Γ obtain. So, there is no fact of the matter about whether or not a consequence relation obtains.[9]

5.2 Sentential Logic

5.2.1 The Atomic Elements

The atomic elements of the logic of states of affairs are first-order states of affairs—ones that do not contain states of affairs or properties of states of affairs. For example, [Uranium, <being radioactive>] and [Janice, <being rich>] are both states of affairs and neither contains a state of affairs. An example of a second-order state of affairs would be [[Jeff Bezos, <being rich>], <being a state of affairs>][10]. This state

[9] More will be said later in this chapter about how modality works in the logic of states of affairs.

[10] Jeff Bezos' being rich is a state of affairs.

of affairs obtains because the state of affairs [Jeff Bezos, <being rich>] is a state of affairs.

A state of affairs obtains when its object instantiates its property or its objects stand in its relation. It unobtains when its object does not instantiate its property or its objects do not stand in its relation. Sometimes, however, a state of affairs is indeterminate because it is indeterminate whether or not its object instantiates its property.

1. [Bill Gates, <being rich>]

2. [Patrick Stewart, <being hirsute>]

3. [The hotdog on my desk, <being a sandwich>]

Bill Gates does instantiate the property <being rich>, so the state of affairs obtains. Patrick Stewart does not instantiate the property <being hirsute>, so 2 unobtains. Finally, the hotdog on my desk indeterminately instantiates the property <being a sandwich>[11]—the state of affairs is indeterminate.

5.2.2 On Choosing Definitions

I will now go through the standard set of logical operators used in an introductory logic textbook and show how to analyze them in the logic of states of affairs. Each operator is a kind of higher-order state of affairs—a state of affairs that attributes a property or relation to states of affairs. In the logic of states of affairs, there are negative, conjunctive, and disjunctive states as well as necessary, possible, and universal states of affairs.

In the next few subsections I will provide definitions for various logical relations that can hold between states of affairs. In choosing these definitions, the goal is to capture something like the ordinary logical meanings of words like "and," "or," and "if...then...". I will discuss various options for the relation corresponding to conjunction, disjunction, and so on. They are grouped into kinds the way we may group

[11] Of course, this is controversial. Mentioning this example to my students inevitably garners a heated debate. However, it is that debate that lends credence to treating hotdogs as indeterminate sandwiches.

species into genera. Conjunction, then, is a genus that contains multiple relations as species. The result of this plurality of logical operators will be that some options will preserve various classical tautologies and others will not. Though multiple logical relations are presented and they correspond to options sometimes taken by theorists about vagueness, I take it that for each group there is one species that is closest to our ordinary meanings for the relevant connective in English. These options will be the ones chosen for the logic of states of affairs in determining how much of classical logic is preserved. An argument will be given in each case for the preferred species of each genus. The important results for the logic of states of affairs discussed in appendix B use these preferred relations and a list of them can be found in that appendix.

Most of the options I will argue for are non-truth-functional. However, in each section, I provide at least one truth-functional option. One proponent of a trivalent approach to vagueness I discussed briefly in chapter 4 is Michael Tye. Tye uses a truth-functional semantics for his trivalent theory of vagueness.[12] For his semantics he uses the Kleene truth-tables. For most of the logical operators, the truth-functional choice will correspond to these tables.

5.2.3 Tautologies and Contradictions

At points in the rest of this chapter I will discuss the preservation of classical tautologies. Here is a rundown of the definition of 'tautology' as it applies in the logic of states of affairs. A tautology in the logic of states of affairs is a state that always obtains by virtue of its logical form. A contradiction is a state of affairs that always unobtains by virtue of its logical form. What is the logical form of a state of affairs?

There are a handful of logical relations that can hold between states of affairs: the conjunction relation, the disjunction relation, the material conditional relation, and the material biconditional relation. There is also the property of negation that can be instantiated by states of affairs. Logical relations have relata that are states of

[12] Tye, 1994, p. 194.

affairs that can have as their relations logical relations with states of affairs as relata. Here is an example:

$$\phi \wedge (\psi \vee \chi) \tag{5.1}$$

(5.1) is an example of a state of affairs with the conjunction relation and with one object being a state of affairs containing the disjunction relation. The logical form of a state of affairs is given by the logical properties and relations that make up the state of affairs and its constituent states of affairs. Given a state of affairs, o, we can get another state of affairs, m, with the same logical structure by only swapping the atomic states of affairs.

A tautology, then, is a state of affairs that always obtains because any state of affairs with the same logical form also obtains. A contradiction is one that always unobtains because any state of affairs with the same logical form unobtains.

5.2.4 Negation

A negative state of affairs consists of a state of affairs and the property of unobtaining. For example,

[[Jeff Bezos, <being rich>], <unobtaining>] (5.2)

(5.2) contains the state of affairs [Jeff Bezos, <being rich>] as its object and the property <unobtaining>. The linguistic representation of this state of affairs would be “Jeff Bezos is not rich." (5.2) will obtain if [Jeff Bezos, <being rich>] instantiates the property <unobtaining>—if [Jeff Bezos, <being rich>] unobtains.

Suppose that Jeff Bezos gives away most of his assets such that he becomes a borderline rich person; then [Jeff Bezos, <being rich>] indeterminately obtains. Since it is indeterminate whether or not [Jeff Bezos, <being rich>] obtains, it is indeterminate whether or not it unobtains, and so (5.2) would be indeterminate. We can glean from the above an obtaining condition and an unobtaining condition for negative states of affairs. If a state of affairs meets neither the obtaining nor the unobtaining conditions, then it is indeterminate.

Negation Obtaining Condition #1: A negative state of affairs obtains iff its object state of affairs unobtains.

Negation Unobtaining Condition #1: A negative state of affairs unobtains iff its object state of affairs obtains.

These simple versions of the obtaining and unobtaining conditions for negation give us a truth-functional[13] negation. They correspond to the Kleene truth-tables that Tye uses. The intuitiveness of this sense of negation is captured by the above example with Jeff Bezos. If a state of affairs unobtains, its negation obtains. If it obtains, then its negation unobtains. In all other cases, the status of the negation is indeterminate.

However, there are problems with the truth-functional approach to a trivalent logic. For example, even if we treat $\phi \vee \sim\phi$ as indeterminate, it seems compelling that $\sim(\phi \vee \sim\phi)$ should unobtain. If negation is taken to mean 'unobtains' instead of 'does not obtain,' then $\sim(\phi \vee \sim\phi)$ says that $\phi \vee \sim\phi$ unobtains. However, it is impossible for excluded middle to unobtain. We will explore this feature of disjunctions in more detail below, but the only way that $\phi \vee \sim\phi$ can unobtain is for both ϕ and $\sim\phi$ to unobtain. However, it is a feature of states of affairs that they cannot both obtain and unobtain at the same time. Tye even uses the fact that excluded middle cannot be false[14] to argue that treating it as sometimes indeterminate is not too high of a theoretical cost.

> I concede that the tables would be mistaken, if they permitted $A \rightarrow A$ to be false and $A \wedge \sim A$[15] to be true. But they do no such thing. $A \rightarrow A$ is a quasi-tautology and $A \wedge \sim A$ is a quasi-contradiction. So, while the former cannot be false and the latter cannot be true, both can be indefinite. This seems to me entirely palatable.[16]

Since $\phi \vee \sim\phi$ completely determinately cannot be false, it should be the case that $\sim(\phi \vee \sim\phi)$ is completely determinately false. This result is not captured by the

[13] obtaining-functional
[14] In Tye's logic they are truth-values and not obtaining-values.
[15] Tye uses $\neg$ instead of $\sim$.
[16] Tye, 1994, p. 194.

Negation Unobtaining Condition #1. Treating negation as non-truth-functional does a better job of capturing the ordinary meaning of negation than the Kleene tables do. This is where I disagree with Tye. It is not 'entirely palatable' that $\phi \wedge \sim\phi$ should count as merely a quasi-contradiction. I will again make use of non-truth-functional semantics to preserve the status of $\phi \wedge \sim\phi$ as a full contradiction.

What would be the conditions for a non-truth-functional negation that treats $\sim(\phi \vee \sim\phi)$ as unobtaining?

Negation Obtaining Condition #2: A negative state of affairs obtains iff its object state of affairs unobtains.

Negation Unobtaining Condition #2: A negative state of affairs unobtains iff its object state of affairs obtains or it is metaphysically impossible for the object state of affairs to unobtain.

The second disjunct on the right side of the biconditional in Negation Unobtaining Condition #2 captures the case we want. Since it is metaphysically impossible for excluded middle to unobtain, the negation of an instance of excluded middle will meet the right disjunct of the unobtaining condition, and will therefore unobtain. Note the use of metaphysical possibility instead of logical possibility. I have opted for the former because evaluating logical possibility will require knowing the semantics for the logical operators, and so logical possibility can't play a role in our definitions of those logical operators. However, I take it as a fundamental feature of states of affairs that no state of affairs can both obtain and unobtain at the same time. It is a metaphysical impossibility.[17] As we will see when we get to the semantics for disjunctions, under each of the options, a negated disjunction that obtains entails the negation of both disjuncts. So, $\sim(\phi \vee \sim\phi)$ entails $\sim\phi$ and $\sim\sim\phi$. If both of these obtain, then ϕ both obtains and unobtains at the same time. This is metaphysically impossible, and so $\sim(\phi \vee \sim\phi)$ is metaphysically impossible.[18]

[17] If there is a convincing argument to the alternative, that would be acceptable. A different logic of states of affairs would then be correct, but a strong enough reason to accept that kind of logic has not arisen yet. The liar's paradox can be resolved without appeal to paradoxical truth-values.

[18] See appendix B for more on DeMorgan's Law in the logic of states of affairs.

The metaphysical impossibility I just mentioned need not be purely logical. Suppose there is a uranium atom, call it u, that is going through *alpha*-decay. Suppose also that the process of losing two protons is at a point where it is indeterminate whether or not the protons are still part of the atom. So, it is indeterminate whether or not the atom is a uranium atom or a thorium atom. The options for this atom at this point in its history include just being a uranium atom or being a thorium atom. It couldn't be a carbon atom. Consider the embedded disjunction in (5.3).

It's not the case that u is either a uranium atom or a thorium atom. (5.3)

The only options for u at the time in question are being uranium or being thorium. It is therefore a metaphysical impossibility for it to both determinately fail to be uranium and determinately fail to be thorium. No other options are open to it. So, it is determinate that 'u is either a uranium atom or a thorium atom' does not unobtain, even though it would be possible for it to be indeterminate. That is to say, $U \vee T$ could be indeterminate when $\sim(U \vee T)$ unobtains. Since it is impossible for 'u is either a uranium atom or a thorium atom' to unobtain, (5.3) unobtains according to Negation Unobtaining Condition #2.

Another option for the obtaining condition would be to add a modal claim similar to the one in Negation Unobtaining Condition #2:

Negation Obtaining Condition #3: A negative state of affairs obtains iff its object state of affairs unobtains or it is metaphysically impossible for the object state of affairs to obtain.

If, like excluded middle, we have a state of affairs, S, that is indeterminate and cannot possibly obtain, since doing so would entail that a state of affairs both obtains and unobtains at the same time, then Negation Obtaining Condition #3 would be met for $\sim S$.

In what sense are these species of the genus "negation?" The answer is that they all focus on the unobtaining of the negated state of affairs. The truth-functional sense is easy to grasp since it would be the same in classical logic. However, when we move to a non-bivalent logic, new complications emerge. We could have a state of affairs that is indeterminate, but since it couldn't possibly unobtain we could put added emphasis on the unobtaining aspect of negation. Since it is completely determinate that $\phi \vee \sim\phi$ does not unobtain, it is completely determinate that $\sim(\phi \vee \sim\phi)$ unobtains. After all, what $\sim(\phi \vee \sim\phi)$ says about the world is completely determinately not the way the world is.

We have three senses of negation here. Which of them best matches our ordinary sense of negation? This is hard to say since we tend to avoid the kind of cases where the distinction here would play a role. I claim that the non-truth-functional sets of conditions are better and that conditions #3 are the best matches for ordinary language. Using the non-truth-functional sets of conditions we can capture some of the intuitions that emerge in the supervaluationist defense of their semantics for disjunctions. Cases like the one above with the atom undergoing $alpha$-decay may give us some reason to treat a state of affairs like 'u is a uranium atom or it is not' as obtaining. However, doing so will give us a semantics for disjunctions that differs from the ordinary meaning of 'or.' If negation follows conditions like #3, then our intuitions about the uranium-thorium case are validated. The disjunction does not obtain, but the negation of it determinately unobtains. 'u is either a uranium atom or it is not' is indeterminate when u is in the border area of its transition, but 'It's not the case that u is either a uranium atom or not' unobtains.

5.2.5 Conjunction

I will begin this subsection with three species of the genus "conjunction." The first will be our truth-functional option and the others will be options for a non-truth-functional semantics. Conjunctions, just like negations, are states of affairs that have other states of affairs as their components. Unlike negations, conjunctions are binary operators, and so they contain two states of affairs and the conjunctive relation. At the core of the genus "conjunction" is the relation of obtaining at the same time. That is, generally, conjunctive relations obtain when both conjuncts obtain.

Conjunction Obtaining Condition #1: A conjunctive state of affairs obtains iff both of its object states of affairs obtain.

Conjunction Unobtaining Condition #1: A conjunctive state of affairs obtains iff at least one of its object states of affairs unobtains.

The truth-functional choice suffers from the same problems as the truth-functional approach to negation. Namely, it attributes indeterminacy to some states of affairs that are not indeterminate. For example, when ϕ is indeterminate $\phi \wedge {\sim}\phi$ should not be indeterminate. Remember that it is a property of states of affairs that they cannot both obtain and unobtain at the same time. $\phi \wedge {\sim}\phi$ attributes the relation of obtaining at the same time to ϕ and ${\sim}\phi$. However, if ${\sim}\phi$ obtains, then either ϕ unobtains or it is metaphysically impossible for ϕ to obtain. For ϕ and ${\sim}\phi$ to stand in the relation of obtaining at the same time, ϕ would have to both obtain and unobtain at the same time, or it would have to be the case that ϕ can obtain while it is also metaphysically impossible for ϕ to obtain. So, it is impossible for ϕ and ${\sim}\phi$ to stand in the conjunctive relation, and so it should be completely determinate that they do not stand in that relation; $\phi \wedge {\sim}\phi$ should unobtain. It is for this reason that we should treat the non-truth-functional alternatives, below, as more accurate representations of what we care about with conjunction.

Conjunction Obtaining Condition #2: A conjunctive state of affairs obtains iff both of its object states of affairs obtain.

Conjunction Unobtaining Condition #2: A conjunctive state of affairs obtains iff at least one of of its object states of affairs unobtains or it's metaphysically impossible for both to obtain.

Conjunction Obtaining Condition #3: A conjunctive state of affairs obtains iff both of its object states of affairs obtain or it is impossible for at least one of them to unobtain.

Conjunction Unobtaining Condition #3: A conjunctive state of affairs obtains iff at least one of its object states of affairs unobtains or it's metaphysically impossible for both to obtain.

Treating $\phi \wedge \sim\phi$ as always unobtaining is accomplished in the second pair of conditions by adding a disjunct to the unobtaining condition. If it is metaphysically impossible for both conjuncts to obtain, then the conjunction unobtains.

Consider the following example of a conjunction where the distinction between the first and second pairs of conditions comes out.

$$\text{Atom } u \text{ is a uranium atom and atom } u \text{ is a thorium atom.} \tag{5.4}$$

Suppose that something cannot be both a uranium atom and a thorium atom. Now, suppose that u is going through alpha decay and losing two protons. Let's suppose that u is in that border area where it is indeterminate whether it is a uranium atom or a thorium atom. According to conditions #1, (5.4) is indeterminate. However, one may also be tempted to say that it completely determinately unobtains. By assumption, an atom cannot be both a uranium atom and a thorium atom at the same time. So it is completely determinate that [u, <being a uranium atom>] and [u, <being a thorium atom>] do not obtain at the same time.

Both conditions #1 and conditions #2 are listed as options for the semantics for conjunction. This is because they both have as central requirements that both conjuncts obtain. As for which is a better match for our ordinary understanding of conjunction, the non-truth-functional option should win out. The non-truth-functional conditions capture the intuition that ϕ and $\sim\phi$ cannot both obtain at the same time. It would be perfectly reasonable for someone to say "I don't know whether or not ϕ, but I know that it's not the case that ϕ and $\sim\phi$."

If we want a more symmetrical set of conditions, option #3 adds a disjunct to the obtaining condition to mirror the unobtaining condition. Here we could have a conjunction with indeterminate conjuncts that still meets the obtaining condition. However, there are reasons to reject this move and stick with the asymmetric conditions. For example, $(\phi \vee \sim\phi) \wedge (\phi \vee \sim\phi)$ will obtain under conditions #3, but, as will be discussed in the next section, $\phi \vee \sim\phi$ can be indeterminate. So, $(\phi \vee \sim\phi) \wedge (\phi \vee \sim\phi) \not\models \phi \vee \sim\phi$. It seems like $\phi \wedge \phi$ should be equivalent to ϕ, and so conditions #3 give us a counterintuitive result. Yet another problem is the invalidation of simplification. I take it that a central part of the meaning of 'and' in English is the validity of simplification. With these problems for conditions #3, conditions #2 seem to be closer to the meaning of 'and.'

5.2.6 Disjunction

We will do the same thing with disjunction that we did with conjunction. To begin, the feature that ties together the species of the genus "disjunction" is a providing of options. Consider inclusive and exclusive disjunctions in classical logic. Both provide options for what can be true. Inclusive disjunctions allow for both options to be true and exclusive disjunctions require that only one of the options obtain. Both of these forms of disjunction allow for at least one of the disjuncts to unobtain/be false. That is to say, the disjuncts can't both unobtain. In conjunctions, the conjucts are not optional. Both need to obtain. Of course, I am working here with Conjunction Conditions #2 since they preserve simplification. All of the following are species of disjunctions because they present options that need not both obtain.

Disjunction Obtaining Condition #1: A disjunctive state of affairs obtains iff at least one of its object states of affairs obtains.

Disjunction Unobtaining Condition #1: A disjunctive state of affairs unobtains iff both of its object states of affairs unobtain.

Within the vagueness literature, much ink has been spilled on the law of excluded middle. The addition of a third truth-value, or in my case obtaining-value, opens up the possibility that a person need not be either bald or not bald. One could lie in between the two options. Some theories of vagueness have just let go of excluded middle. Tye's trivalent theory allows for sentences like "Winston is bald or he is not bald" to be indeterminate—though, they can never be false. In contrast, supevaluationism preserves excluded middle. It is supertrue that "Winston is bald or he is not bald" since no matter how we precisify "Winston" and "bald," the sentence will be true. One complaint against supervaluationism, discussed in chapter

3, is that it provides the wrong semantics for disjunctions. The objection is that for a disjunction to be true at least one of its disjuncts must be true. However, if Winston is a borderline bald man, then neither "Winston is bald" nor "Winston is not bald" is supertrue.

Again, our first pair of conditions are the truth-functional conditions that correspond to the Kleene semantics used by Tye. They invalidate excluded middle because a disjunction with indeterminate disjuncts is itself indeterminate. So, $\phi \vee \sim\phi$ is indeterminate when ϕ is indeterminate.

Disjunction Obtaining Condition #2: A disjunctive state of affairs obtains iff at least one of its object states of affairs obtains.

Disjunction Unobtaining Condition #2: A disjunctive state of affairs unobtains iff both of its object states of affairs unobtain or it's metaphysically impossible for at least one of its object states of affairs to obtain.

Disjunction Obtaining Condition #3: A disjunctive state of affairs obtains iff at least one of its object states of affairs obtains or it is metaphysically impossible for both to unobtain.

Disjunction Unobtaining Condition #3: A disjunctive state of affairs unobtains iff both of its object states of affairs unobtain or it's metaphysically impossible for at least one of its object states of affairs to obtain.

Moving on to the non-truth-functional alternatives, #2 is the mirror to #2 in the conjunction subsection. However, in this case the addition to the unobtaining condition does not do much to capture additional cases. The interesting results come from making an addition to the obtaining condition. According to any ordinary sense of disjunction, it is impossible for $\phi \vee \sim\phi$ to unobtain. If $\phi \vee \sim\phi$ unobtains, then both disjuncts must unobtain. However, it is impossible for both ϕ and $\sim\phi$ to unobtain. If $\sim\phi$ unobtains, then either ϕ obtains or it is metaphysically impossible for ϕ to unobtain. Under the assumption that ϕ unobtains, it cannot be the case that ϕ obtains or that it is metaphysically impossible for ϕ to unobtain.

Since it is impossible for $\phi \vee \sim\phi$ to unobtain, we may be tempted to say that it obtains. This is similar to the move made in the case of $\phi \wedge \sim\phi$, which could not obtain. This is, in part, what supervaluationists do. In the case of the logic of states of affairs, the move would again rely on the property of states of affairs that restricts them from both obtaining and unobtaining at the same time. If $\phi \vee \sim\phi$ unobtained, then, given the unobtaining conditions #2/3, both ϕ and $\sim\phi$ must unobtain. However, that would entail that ϕ obtains and unobtains at the same time or that ϕ obtains and cannot possibly obtain.

So, which of these semantics is better? The emphasis on absolutely requiring at least one disjunct to obtain before the disjunction can obtain does a better job of grasping our ordinary meaning behind "or." In everyday life, when presented with a true disjunction, it is reasonable to ask which of the disjuncts is true. If the answer is that neither of them is true, then it is also reasonable to doubt whether or not the disjunction was actually true. Supervaluationists have a different focus when evaluating disjunctions. Instead of focusing on the condition that at least one disjunct is true, they focus on the condition that it's not the case that both are false. Consider again the uranium example from earlier in this chapter. The uranium atom u cannot be both determinately a uranium atom and determinately a thorium atom. As such, the disjunction 'u is a uranium atom or a thorium atom' cannot be false—and according to the supervaluationist it is true. While it may seem intuitive that 'u is either uranium or thorium' obtains when both disjuncts are indeterminate, it is hard to say whether this intuition comes from a different meaning for disjunction or from our propensity to favor classical logic.

Ultimately, preference should be given to #1 for two reasons. First, it preserves our intuition that at least one disjunct must be true. In asserting a disjunction we are giving options, at least one of which must be chosen. Second, #1 is in better accordance with treating indeterminacy as a third value. When there are three options, the middle should not be excluded. These points work equally well against conditions #3.

5.2.7 Quick Summary

Though each of these species of the various genera are relations that can hold between states of affairs, one can still ask which ones best match our ordinary usage of logical operators. In the case of negation and conjunction, I believe that the non-truth-functional versions are a better bet. Among the non-truth-functional versions, #3 for negation is likely preferable because of the symmetry between the obtaining and unobtaining conditions, and #2 for conjunction is preferable because it preserves simplification. In the case of disjunction, the truth-functional option is preferable. What would justify this different treatment of conjunction and disjunction? The answer is, in some ways, an empirical one. How do we actually use these logical operators? However, it seems likely that we will treat sentences like "Winston is bald and he is not bald" as patently false. On the other hand, we will be more reticent in cases like "u is a uranium atom or it's a thorium atom." At first one might quickly reply that "yes, u is either a uranium atom or a thorium atom," but pull back when asked to say which one. Which obtaining and unobtaining conditions for the operators best match our English uses of the connectives will mirror the asymmetry between conjunction and disjunction in English.

What about the logical truths? Is $\phi \vee \sim\phi$ a tautology and is $\phi \wedge \sim\phi$ a contradiction? According to the logic of states of affairs that is formed from the semantics I have given preference to, excluded middle is not preserved and $\phi \wedge \sim\phi$ is a contradiction. The other combinations of logical relations will give different results. These other options will result in logical forms that are tautologous and resemble excluded middle. However, since they do not match the meanings for the operators that form sentences like 'Winston is bald or he is not' they do not count as preservations of classical tautologies.

5.2.8 Conditionals

There are a number of different species of conditionals that are already present in discussions of classical logic. Material conditionals are true when their antecedents are false or their consequents are true. A counterfactual conditional is true if its consequent is true in the closest possible world where its antecedent is true.

Much discussion has gone into the question of what kind of conditionals best represent our ordinary uses of conditional structures. Here I will deal with just the question of what semantics best fit the material conditional when we add in indeterminacy.

Before delving into the various options for material conditionals, let's consider the defining feature of the genus. Conditionals have as a requirement that whenever the antecedent is true the consequent is as well. That is to say, argument forms like *modus ponens* will be valid. Each of the following options for the semantics of material conditionals captures this essential feature while dealing with indeterminacy in different ways.

Material Conditional Obtaining Condition #1: A material conditional state of affairs obtains iff either its antecedent state of affairs unobtains or its consequent state of affairs obtains.

Material Conditional Unobtaining Condition #1: A material conditional state of affairs unobtains iff its antecedent state of affairs obtains and its consequent state of affairs unobtains.

Material Conditional Obtaining Condition #2: A material conditional state of affairs obtains iff it's not the case that its antecendent state of affairs obtains and its consequent state of affairs unobtains.

Material Conditional Unobtaining Condition #2: A material conditional state of affairs unobtains iff its antecedent state of affairs obtains and its consequent state of affairs unobtains.

Material Conditional Obtaining Condition #3: A material conditional state of affairs obtains iff either its antecedent state of affairs unobtains or its consequent state of affairs obtains or it is metaphysically impossible for the antecendent state of affairs to obtain when the consequent state of affairs unobtains.

Material Conditional Unobtaining Condition #3: A material conditional state of affairs unobtains iff its antecedent state of affairs obtains and its consequent state of affairs unobtains.

Pairs #1 and #2 are both truth-functional options. The first gives us the Kleene semantics. Though the second pair looks like it would be equivalent to the first, when we include indeterminacy, it provides different results. For example, $\phi \rightarrow \phi$ is indeterminate when ϕ is indeterminate according to pair #1. According to pair #2, $\phi \rightarrow \phi$ obtains. After all, it is determinately not the case that ϕ obtains and ϕ unobtains since it is impossible for a state of affairs to obtain and unobtain at the same time. We can achieve the same thing with pair #1 by amending the obtaining condition with a disjunction like the ones seen with previous logical operators. Since it is impossible for ϕ to have a different value than ϕ, the obtaining condition for #3 is met in the case of $\phi \rightarrow \phi$.

Which of these options best fits our ordinary material conditional? There may be no good answer to this question. If a defining feature of the material conditional is that it is true when the antecedent is not met, then it should be indeterminate when the antecedent is indeterminate and the consequent is indeterminate or unobtains. However, if the defining feature of a material conditional is that the antecedent does not obtain when the consequent unobtains, then $\phi \rightarrow \phi$ should obtain. After all, ϕ cannot obtain when ϕ unobtains. Both #2 and #3 are good candidates for our ordinary material conditional. Regardless of what kind of conditional the English one is, there is good reason to think that that kind of conditional will treat $\phi \rightarrow \phi$ as tautologous. #3, however, has an advantage over #2. Suppose we have a state of affairs $\phi \rightarrow \psi$ where both ϕ and ψ are indeterminate. It is not clear that the state of affairs should obtain unless there is some special connection between the antecedent and consequent. According to #2 this state of affairs would obtain and according to #3 it would be indeterminate.

5.2.9 Biconditionals

As in the previous subsection, I will be focused on material biconditionals here. The defining feature of biconditionals is that the two states of affairs they contain have the same obtaining-value.

Material Biconditional Obtaining Condition #1: A material biconditional state of affairs obtains iff its object states of affairs both obtain or both unobtain.

Material Biconditional Unobtaining Condition #1: A material biconditional state of affairs unobtains iff one of its object states of affairs obtains and the other unobtains.

Material Biconditional Obtaining Condition #2: A material biconditional state of affairs obtains iff its object states of affairs have the same obtaining-value.

Material Biconditional Unobtaining Condition #2: A material biconditional state of affairs unobtains iff its object states of affairs have different obtaining-values.

Material Biconditional Obtaining Condition #3: A material biconditional state of affairs obtains iff its object states of affairs both obtain or both unobtain or it's metaphysically impossible for its object states of affairs to have different obtaining-values.

Material Biconditional Unobtaining Condition #3: A material biconditional state of affairs unobtains iff one of its object states of affairs obtains and the other unobtains.

According to #1, $\phi \longleftrightarrow \phi$ is indeterminate when ϕ is indeterminate. According to #2, $\phi \longleftrightarrow \phi$ obtains. After all, ϕ has the same obtaining value as ϕ. The modal condition in #3 ensures that $\phi \longleftrightarrow \phi$ obtains since it is impossible for ϕ to obtain and unobtain at the same time.

I take it that $\phi \longleftrightarrow \phi$ should obtain. ϕ cannot have a different obtaining-value than itself, and so there will never be a situation where the two sides of $\phi \longleftrightarrow \phi$ have different obtaining-values. This rules out #1. As for #2, suppose we have a state of affairs $\phi \longleftrightarrow \psi$ where ϕ and ψ are both indeterminate. According to conditions #2, $\phi \longleftrightarrow \psi$ would obtain—both sides have the same obtaining-value. According to #3, it is indeterminate unless there is some special connection between the left and right sides—a connection like identity in the case of $\phi \longleftrightarrow \phi$.

5.2.10 The Operators in the Metalanguage

Read any text on non-classical logics and you are bound to see words like 'not,' 'and,' 'if,' etc. used throughout the text. Often it feels like the proponents of these non-classical logics are using classical logic to explicate their own logic. The abundant use of negations, biconditionals, and disjunctions in the conditions I have spent the last 18 pages listing should cause some pause. I have just argued that the standard logical operators are merely genera of which there are many species. Which species are playing a role in the conditions?

The answer is that they are the ordinary senses of the operators. I have argued that certain of the options for each kind of relation best match our ordinary senses, but there may be some added complexity that is not captured by the obtaining and unobtaining conditions that comprise the semantics of the logic of states of affairs.

Note that it is important that we use the material biconditional conditions #3 instead of #2 when defining the obtaining and unobtaining conditions for the operators. Consider disjunction obtaining condition #1. Let $\phi \vee \psi$ be a disjunction where ϕ and ψ are indeterminate. Since ϕ and ψ are indeterminate it is indeterminate whether or not at least one of $(\phi \vee \psi)$'s disjuncts obtains. If $\phi \vee \psi$ is also indeterminate, then both sides of the biconditional in disjunction obtaining condition #1 have the same obtaining-value, and so, by material biconditional obtaining condition #2, disjunction obtaining condition #1 is met. This would mean that when $\phi \vee \psi$ is indeterminate it obtains. This problem does not occur with material biconditional obtaining condition #3 since it is not met when both sides are indeterminate unless it is impossible for them to have different obtaining values, which is not the case with the obtaining conditions for the operators.

5.3 Quantificational Logic

The most interesting aspects of the logic of states of affairs appear when we consider the meanings of the quantifiers. In chapter 4, I mentioned the possibility of indeterminately existing objects. For example, when S walks out of the personal identity machine, it is indeterminate whether or not S is identical to P, the person who entered it, and so it is indeterminate whether or not P still exists. This complicates matters in the logic of states of affairs. Any atomic state of affairs that attributes a property to an indeterminately existing object will itself be indeterminate, as was discussed at the end of chapter 4.

However, indeterminately existing objects can be parts of obtaining states of affairs. For example, supposing that P indeterminately exists, the following state of affairs can still obtain: "P is bald or the Eiffel Tower is in France." The left disjunct is indeterminate since it attributes a property to an indeterminately existing object. Because the right disjunct obtains, the disjunctive relation holds between two states and so the disjunctive state of affairs obtains.

5.3.1 The Domain

The domain of quantification will be all of the objects that exist. As for objects that do not actually exist but are merely possible, I will leave them to the side for now. I will avoid discussions of quantified modal logic since they will take us too far afield. Though I am leaving out discussions of merely possible objects, I will consider objects that indeterminately exist. If an object indeterminately exists, then it is indeterminate whether or not it is a member of the domain.

5.3.2 Universal Quantification

To avoid worries about the size of states of affairs, I will treat universally quantified states of affairs as being composed of the domain of quantification, as the object, and a complex property, as the property. The complex property is best explained through an example:

$$\forall x F x \tag{5.5}$$

The object of this state of affairs is the domain of quantification, and the state of affairs says that the domain instantiates the property that all of its members are such that they are F. The state of affairs does not have another state of affairs as an object, but whether or not its property is instantiated by its object depends on a number of states of affairs. The states of affairs that determine whether or not the domain has such a property would normally be considered instantiations of the universal quantifier. The property of a universally quantified state of affairs gives us a logical form and the instantiations of that quantifier are the states of affairs with that logical form. When we move from propositional logic to predicate logic, the logical form includes more than just the logical operators, it also includes the properties. So, two states of affairs have the same logical form if they differ only in their atomic states of affairs and only with regard to the objects in those atomic states.

Universal Generalization Obtaining Condition #1: A universally quantified state of affairs obtains iff every instantiation of the quantifier obtains.

Universal Generalization Unobtaining Condition #1: A universally quantified state of affairs unobtains iff at least one instantiation of the quantifier unobtains.

$$\forall x(Fx \rightarrow Gx) \tag{5.6}$$

(5.6) is a state of affairs with the domain as its object and the property that all of the domain's members are such that if they are F then they are G. If every member of the domain is such that they are G if F, then (5.6) obtains, since the domain instantiates the property of the quantified state of affairs. If at least one member of the domain is F and not G, then the domain does not instantiate the property in (5.6) and so it unobtains. Finally, suppose that there is a member of the domain, m, that is indeterminately F and determinately not G. Suppose also that it is possible that

m is F and not G, then $Fm \rightarrow Gm$ will be indeterminate. If there are no members of the domain that are F and not G, then the (5.6) will be indeterminate.

So far, I have only considered the cases where the objects in the domain determinately exist. Objects that indeterminately exist are indeterminately members of the domain. Suppose that Rex the dog indeterminately exists.

$$\text{All dogs are golden retrievers.} \tag{5.7}$$

Suppose we are in a world where every determinately existing dog is a golden retriever, but that Rex is an indeterminately existing Aussie (Australian Shepherd the dog, not the kind of person). It is indeterminate whether or not all dogs are golden retrievers because it is indeterminate whether or not Rex exists. If Rex existed, then he would instantiate the property <being an Aussie> and he would not instantiate the property <being a golden retriever>. Rex's indeterminate status as a member of the domain makes it such that it is indeterminate whether or not there is a member of the domain that instantiates the property <being a dog> but does not instantiate the property <being a golden retriever>.

Given the modal conditions that apply to the other logical operators, it seems like there should be something special to say about universal quantification. However, much of this work will be done with the other operators. For example, $\forall x {\sim}(Fx \wedge {\sim}Fx)$ will obtain, but this is because ${\sim}(\phi \wedge {\sim}\phi)$ always obtains according to the conjunction conditions #3.

I take it that a central part of the meaning of the universal quantifier is that universal instantiation is valid. However, just like conjunctions, it seems like were it impossible for every instantiation of a universal quantifier to obtain, then it is determinately the case that that quantified state of affairs unobtains. For example, suppose that objects a and b cannot be F at the same time—it's a metaphysical impossibility. Suppose also that they are both indeterminately F. According to the universal generalization conditions #1, $\forall x Fx$ would be indeterminate in this case. However, we have a reason to think that it should unobtain. It is metaphysically impossible for every member of the domain to instantiate the property <being F>, and so it is determinate that not every object of the domain is F. We can add a modal condition to the universal generalization conditions to capture this reasoning.

Universal Generalization Obtaining Condition #2: A universally quantified state of affairs obtains iff every instantiation of the quantifier obtains.

Universal Generalization Unobtaining Condition #2: A universally quantified state of affairs unobtains iff at least one instantiation of the quantifier unobtains or it is metaphysically impossible for every instantiation to obtain.

5.3.3 Existential Quantification

Existential quantification gets a similar treatment to the one given to universal quantification. The object of an existentially quantified state of affairs is the domain of quantification. The property of such a state is a complex property that says that at least one member of the domain is such that the state of affairs of its being ϕ obtains.

$$\exists x F x \tag{5.8}$$

The above state of affairs has as its object the domain of quantification and its property is the property of having at least one member such that that member is an F.

Existential Generalization Obtaining Condition: An existentially quantified state of affairs obtains iff at least one instantiation of the quantifier obtains.

Existential Generalization Unobtaining Condition: An existentially quantified state of affairs unobtains iff every instantiation of the quantifier unobtains.

If at least one of the instantiations of (5.8) obtains, then the domain has the property contained in (5.8) and so (5.8) obtains. If none of the members of the domain are F, then the domain does not instantiate the property of (5.8), and so (5.8) unobtains. Suppose there is a member of the domain m such that m is indeterminately F and no object in the domain is determinately F, then it is indeterminate whether or not the domain instantiates the property in (5.8). So, (5.8) is indeterminate.

$$\text{There is a golden retriever.} \tag{5.9}$$

Suppose that Rex is an indeterminately existing golden retriever; he indeterminately exists, but if he determinately existed he would be a golden retriever. If Rex determinately existed, then it would not affect the status of the above existential. Now, suppose that Rex is instead an indeterminately existing golden retriever and that otherwise there are no determinate golden retrievers. If Rex determinately existed, then the existential state would obtain, and so Rex's indeterminately existing makes it such that the existential state is indeterminate.

As with universal generalizations, we can add a modal component to the obtaining and unobtaining conditions for existential generalizations. Doing so will result in an operator that diverges from our ordinary meaning of 'exists.' Below is a case where such a modal condition may matter. Again, imagine that we have a series of color patches arranged from red to orange.

$$\exists x(x \text{ is the cutoff between red and orange.}) \tag{5.10}$$

One may think that it cannot be the case that each patch is determinately not the cutoff, and so (5.10) should obtain. This is akin to the move made by supervaluationists for whom (5.10) will be supertrue. For the reasons described in chapter 3, this sense of 'exist' does not match or ordinary sense. So, adding a modal condition to capture cases like (5.10) should be avoided. Instead, (5.10) should be indeterminate because all of the instantiations will either unobtain or be indeterminate. If x is on the determinately red side, then it will not be the cutoff. If it is on the determinately orange side, then it will not be the cutoff. Finally, if it is in the border area between red and orange, then it will indeterminately be the cutoff. By the existential generalization conditions, (5.10) is indeterminate.

Similarly to the case against adding a modal condition to the disjunction conditions, avoiding doing so in the case of existential generalizations makes the logic of states of affairs fit better with ordinary language than supervaluationism.

5.3.4 Note on Indeterminate Existence

There can be states of affairs that determinately obtain, even though they contain indeterminately existing objects. Indeterminate objects introduce indeterminacy into a state of affairs when the state of affairs implies that the object instantiates a property, something that indeterminately existing objects cannot do.

Consider the state of affairs [It's not the case that Tibbles is both a cat and not a cat.] where Tibbles is an indeterminately existing object. This state of affairs determinately obtains since $\sim(\phi \wedge \sim\phi)$ always obtains in the logic of states of affairs. It obtains even when Tibbles does not determinately exist. However, a state of affairs corresponding to the lambda abstraction of the above state of affairs[19] would be indeterminate since Tibbles indeterminately exists.

5.4 Modal Logic

A certain kind of modality has already appeared in some of the semantics for the logical operators. As I mentioned in the previous section, I will not be taking a position in the ongoing debate on the correct quantified modal logic; doing so would take us too far afield. So, this section will be just a snapshot of the meanings of the modal operators $\Box$ and $\Diamond$.

5.4.1 Necessity

As with the other logical operators, necessity is a property that states of affairs can instantiate.

Necessity Obtaining Condition: A state of affairs that attributes necessity obtains iff its object state of affairs obtains in all possible worlds.

Necessity Unobtaining Condition: A state of affairs that attributes necessity unobtains iff its object state of affairs unobtains in some possible world.

[19] $\lambda x(\sim(Cat(x) \wedge \sim Cat(x))(Tibbles)$

If a state of affairs obtains in every possible world, then it is necessary. If it unobtains in some world, then it is not necessary. Finally, if it is indeterminate in some world and either obtains or is indeterminate in all others, then it is indeterminate whether or not the state of affairs is necessary.

5.4.2 Possibility

Possibility Obtaining Condition: A state of affairs that attributes possibility obtains iff its object state of affairs obtains in some possible worlds.

Possibility Unobtaining Condition: A state of affairs that attributes possibility unobtains iff its object state of affairs unobtains in all possible worlds.

If a state of affairs obtains in some possible world, then it is possible. If it unobtains in all worlds, then it is not possible. Finally, if it is indeterminate in some world and is either indeterminate or unobtains in all other possible worlds, then it is indeterminate whether or not it is possible.

5.4.3 Equivalences

In classical logic, $\Box\phi$ is equivalent to $\sim\Diamond\sim\phi$. This equivalence, like most elements of classical logic, holds with some species of the logical operators and not for others. Using the truth-functional species, the equivalence holds. Suppose $\sim\Diamond\sim\phi$, then $\Diamond\sim\phi$ unobtains. By the conditions given in the previous subsection, $\sim\phi$ unobtains in all possible worlds, and so ϕ obtains in all possible worlds. Now, suppose that ϕ is the state of affairs "It's raining or it's not raining." Suppose also that we are using disjunction conditions #3 and negation conditions #2. Once we get to $\sim\phi$ in the proof just given, $\sim\phi$ will unobtain in all possible worlds, but ϕ will not—it will be indeterminate in worlds where it is indeterminate whether or not it is raining.

5.4.4 Metaphysical Possibility

Many of the logical relations described in section 5.2 make use of metaphysical possibility. The definitions in 5.4.1 and 5.4.2 can be expanded to include metaphysically possible worlds.

Possibility Obtaining Condition: A state of affairs that attributes metaphysical possibility obtains iff its object state of affairs obtains in some metaphysically possible worlds.

Possibility Unobtaining Condition: A state of affairs that attributes metaphysical possibility unobtains iff its object state of affairs unobtains in all metaphysically possible worlds.

How we determine which worlds are metaphysically possible depends on which metaphysical principles are correct. I take it, for example, that properties are such that an object cannot both determinately have and determinately lack a property at the same time. This is the reason that it is metaphysically impossible for ϕ and $\sim\phi$ to both obtain at the same time. Another example I've used in this chapter is that an atom cannot be both a uranium atom and a thorium atom at the same time. It can be indeterminate which element a particular atom is, but it cannot be both.

As for other options for the obtaining and unobtaining conditions of modal operators, we cannot add modal conditions like we did for the other operators. After all, the conditions are already modal. It may be complained that the conditions we have implicitly appeal to conjunctions and disjunctions. That is, the object state of affairs obtains in w_1 and w_2 etc. If a conjunction is needed here, then it will have to be the truth-functional choice, #1. The same holds in the case of disjunctions implicit in the condititions for possibility.

5.5 Higher-Order Vagueness

If the property <being bald> is not precise enough to draw a sharp border between when an object has that property and when it does not, then it probably shouldn't be precise enough to draw a sharp border between when an object determinately instantiates the property and when it is indeterminate whether the object instantiates the property. We've already seen the phenomenon of higher-order vagueness a few times now. In this section, I will provide semantics for the determinacy operator ($\triangle$) and the indeterminacy operator ($\bigtriangledown$). These semantics along with a definition of consequence will act as an account of higher-order vagueness in the logic of states of affairs.

5.5.1 The Operators

$$\text{It is determinate that Jeff Bezos is rich.} \tag{5.11}$$

(5.11) is a state of affairs that has as its components the state of affairs [Jeff Bezos is rich.] and the property <being determinate>. (5.11) will obtain iff [Jeff Bezos is rich.] is determinate—that is, Jeff Bezos instantiates the property <being rich>. The determinacy operator is redundant in the logic of states of affairs. If a state of affairs is determinate, then it obtains, and vice versa.

$$\triangle\phi \dashv\vDash \phi \tag{5.12}$$

Indeterminacy is where higher-order vagueness begins to emerge.

$$\text{It is indeterminate whether or not Donald Trump is bald.} \tag{5.13}$$

(5.13) is a state of affairs that has as its components the state of affairs [Donald Trump is bald.] and the property <being indeterminate>. (5.13) will obtain iff [Donald Trump is bald.] is indeterminate.

Normally, indeterminacy is defined in terms of determinacy using the following definition.

$$\triangledown\phi \dashv\vDash {\sim}\triangle\phi \wedge {\sim}\triangle{\sim}\phi \tag{5.14}$$

If it is indeterminate that ϕ, then it is not determinate that ϕ and it is not determinate that not ϕ. This definition amounts to saying that ϕ's being indeterminate is the same as its not having the other obtaining values.

(5.12) may seem problematic, given (5.14). If we can substitute $\triangle\phi$ for ϕ on the right-hand side of (5.14), then the following equivalence should hold.

$$\triangledown\phi \dashv\vDash {\sim}\phi \wedge {\sim}{\sim}\phi \tag{5.15}$$

But the right side of (5.15) is a contradiction. Though ϕ and $\triangle\phi$ entail one another, they are not intersubstitutable.

$$\sim\triangle\phi \nvDash \sim\phi \tag{5.16}$$

Suppose ϕ is indeterminate. It follows that $\triangle\phi$ unobtains, and so $\sim\triangle\phi$ obtains. However, if ϕ is indeterminate, then $\sim\phi$ is also indeterminate. Therefore, it is only indeterminate whether or not a consequence relation holds between the states of affairs in (5.16).

That (5.12) holds in the logic of states of affairs is clear when we consider the situation when ϕ obtains. That ϕ obtains is enough to tell us that it is not indeterminate whether or not it obtains. So, it is determinate that ϕ. $\triangle\phi$ is taken simply to mean that ϕ obtains and $\triangledown\phi$ is taken to mean that ϕ is indeterminate.

5.5.2 Tripartite Division

The last subsection may appear to completely ignore the worry that started off the section. Namely, it appears to create a tripartite distinction. If we're just going to end up with two sharp borders instead of one, it seems that we should give in to Williamson's abductive argument and accept a single sharp border. However, I hold that the borders between the determinate ϕs and the indeterminate ϕs are not sharp. The state of affairs $\triangle\phi$ can be indeterminate. That is, we can have a state of affairs $\triangledown\triangle\phi$ that obtains, which would happen when there is no fact of the matter about whether or not there is a fact of the matter about the status of ϕ. Similarly, it can be indeterminate whether or not it is indeterminate whether or not it is determinate that ϕ.

Suzanne Bobzien[20] argues for two axioms for the logic of vagueness that give a different structure to higher-order vagueness than the one described in the previous paragraph.

$$\triangle\phi \vDash \triangle\triangle\phi \tag{5.17}$$

[20]Bobzien, 2015.

$$\triangledown\phi \vDash \triangledown \triangledown \phi \tag{5.18}$$

Of course, I accept (5.17)—it follows from (5.12). However, I reject (5.18). It can be determinate that a state of affairs is indeterminate. If we have a sorites series composed of people at different stages of balding, there will be ones in the middle for which we know that we do not feel comfortable asserting them to be bald or asserting them to be not bald. We can be certain of our uncertainty. Though this example makes use of our epistemic status, in the realm of states of affairs I do not see a reason why there could not be states of affairs that are determinately indeterminate. Cases like the clearly borderline bald man, let's call him Francis, could play a role in states of affairs like [Francis, <being bald>]. Such a state of affairs would be determinately indeterminate, and part of the evidence for this status is our certainty that we are uncertain about Francis' status.

If baldness does not seem like a good enough example, perhaps because the border area is too small for there to be enough cases such that the ones in the middle are determinately indeterminate, consider predicates with larger border areas. For example, the predicate "cool" when it is used to describe movies or people, not the temperature. The border area for such a predicate, regardless of the dimension we are using, should be large. Cases that are very much in the middle of that border area are ones that we can be confident are, in fact, in the border area.

Another serious problem with Bobzien's account of higher-order vagueness is that it very clearly commits one to a tripartite division. Sorites series are divided into the $\triangle ... \triangle \phi$ cases, the $\triangledown ... \triangledown \phi$ cases, and the $\triangle ... \triangle {\sim}\phi$ cases.

5.5.3 Indeterminately Indeterminate

It can be indeterminate whether or not a state of affairs is indeterminate. Consider a sorites series of groupings of sand. The indeterminately indeterminate heaps are the ones that are in the border area between the indeterminate heaps and the determinate heaps. The indeterminately indeterminate non-heaps are in the border area between the indeterminate non-heaps and the determinate non-heaps. Similarly, it can be determinate that a state of affairs is indeterminate. In our sorites series of groupings of sand, these are the ones in the middle of the border area between the heaps and the non-heaps.

What exactly does it mean for it to be indeterminate whether or not a state of affairs is indeterminate? Just like atomic states of affairs, states of affairs containing logical components like $\bigtriangledown$ are states of affairs. States of affairs can be indeterminate in status between obtaining and unobtaining; this stays the same when we introduce determinacy and indeterminacy operators. It might be indeterminate whether or not an object instantiates a given property. It might also be indeterminate whether or not it is indeterminate whether or not an object instantiates a property. In such cases, it is indeterminate whether or not it is indeterminate whether or not the state of affairs obtains.

For an analogy to epistemic indeterminacy—ignorance related to the border areas of applicability of vague predicates—it might be unclear whether or not someone is bald. It might also be unclear whether or not it is unclear whether or not someone is bald. Remember, though, that this is just an analogy. The indeterminacy present in the logic of states of affairs is metaphysical, not epistemic. However, epistemic indeterminacy can piggyback on metaphysical indeterminacy. So, when we are unclear about whether or not something is unclear, it may be because it is generally indeterminate whether or not something is indeterminate. That is to say, our ignorance of our ignorance may not be a result of something like margins for error as Williamson would argue, but rather this ignorance is a result of a state of affairs being indeterminately indeterminate.

With regard to the logical relations, higher-order vagueness can affect whether or not a complex state of affairs is indeterminate or indeterminately indeterminate.

$$\phi \wedge \psi \tag{5.19}$$

Suppose that ϕ determinately obtains and it is indeterminate whether or not ψ is indeterminate. In order for ϕ and ψ to stand in the conjunction relation they must both obtain. If ψ were indeterminate, then the conjunction would be indeterminate, since it would be indeterminate whether or not they both obtain. If ψ were determinate, then the conjunction would be determinate. Since it is indeterminate whether or not ψ is indeterminate, it is indeterminate whether or not the conjunction is indeterminate.

5.5.4 Is this Really Trivalent?

No, the logic of states of affairs is not really trivalent. It effectively has an infinite number of obtaining values. After all, there are an infinite number of ways the world could fail to determine the status of a state of affairs. It could fail to determine whether or not it is has determined a state of affairs. It could determine that it has failed to determine that it has determined that ϕ.

Though there are an infinite number of obtaining-values, the logic of states of affairs is very different from a degree-theoretic approach to vagueness. In the next chapter, I will argue that these differences make the logic of states of affairs the preferable choice.

5.6 Advantages

I argued at the end of chapter 4 that theories of vagueness that treat vagueness as a purely linguistic or epistemic phenomenon will struggle to handle metaphysical vagueness. Supervaluationism can be amended to handle vagueness in the world, so it, along with degree theories, stand as the primary competitors to the logic of states of affairs.[21] For the remainder of this chapter, I will describe the advantages of the logic of states of affairs over supervaluationism. We have already seen that the logic of states of affairs has advantages with regard to the semantics for disjunctions and for existential generalizations. There are also advantages in the realm of higher-order vagueness.

[21] See Elizabeth Barnes' view, described briefly at the end of chapter 4. Barnes, 2010

Supervaluationism has a problem with higher-order vagueness. In order to get a truth-value for a sentence, we must check its truth-value on every admissible precisification. However, what counts as an admissible precisification is vague. It appears, then, that we need to precisify "admissible precisification." However, doing so will still use "admissible precisification." We start off needing to know what an admissible precisification of 'bald' is, but then we need to know what an admissible precisifiction of "admissible precisification of 'bald'" is. We never get a truth-value since we can never completely determine the set of precisifications on which to evaluate the sentence. This problem is exacerbated with supervaluationist views that treat vagueness as not just a linguistic phenomenon but also a metaphysical one.

Elizabeth Barnes proposes to treat determinacy and indeterminacy operators like modal operators, and precisifications as worlds. The determinacy operator, then, ranges over a set of worlds in the same way that the necessity operator does.

> For the purposes here I take precisifications to be possible worlds—not just *like* possible worlds, they *are* possible worlds.[22]

> My proposal is this: that every possible world is fully precise, but that if there is ontological indeterminacy it is indeterminate which of the possible worlds is the actualized world—that is, it is indeterminate which world, out of the many worlds that represent things to be a precise way, is the one that represents the way the actual world is.[23]

Unless the properties that Barnes' view of vagueness is meant to deal with are unlike our everyday properties, then it should be indeterminate which worlds we should include in the set of possible precise worlds. That is, if we have any sense of an admissible precisification in this account, then it will be vague which worlds are captured by the determinacy and indeterminacy operators.

Higher-order vagueness does not pose this problem for the logic of states of affairs. Since indeterminacy and indeterminate indeterminacy and indeterminate determinate indeterminacy are akin to obtaining-values, the vagueness of vagueness does not present a problem with evaluating states of affairs.

[22] Barnes, 2010, p. 613.
[23] Barnes, 2010, p. 613.

The logic of states of affairs has a clear advantage over a straightforward trivalent approach like Tye's. By avoiding truth-functionality, we can preserve some classical tautologies. Tye gives a truth-functional semantics for his trivalent logic, where indeterminacy is a truth-value gap. According to his semantics, an interpretation on which every atomic sentence is assigned indeterminacy, every sentence will receive a valuation of indeterminate, including excluded middle and non-contradiction. Tye defends this result by claiming that at least the classical tautologies can never be false,[24] but we can allow for some tautologies by moving away from truth-functionality.

As for the added complexity to the theory, in chapter 1, I discussed Williamson's argument that higher-order vagueness greatly increases the complexity of view that reject bivalence. The logic of states of affairs may seem incredibly complex at this point. The logic of states of affairs is meant to give an account of indeterminacy in the world, the possibility of which was argued for in chapter 4. Simplicity is a theoretical virtue, but it does not always outweigh considerations of explanatory power. Amongst the theories that account for metaphysical vagueness, simplicity will be a more valuable guide than it is between, say, epistemicism and the logic of states of affairs. Ultimately, the argument in chapter 1 is that epistemicism is not as simple as it might first appear. This is because of the additional linguistic claims that need to be made about how our uses of words determine precise boundaries for every vague word. In addition, not all of the complexity of the logic of states of affairs is unique to it. In the next chapter we will see Dorothy Edgington's arguments that we should adopt non-truth-functional semantics for a degree theory. Here arguments will apply equally well to classical logic. Epistemicism's advantages in simplicity are not enough to outweigh its disadvantages with regard to explanatory power and fit—its simplicity is not a strong enough reason to reject the arguments for vagueness in the world.

[24]Tye, 1994, p. 194.

Chapter 6

Degree Theories

We have seen in chapter 4 that it is possible for there to be vagueness in the world. In chapter 5, I provided the framework for a logic of states of affairs that I argue approximates the correct logic for the relations between states of affairs. In doing so, it handles the possibility of states of affairs that are indeterminate between obtaining and unobtaining. In this final chapter, I will consider the main competition for the logic of states of affairs in dealing with vagueness in the world. The final question I will consider in this dissertation is whether or not states of affairs obtain to varying degrees.

The concepts one is most likely to see playing a role in a sorites argument are gradable predicates like "rich," "bald," or "small." Someone can be very rich or somewhat rich or not rich at all—Jeff Bezos is extremely rich and Bernie Sanders is just somewhat rich. Similarly Patrick Stewart is very bald and David Letterman is somewhat bald. This fact about sorites series prompts some to treat truth as gradable as well. A sentence can be completely true or somewhat true, or mostly true—and anything in between. This gradability of truth is often treated by associating the degrees of truth of sentences with real numbers between 1 and 0. For example, supposing that truth to degree 1 is complete truth and truth to degree 0 is complete falsity, "Letterman is bald" might be true to degree .5 while "Patrick Stewart is bald" might be true to degree .874. In this chapter, I first describe various degree theories along with specific objections to those theories. I then move on to general objections to degree theories. Finally, I will return to the question of whether or not states of affairs can obtain in degrees and argue that the logic of states of affairs is a better account of vagueness than a degree theory.

6.1 An Argument for Degrees

The primary reason offered in favor of degree theories is the gradability of vague predicates. Since one can be bald to varying degrees, it makes sense that the sentence "Ivan is bald" can be true to different degrees. There is, however, no entailment from the gradability of predicates to the gradability of truth. We could have a gradable predicate that, nonetheless, has a sharp cutoff for its applicability. However, the gradability of the predicates does lend itself well to a gradability of truth. Nicholas J. J. Smith[1] provides an argument that the best definition of vagueness involves a principle called CLOSENESS, which resembles the TOLERANCE principle below. CLOSENESS hinges on there being multiple degrees of truth, and so if vagueness requires CLOSENESS, then there will be a reason to accept the gradability of truth.

TOLERANCE: If *a* and *b* are very close in F-relevant respects, then *Fa* and *Fb* have the same truth value.

TOLERANCE holds that small changes will not change the applicability of a predicate. It is generally accepted that we feel an intuitive pull towards TOLERANCE with regard to vague terms. Of course, most theories of vagueness are going to say that TOLERANCE is false. If one accepts TOLERANCE, then one is going to have to follow through with the sorites argument and accept that the conclusion of such an argument is true. After all, since each member, *a*, of a sorites series is very close in *F*-relevant respects to its predecessor, *b*, then *Fa* and *Fb* must have the same truth-value. This should chain together all the way to the conclusion. People who actually follow through with sorites arguments, like Peter Unger, as we saw in chapter 3, may not need to reject TOLERANCE, but they do end up rejecting things like ordinary objects. Those who do reject TOLERANCE often provide explanations for why we think that TOLERANCE is true. Williamson[2], as we saw in chapter 1, rejects TOLERANCE on the grounds that *a* and *b* can be arbitrarily close and still

[1]Smith, 2008.
[2]Williamson, 2002b.

fall on either side of the cutoff for the application of F. Williamson's answer for why we think that TOLERANCE is true comes down to our inability to know where that cutoff is, and combined with this epistemic inability is an inability to imagine a transition from a predicate applying to it not applying along a sorites series.[3]

Smith argues that the kinds of examples that draw us towards believing TOLERANCE provide equal support for a principle he claims is actually true, CLOSENESS.

CLOSENESS: If a and b are very close in F-relevant respects, then Fa and Fb are very close in respect of truth.

> Closeness says that very small differences between objects in F-relevant respects never make a big difference to the application of the predicate F —but (in contrast to tolerance) they may make a very small difference.[4]

We will consider Smith's argument with regard to the predicate "heap." He begins by citing Crispin Wright's argument that TOLERANCE is a part of the meaning of the predicate.

> Our conception of the conditions which justify calling something a heap of sand is such that the justice of the description will be unaffected by any change which cannot be detected by *casual* observation.[5]

Generally, a change in one grain of sand will not be casually observable. As such, the concept "heap" will not be sensitive enough to apply to a pile with n grains and not to a pile with $n-1$ grains. In terms of our use of the predicate "heap," TOLERANCE seems to be part of the meaning. Obviously, if we actually accept TOLERANCE as part of the meaning of "heap," then the concept will be paradoxical. Smith argues that defining vagueness in terms of TOLERANCE amounts to including as a condition for the competent use of a vague word a disposition to accept TOLERANCE with regard to that word. This is a position accepted by Matti Eklund.

> I have argued that what distinguishes vagueness is that semantic competence with a vague predicate partly consists in a disposition to accept

[3] Williamson, 1997.
[4] Smith, 2008, p. 146.
[5] Wright, 1975, p. 335.

> that this predicate is tolerant (a disposition that can be overridden, for example when it is learned that tolerance principles can never be satisfied). In slogan form, the view is that the meanings of vague expressions are inconsistent.[6]

A person who only uses vague words as if they had sharp borders might be deemed a less competent user of those words than someone who acted as if the words were tolerant. Consider the incredulous stare levelled against epistemicism. Epistemicists claim that there is a last noonish nanosecond, but many people have a intuitive reaction that drawing such a sharp border doesn't capture the inherent vagueness of "noonish." The disposition Eklund is talking about is somewhat finkish. If one actively attempts to back up their acceptance of TOLERANCE with regard to some vague word, then they will temporarily lose the disposition to accept TOLERANCE. However, when they leave the philosophy room and go back to talking about their bald uncle at Thanksgiving, they will return to the disposition to treat "bald" as if it were tolerant.

Smith argues that a small change in the applicability of "heap" would be acceptable, even if the difference between two piles is not casually observable.

> If 'heap' is a predicate of casual observation, then certainly there cannot be a difference of just one grain between a thing to which 'heap' clearly applies and a thing to which 'heap' clearly does not apply, for such a difference would not be noticeable to casual observers. It could be the case, however, that a negligible or insignificant difference (say, of one grain) between two objects makes a negligible or insignificant difference (i.e. one which we are entitled to ignore for all practical purposes) to the applicability of the word 'heap', and also that many insignificant differences add up to a significant one: this does not conflict with casual observationality, because many insignificant differences put together are noticeable to casual observers.[7]

Throughout this dissertation I have been defining vagueness in terms of tolerance. A predicate is vague iff it is sorites-susceptible, and to be sorites-susceptible, a predicate must appear to adhere to TOLERANCE. Smith claims that we should instead define vagueness in terms of CLOSENESS. Of course, there is some advantage in doing so. If our vague words adhere to CLOSENESS, then they do not include a

[6] Eklund, 2005, p. 41.
[7] Smith, 2008, p. 161.

clearly paradoxical condition. Smith's definition of vagueness in terms of CLOSENESS makes it such that having a disposition to accept an inconsistency is not part of what it means to be a competent user of most of our words. There is certainly an advantage here over Eklund's understanding of vagueness. CLOSENESS requires truth-values that are very close—close enough to capture the closeness between two objects in F-relevant respects. Degrees of truth fill this role. If true and false were the only truth-values, then CLOSENESS would either be false since there would be a transition from true to false along a sorites series, or it would be true but only because every member along a sorites series gets the same evaluation. That is to say, if true and false are the only values, then CLOSENESS will have the same problems as TOLERANCE. So, if we remain bivalent and think that CLOSENESS is a central part of the meaning of vague words, then we will still be requiring one to be disposed to accept an inconsistency to be a competent user of a vague word.

It is hard to determine what evidence about competent language use would adjudicate between TOLERANCE and CLOSENESS. We sometimes speak as if truth comes in degrees and sometimes we don't. Empirical research on the topic may not give us much direction regarding which of these principles is part of the competent use of vague predicates. The strongest point in favor of CLOSENESS is that, when coupled with a degree theory, it does not provide a clearly paradoxical condition on competent language use—we are not required to accept a paradox in order to be competent users of vague words. However, most people find sorites arguments paradoxical. This is at least some indication that they are not treating truth as coming in degrees. Treating CLOSENESS as a part of the meanings of vague words only resolves the problems with TOLERANCE when it is coupled with a degree theory. Since people do not always treat truth as coming in degrees, it is not clear that having CLOSENESS as part of the meanings of vague words is better than having TOLERANCE as part of the meanings.

Regardless of what the conditions are for competent use of vague predicates, with regard to properties, the logic of states of affairs can accept CLOSENESS. By virtue of the infinite obtaining-values that emerge from higher-order vagueness, if a and b are similar in F-relevant ways, then Fa and Fb will have similar higher-order vagueness statuses. For example, it might be completely determinate that a is F and indeterminately determinately ... determinate that b is F. Some higher-order vagueness statuses are closer to one another than others. Someone that is completely determinately bald is closer in status to someone that is indeterminately determinately bald than they are to someone who is indeterminately determinately not bald. For a more complete picture of the structure of higher-order vagueness, reference the chart in appendix A. This shows that a degree theory is not the only way to allow for CLOSENESS. It could be the case that CLOSENESS is part of the meanings of vague words because those words connect to states of affairs in the world and a competent use of those words would require us to treat sufficiently similar cases as close in higher-order vagueness status. So, a complete logic of states of affairs that connects with our language would also be able to avoid treating competent speakers as accepting a paradoxical principle.

One final point on CLOSENESS: the argument that it is part of the meaning of vague words instead of TOLERANCE may be explained away by epistemicists. The fact that we find CLOSENESS intuitive is a result of our inability to imagine a sharp transition in the case of a vague word. Our epistemic failures explain why we find both TOLERANCE and CLOSENESS are true, even though neither is. We feel that competent use of vague words incorporates TOLERANCE or CLOSENESS because of our lack of imagination.

6.2 Degree-Functional Semantics

In this section, I will describe one of the two main types of semantics for a degree theory. This first approach treats the logical connectives as degree-functional; the degree of truth of the complex formula is a function of the degrees of truth of its component formulae. Below are the functions used by the Łukasiewicz semantics for a degree theory.[8]

- $V(\sim\phi) = 1 - V(\phi)$
- $V(\phi \wedge \psi) = min(V(\phi), V(\psi))$

[8] Lukasiewicz, 1930.

- $V(\phi \vee \psi) = max(V(\phi), V(\psi))$
- $V(\phi \rightarrow \psi) = \begin{cases} 1, & \text{if } V(\phi) \leq V(\psi) \\ 1 - (V(\phi) - V(\psi)), & \text{otherwise} \end{cases}$

A disjunction can't have a value lower than its highest-valued disjunct. Similarly, a conjunction can't have a value greater than its lowest-valued conjunct. Both of these constraints match the classical understanding of disjunctions and conjunctions, so the degree-functional semantics given here do capture some elements of our intuitive understanding of the connectives.

6.2.1 Some Approaches to Validity

There are a number of ways to handle validity in a degree-functional degree theory. Here I will describe two and in subsection 6.2.3 I will describe another definition, given by Nicholas J.J. Smith.

Validity as 1-Preservation: An argument is valid iff it is impossible for the premises to all be true to degree 1 while the conclusion is true to a degree less than 1.

Validity as Degree-Preservation: An argument is valid iff it is impossible for the conclusion to have a degree of truth less than the least true premise.

I will not go through all of the classical inferences that these definitions invalidate. For now, we will just consider disjunctive syllogism.

P1. $\phi \vee \psi$

P2. $\sim\phi$

C. ψ

Let $\mathscr{I}$ be an interpretation such that $V_{\mathscr{I}}(\phi) = .5$ and $V_{\mathscr{I}}(\psi) = 0$. According to the semantics given earlier, $V_{\mathscr{I}}(\phi \vee \psi) = max(V_{\mathscr{I}}(\phi), V_{\mathscr{I}}(\psi)) = .5$ and $V_{\mathscr{I}}(\sim\phi) = .5$. So, the premises each have a value of .5 and the conclusion has a value of 0. So, according to the degree-preservation definition of validity, disjunctive syllogism is invalid.

Disjunctive syllogism is valid according to the 1-preservation definition, but classical equivalences like $\phi \rightarrow \psi \vDash \sim\phi \vee \psi$ fail. We can give a conditional a value of 1 even though neither the negation of the antecedent nor the consequent has a value of 1. But, if neither disjunct of a disjunction has a value of 1, then according to the semantics above, the disjunction cannot have a value of 1. Suppose that for some interpretation, $\mathscr{I}$, $V_{\mathscr{I}}(\phi) = .4$ and $V_{\mathscr{I}}(\psi) = .5$. It follows from the above functions that $V_{\mathscr{I}}(\phi \rightarrow \psi) = 1$ and $V_{\mathscr{I}}(\sim\phi \vee \psi) = .6$. This classical equivalence is also invalid according to validity as degree-preservation using the same interpretation results in a conclusion with a degree of truth lower than the degree of truth of the least true premise.

Smith wants to preserve the equivalance between $\phi \rightarrow \psi$ and $\sim\phi \vee \psi$, and so he rejects the Łukasiewicz conditional for one that is defined as $\sim\phi \vee \psi$.

> The Łukasiewicz conditional lacks both these properties (relative to the definition of consequence that I proposed above).[9] I therefore propose simply to define the truth conditions for the conditional so that $\mathscr{A} \rightarrow \mathscr{B}$ is equivalent to $\sim\mathscr{A} \vee \mathscr{B}$[10] and $\sim(\mathscr{A} \wedge \sim\mathscr{B})$ (the latter two are already equivalent).[11]

There are some worries with defining $\phi \rightarrow \psi$ as $\sim\phi \vee \psi$. For example, consider the sentence $\phi \rightarrow \phi$. One advantage of the Łukasiewicz conditional is that $\phi \rightarrow \phi$ will always be true to degree 1, thereby preserving one of the most intuitive classical tautologies. After all, the consequent of this conditional can never have a value lower than the antecedent. However, according to a semantics that defines the conditional in terms of a disjunction and where a disjunction is defined as it is above, the value of $\phi \rightarrow \phi$ is equal to the value of $\sim\phi \vee \phi$, which is equal to $max(\sim\phi, \phi)$. Unlike the Łukasiewicz conditional, $max(\sim\phi, \phi)$ need not equal 1, though it must be greater than or equal to .5. The fact that $\phi \rightarrow \phi$ is not always true to degree 1 according to Smith's degree theory does not mean that he excludes it from the list of tautologies. As we will see in 6.2.3, his definition of "tautology" will include $\phi \rightarrow \phi$, even though it is not constrained to having a degree of truth of 1.

[9] This definition is to be discussed momentarily.
[10] Smith uses $\neg$ instead of $\sim$.
[11] Smith, 2008, p. 268.

6.2.2 Solving the Sorites Paradox

With these definitions of validity in hand, we will now examine how a degree theorist might solve the sorites paradox.

Base Premise: 0 grains of sand does not make a heap.

Inductive Premise: For all n, if n grains do not make a heap, then neither do $n + 1$ grains.

Conclusion: 10^{10} grains of sand do not make a heap.

Degree theorists offer different resolutions to the sorites paradox. The way that a degree theorist solves the paradox hinges on the semantics they choose for the basic connectives. To begin, consider the Łukasiewicz semantics for the conditional.

$$V(\phi \to \psi) = \begin{cases} 1, & \text{if } V(\phi) \leq V(\psi) \\ 1 - (V(\phi) - V(\psi)), & \text{otherwise} \end{cases}$$

A conditional is true to degree 1 if the degree of truth of the antecedent is less than or equal to the degree of truth of the consequent. Otherwise, its degree of truth is the difference in degrees of truth between the antecedent and consequent subtracted from 1.

Many of the instantiations of the inductive premise in the first argument will be true to degree 1. These include the ones where the antecedent and consequent are true to degree 1 and those where the antecedent and consequent are true to degree 0. However, many of the instances in the middle will be true to a degree very close to 1. This is because the antecedent will be true to a slightly higher degree than the consequent. Any acceptable semantics for the universal quantifier will result in the inductive premise being true to a degree less than 1. It could be true to the same degree as its lowest valued instantiation or it could be true to a degree close to zero because of the large number of instantiations with values less than 1. The base premise is true to degree 1, and the conclusion is true to degree 0. Without going into a discussion of validity in a degree theory, we already have a response to the sorites paradox. The inductive premise is not true to degree 1, and so it is no wonder that the conclusion is not true to degree 1. Essentially, the argument is unsound.

According to validity as degree-preservation, the sorites argument is invalid. The conclusion is true to degree 0, but the inductive premise is true to a degree less than 1 but greater than 0. As such, the conclusion is true to a degree less than the lowest-valued premise.

According to validity as 1-preservation, the argument is valid. If we assume that both premises are true to degree 1, then it follows that the conclusion must also be true to degree 1. This is because each of the instantiations of the inductive premise must be true to degree 1 for the premise to be true to degree 1, and *Modus Ponens* is valid according to validity as 1-preservation. Though it is valid, the inductive premise is not true to degree 1 and so the argument is unsound.

One intuitive way the degree theorist can explain our propensity to think that the inductive premise is true is that it is very nearly true—it is true to a degree very close to 1. If the semantics for a universally quantified sentence are such that the value of the sentence equals the value of its lowest valued instantiation, then the inductive premise of the sorites argument should be close to 1, since the drops in truth between any antecedent and consequent in the series will be miniscule.

This explanation will not work in the case of other formulations of sorites arguments.

Base Premise: 0 grains of sand does not make a heap.

Inductive Premise: For all n, it is not the case that both n grains of sand do not make a heap and $n+1$ grains of sand do.

Conclusion: 10^{10} grains of sand do not make a heap.

Suppose that, for some value of n, 'n grains of sand do not make a heap' is true to degree .5 and '$n+1$ grains of sand do not make a heap' is true to degree .49. It follows from the semantics above that 'n grains of sand do not make a heap but $n+1$ grains do' is true to degree .49, and its negation is true to degree .51. The universally quantified sentence is still not true to degree 1, so the argument is still unsound. However, the inductive premise is no longer true to a degree close to 1, so we lose out on the good explanation for the intuitiveness of the premise. If one accepts Smith's semantics where $\phi \rightarrow \psi$ is equivalent to $\sim\phi \vee \psi$, then this objection will bear the same weight for the conditional version of the inductive premise.

6.2.3 Smith's Approach to Validity

Nicholas J.J. Smith provides a different definition of validity for his degree theory than the two we've seen so far. His definition, unlike the other two, successfully preserves classically valid inferences.

Validity as .5 Preservation: An argument is valid iff it is impossible for the conclusion to have a degree of truth less than .5 while all of the premises have values greater than .5.

> B is a consequence of Γ just in case on any interpretation on which the value assigned to every A in Γ is strictly greater than 0.5, the value assigned to B is greater than or equal to 0.5.[12]

> A wff is a tautology iff it is impossible for the wff to have a degree of truth less than .5.

Consider the disjunctive syllogism in 6.2.1 again. Now, the particular interpretation used to show that the argument was invalid fails since it is not an example where all of the premises have values greater than .5. Smith provides a proof that his definition preserves classical logic inferences and tautologies. Of course, this preservation requires that we jettison the semantics for the conditional used above.

Smith provides a proof that his definition of validity is identical to a classical one.

> It is not hard to show that the fuzzy consequence relation just defined on our standard first-order language is *identical* to the classical consequence relation on that language.[13]

The proof for this conclusion is provided in a footnote attached to the above quote. Sorites arguments are classically valid. Sorensen uses this validity in his

[12]Smith, 2008, p. 222.
[13]Smith, 2008, p. 222.

argument that the inductive premises of such arguments must be false. If the sorites argument is valid, then it must be unsound, or else we must accept their absurd conclusions. In order for a sorites argument to be unsound for Smith, it must be the case that one of the premises of the argument has a degree of truth less than or equal to .5.

Let's look at a particular sorites argument for heap that uses a finite number of conditionals as premises.

P1. 100 grains of sand make a heap.

P2. If 100 grains of sand make a heap, then 99 grains do too.

P3. If 99 grains of sand make a heap, then 98 grains do too.

...

P100 If 2 grains of sand make a heap, then 1 grain does too.

C. 1 grain of sand makes a heap.

According to Smith's semantics, $\phi \rightarrow \psi$ is equivalent to $\sim\phi \vee \psi$, and so the conditionals comprising P2-P100 are all disjunctions of this form. This argument is classically valid, and so it is valid in Smith's degree theory. So, if all of the premises are true to degrees strictly greater than .5, then the conclusion is true to a degree greater than or equal to .5. But, the conclusion is true to degree 0, and so at least one of the premises must be true to a degree less than or equal to .5. Imagine that P1 is true to degree 1 and for each grain we remove the degree of truth goes down by .01. The offending premise will then be P51: "If 51 grains of sand make a heap, then 50 grains do too." The antecedent is true to degree .51 and the consequent is true to degree .5. The value of the premise as a whole is the maximum between the negation of the antecedent and the consequent—$max(.49, .5)$, which equals .5.

One reason to move away from classical logic is to avoid the counterintuitiveness of epistemicism, which holds that there is an exact point along a sorites series at which the predicate no longer applies. It is not clear that we get away from that worry with Smith's degree theory. For every sorites argument, there will be a precise point at which we have to stop applying *Modus Ponens* because one of the premises of that *Modus Ponens* is only true to degree .5. This isn't to say that the predicate no longer applies—it still applies to some degree from that point on—but it does create a first-order stopping point in a sorites series that seems too sharp. If we're going to draw such a line, why not just say that everything above .5 is true and everything else is false? The issue brought up here is related to the problem of higher-order vagueness. The degree theorist may claim that 'being true to degree .5' is not always sharp, and so they do not bear the same counterintuitiveness that epistemicism does. I will explore the difficulties degree theories face in the realm of higher-order vagueness later in this chapter.

A Problem

Smith does good to find a definition of validity that preserves classically valid inferences. However, there is a philosophical problem with this definition. The particular problem comes from the definition of a tautology. $\phi \vee {\sim}\phi$ is a tautology according to Smith's definition, even though it can be the case that $V(\phi) = .50001$ and so $V(\phi \vee {\sim}\phi) = .50001$. Lowering the bar for a tautology from 1 to .5 may solve the formal problem of preserving classical logic, but this is accomplished by loosening the meaning of tautology. 'Jason Statham is bald' may come out to have the same value as 'Jason Statham is bald or he is not.' Since Jason Statham is a borderline case of 'bald' there seems to be something wrong with saying that a tautology can have the same truth-value as a borderline case.

Smith does attempt to justify his use of the definition. He argues the minimum value required for an *assertion grade* statement is .5. That is to say, it is acceptable to assert a statement only if it is at least .5 true.

> If a sentence S is at least 0.5 true, then one cannot make a truer statement by asserting the negation of S than by asserting S.[14]

[14] Smith, 2008, p. 224.

An *inference grade* statement is one that is acceptable for use in an inference. He claims that this value must be strictly greater than the degree required for assertibility. So, it must be greater than .5.

But it is incorrect to think that .5 is the threshold for assertibility. If a statement and its negation are equally true, then that should be a prime example of **un**assertibility. After all, if a sentence is .5 true, then it is genuinely a borderline case. It would be wrong to assert a borderline case. Part of the purpose in treating truth as coming in degrees was to capture the gradability of vague predicates. To say that someone is .5 bald is to say that they are halfway between being clearly bald and being clearly not bald—they're somewhere in the middle of the border area. Applying this to truth, saying that a proposition is true to degree .5 indicates that the sentence is half true and half false. We shouldn't assert something that is false to such a high degree.

Just as it is wrong to assert a proposition that is true to degree .5, it is wrong to use such borderline cases as premises in your reasoning, and thus they also should not be inference grade. If .5 is not assertion grade since it is the epitome of a borderline case, then .5000000001 should not fare much better with regard to inference grade.

The degree theorist is left with a difficult decision here. They must either undermine elements of classical logic or accept unintuitive meanings for 'validity' and 'tautology.' There may be a third route similar to Smith's that, unlike his, avoids weakening the conditions for a tautology, but it is hard to see what such a theory would look like. Given that many degree theorists already abandon elements of classical logic, it seems that their best option is to accept one of the first two definitions of validity.

The logic of states of affairs does not suffer from a problem with its definition of 'tautology.' A tautology is a state of affairs that cannot unobtain. By virtue of the fact that it cannot unobtain, we also rule out its being indeterminate, and so we do not end up with phenomena where a borderline case is just as true as a tautology.

6.3 Non-degree-functional Approach

So far we have seen a semantics for a degree theory that provides degree-functional connectives. We have also seen a number of definitions for validity. In this subsection, we will turn to a non-degree-functional approach to the semantics for degree theories.

6.3.1 Edgington's Semantics

Edgington argues that a better semantics for a degree theory would piggyback on probabilistic calculations. The probabilities of $A \wedge B$ and $A \vee B$ are not determined solely by the probabilities of A and B since the probability of a conjunction or of a disjunction depends on the other relations between A and B. However, the probabilities of $A \wedge B$ and $A \vee B$ are constrained by the probabilities of their components. The maximum value of $A \wedge B$ is the minimum value of its components and its minimum value is 0. Similarly, the minimum value of $A \vee B$ is the maximum value of its components and its maximum value is 1. Semantics like this are useful because they go some distance in capturing the dependence and independence of different components of a complex sentence. For example, $P \vee \sim P$ should get a value of 1. After all, the value of $\sim P$ is dependent on the value of P. On the degree-functional semantics described in section 6.2, supposing that $v(P) = .5$, the value of the disjunction would be .5. Edgington's semantics provides upper and lower limits for the degrees of truth of compound sentences. To calculate a precise value would require knowing the connection between dependent truth-bearers.

> I do not say that the indeterminacy of vague concepts is an epistemic matter. There exist different applications of probabilistic structure. Objective chances apply if and when the future is physically undetermined by the past. Relative frequencies also satisfy the principles of probability. I propose that it is also the right structure for theorizing about the indeterminacy of application of vague concepts.[15]

[15] Edgington, 1992, p. 201.

One proposal Edgington puts forth regarding how to use relative frequencies for describing the indeterminacy of vague concepts is the use of permissible precisifications to determine the degree of truth of a sentence.

> We could interpret the degree of truth of 'x is red' as the proportion of permissible sharpenings in which this sentence is true. A weighted proportion of permissible sharpenings may be better — a 'bell-shaped curve' over permissible sharpenings.[16]

We simply take the proportion of true to false evaluations over a set of permissible precisifications, perhaps also weighted by their permissibility, as a probabilistic structure for degrees of truth. One can see such a view as a combination of supervaluationism with a degree theory. However, instead of treating the sentences that have some true and some false precisifications as indeterminate, we give them a truth value corresponding to the proportion of true to false precisifications.

6.3.2 Objections to Degree-Functionality

Dorothy Edgington offers a number of objections to the degree-functional approach and provides guidelines for a non-degree-functional alternative. Edgington argues that the correct approach to the sorites paradox is to treat the sorites argument as valid, but unsound. She claims that this follows from the **Uncertainty Principle**, which states that the uncertainty of the conclusion of a valid argument is less than or equal to the uncertainty of the conjunction of the premises of the argument, which is less than or equal to the sum of the uncertainties of the premises.

$$u(C) \leq u(A_1 \wedge ... \wedge A_n) \leq u(A_1) + ... + u(A_n)^{17} \tag{6.1}$$

To put Edgington's point about the sorites argument succinctly, to say that the argument is invalid, it must be the case that the uncertainty of the conclusion is greater than the uncertainty of the conjunction of the premises. It may turn out that we should expect the conclusion to have a degree less than 1 because we have premises in a valid argument with degrees less than 1.

[16] Edgington, 1992, p. 201.
[17] Edgington, 1992, p. 193.

Edgington objects to the semantics held by proponents of the degree-functional approach. In particular, the degree-functional semantics given above do not consider the dependence and independence of the components of complex statements.

- $v(x$ is red$) = 1$, $v(x$ is small$) = .5$
- $v(y$ is red$) = .5$, $v(y$ is small$) = .5$
- $v(z$ is red$) = .5$, $v(z$ is small$) = 0$

1. x is red **and** small.
2. y is red **and** small.
3. y is red **or** small.
4. z is red **or** small.

According to the semantics of the degree-functional approach, all four of these sentences have a value of .5. However, intuitively $v(1) > v(2)$ and $v(3) > v(4)$.

> In trying to fulfill the command 'Bring me a ball which is red and small', would I not do better to bring x (than which nothing redder can be conceived) than y? And for the command 'Bring me a ball which is either red or small' would not y be a better choice than z?[18]

Suppose now that $v(x$ is red$) = .5$ and $v(y$ is red$) = .49$. It may make sense for the conditional "if x is red, then y is red" to have a high value, since it may be the case that x and y are close enough in shade such that if one is red, then the other would be too. However, this may not always be the case; for example, x may be close to red on the orange side of the spectrum and y may be close to red on the violet side of the spectrum. Of course, this may just be chalked up to a problem with the material conditional in general. The antecedent and the consequent in a material conditional do not have the right kind of modal connection that most English conditionals seem to have.

6.3.3 Smith's Response

Smith defends the degree-functional approach against Edgington's objections. In his response to the ball example above, Smith concedes that it would be wiser to

[18] Edgington, 1992, p. 201.

bring x than y. Less would have to change about the world to make x the correct ball than to make y the correct ball. He, however, argues that the fact that we may have to change more about the world to make y the correct ball than we would with x does not entail that there is a difference in the truth values of sentences 1 and 2 from 6.3.

> We need to distinguish two things: the distance of a sentence from the truth, in the sense of how much would have to change about its subject matter to render the sentence true; and the distance of a sentence from the truth, in the sense of how far its actual degree of truth is from the maximum truth value. A sentence might be very close to the truth in the first sense, and yet very far from the truth in the second sense. For example, if Bob is 6 ft 1/64 in., then although we would not have to change Bob much to render the sentence 'Bob is six feet tall' true, as things actually stand, the sentence is simply false.[19]

The disagreement here is about a central issue to any view of vagueness—the meaning of 'borderline.' For Smith, a borderline case is one that has a definite truth value, the value is just a degree. Smith claims that the degree of truth of a sentence comes from the degree to which the object actually instantiates its property. For Edgington, borderline cases are undetermined and, in some sense, could go either way. The degree of truth is the probability of the sentence being true. Other complaints with degree-functionality notwithstanding, the better semantics for a degree theory will depend on this central issue with the meaning of 'borderline.'

6.3.4 Problems with the Non-Degree-Functional Approach

The problems faced by the non-degree-functional approach depend on how one treats the probabilities used in truth-value calculations. For Edgington the probabilistic structure comes from the permissible precisifications. Herein lies the first problem for a degree theory in general, but more specifically for Edgington's. As was discussed in chapter 3, the supervaluationist definition of "supertrue" is itself vague. This vagueness is a consequence of the vagueness of "permissible." This is one of the origins of higher-order vagueness for supervaluationists. The other,

[19] Smith, 2008, p. 265.

equally problematic origin is the vagueness of “vague,” which results in situations where it is unclear whether or not a concept can be precisified. All of these problems of supervaluationism are also present in a degree theory that assigns values to sentences based on the proportion of true to false permissible precisifications.

In order to get a value for a sentence, we first need to check the proportion of true to false admissible precisifications. However, to do that, we must have a set of admissible precisifications. Which precisifications should be included in this set is vague, and so we need to check the precisifications of “admissible precisification.” Again, “admissible precisification” is still vague, and so we need an admissible precisification for “admissible precisification of ‘admissible precisification.’" We never get an actual value for the sentence since at every level we cannot determine the extension of the set of admissible precisifications. This problem with supervaluationism is exacerbated by the degree theoretic move. Supervaluationists can successfully evaluate clear cases and clearly indeterminate cases. They need only look at the clearly admissible precisifications to get a number of evaluations. A degree theoretic supervaluationist never gets any value, since we need the proportion of the set of admissible precisifications to get a value. Without a set of admissible precisifications, we do not get a proportion, and therefore do not get a value.

6.4 General Problems with Degree Theories

6.4.1 Higher Order Vagueness

If one does accept a different understanding of degrees of truth than the one Edgington does, regardless of which semantics they accept, one must accept a very complex understanding of higher-order vagueness. If sentences have degrees of truth and it is not vague what degree of truth a given sentence has, then there are infinitely many sharp borders along a sorites series. If it is counterintuitive that there is a last noonish nanosecond, then it should be even more counterintuitive that there is a last nanosecond that is noonish to degree .64927328—or whatever real number between 1 and 0 with any number of significant digits.

The degree of truth of a sentence is itself vague, and so it is not always clear what degree of truth a sentence has. It would be best to apply the same account of vagueness at all levels of higher-order vagueness, and so it can be true to degree n that a sentence is true to degree m. For any sentence we can construct a set of ordered pairs like the following $(0, a)...(.1, b), ...$, where the first member of each ordered pair is a real number between 0 and 1 and the second member is the degree to which the sentence is true to the degree represented by the first member.

Again, the degree to which a sentence is true to some other degree is also vague, and so it can be true to some degree m that it is true to degree n that a sentence is true to degree o. This process will continue *ad infinitum*. This gives us an extremely complex account of higher-order vagueness. For each sentence and for each of the infinite number of truth values there is a degree to which the sentence has that truth value and a degree to which that sentence has that truth value to the degree d, and so on *ad infinitum*.

6.4.2 Plurivaluationism

Smith defends his degree theory against this particular worry by accepting plurivaluationism. We first saw plurivaluationism in chapter 2 with Diana Raffman's most recent theory of vagueness. Raffman's plurivaluationism relativizes truth to precisifications. Smith's plurivaluationism relativizes degrees of truth to interpretations. Smith takes it that he has established that truth comes in degrees. However, this fact alone is not enough to resolve problems caused by vagueness. When we relativize gradable truth to precisification-like structures, we get a plurivaluationist degree theory. Precisifications do not divide people into bald and not-bald, they divide people into .56 bald and .1729 bald and so on.

One of the objections I leveled against plurivaluationism is that it has a serious problem if vagueness never bottoms out—that is, if we can never fully precisify our vague predicates/properties. We may attempt to give precise degrees of baldness by appeal to numbers of hairs and arrangements of hairs, but what counts as a hair is vague. So, we may attempt to give a degree of hairness by appeal to numbers of keratin molecules, but being a keratin molecule is vague. If we never reach a precise stopping point, something I argued is metaphysically possible in chapter 4, then it is possible to have a person whose degree of baldness is left indeterminate.

If Smith wants to include and exclude certain interpretations based on their acceptability as interpretations/precisifications of our natural language, then he will still have a problem with higher-order vagueness. An interpretation that treats 'Jeff Bezos is rich' as true to degree .1 clearly fails to match ordinary language. However, the boundary for what counts as an acceptable interpretation will be vague. If we have a completely indeterminately acceptable interpretation, and truth is only relativized to acceptable interpretations, then it may be unclear whether or not we can relativize truth to that interpretation.

6.4.3 Non-Gradable Indeterminacy

Consider the common words used in sorites arguments, bald, rich, and small. All of these are examples of gradable predicates. Part of the appeal of degree theoretic approaches to vagueness is that the kinds of words and properties that are vague are gradable, and so a real number value fits well. This is not the case with all kinds of indeterminacy. The quantum mechanical indeterminacy described in chapter 4 is a good example of non-gradable indeterminacy.

It would be wrong to say that electron a and electron b are identical to degree .5. They're not halfway to being identical in the way that someone who has half a head of hair is halfway between being clearly bald and clearly not bald. Certainly no other value between 0 and 1 fits the case of quantum indeterminacy. One advantage the logic of states of affairs holds over a degree theoretic approach is that its value of indeterminacy fits both the cases of vagueness and the cases of quantum mechanical indeterminacy.

6.5 The Logic of States of Affairs and Degree Theories

In this section, I will argue that the logic of states of affairs is preferable to a degree theory as a theory of vagueness. To begin, I will head off an objection that the logic of states of affairs is a kind of degree theory.

6.5.1 What is a Truth Value?

Suppose we have a state of affairs ϕ and the following state of affairs obtains: $\bigtriangledown\bigtriangledown\phi$. What does this tell us about ϕ? First, we know that $\sim\bigtriangleup\bigtriangledown\phi \wedge \sim\bigtriangleup\sim\bigtriangledown\phi$. Neither of these conjuncts tell us the status of ϕ. It could be the case that it obtains and it could be the case that it is indeterminate. In such a situation, all we can say about ϕ is that it is indeterminately indeterminate; we cannot conclude from this that it is indeterminate.

$\bigtriangledown\bigtriangledown\phi \nvDash \bigtriangledown\phi$

"Indeterminately indeterminate" seems like a truth value, and so would all of the other combinations of determinate and indeterminate operators. This seems even more powerful of an objection given that we cannot say anything about ϕ other than it being indeterminately indeterminate.

These combinations look even more like truth values when we consider how they affect complex states of affairs. What should we make of a conjunctive state of affairs where one conjunct obtains and the other indeterminately indeterminately obtains? The answer seems to be that the conjunction is also indeterminately indeterminate. As long as there is enough information to determine the value of the complex state of affairs, then it has that value. It gets more complicated when we don't have enough information.

For example, consider $\phi \rightarrow \psi$ where ϕ is indeterminately indeterminate and ψ is indeterminate. ϕ is in the border area between obtaining and being indeterminate. If ϕ obtained, then the conditional would be indeterminate. If ϕ were indeterminate, then the conditional would be indeterminate. This tells us that the conditional is clearly indeterminate.

Since it is open for a state of affairs ϕ that it might indeterminately determinately obtain, indeterminately determinately indeterminately obtain, or similarly for any other combination of determinate and indeterminate operators obtain, there appear to be an infinite number of truth-value-like values that a state of affairs can have.

We can even order these values along a number line from 1 to 0 where completely determinate ϕ is 1 and completely determinate $\sim\phi$ is 0.

6.5.2 States of Affairs and Degrees of Obtaining

I will now return to the original question from the beginning of this chapter. In chapter 4, I argued that vagueness is a metaphysical phenomenon and not just a linguistic one. In chapter 5, I described a logic of states of affairs that captures the vagueness in the world. The possibility of vagueness in the world poses problems for most theories of vagueness, but degree theories stand out as the other prime contender for describing vagueness in the world. In order for vagueness in the world to be described by a degree theory, it must be the case that a state of affairs can obtain to a degree.

So, can a state of affairs obtain to a degree? One could make a coherent degree theoretic logic for states of affairs, albeit a more complicated one than I give in chapter 5. We have seen already that degree theories will have a significantly more complicated account of higher-order vagueness. Vagueness is just a kind of indeterminacy. I argued above that degrees of truth are not a good choice for capturing non-gradable indeterminacy. So, a degree theoretic logic of states of affairs would need to accommodate both kinds of indeterminacy with different kinds of obtaining values—degrees and non-degree theoretic obtaining values.

Finally, this added complication is not necessary. Consider the point made by Rosanna Keefe.

> But although the measure of the underlying quantity may determine the applicability of the vague predicate, it does not follow that this measure is reflected in nonclassical numerical truth-values.[20]

Just because an object can be tall to varying degrees does not entail that it can instantiate the property <being tall> to greater and lesser extents. Gradable predicates can be precise as well. There could be a predicate that is gradable but where there is a sharp cutoff for its applicability. A more compact theory would hold that

[20] Keefe, 1998, pp. 575-576.

vague properties only differ from precise ones with regard to the blurriness of their border. That is, precise properties have cutoffs for applicability and so do vague ones—however, vague properties have vague cutoffs.

The fact that we have no positive reason for accepting states of affairs that obtain to degrees, combined with the greater simplicity of an account that does not include degrees of obtaining, points towards accepting the logic of states of affairs over a degree theory.

Appendix A

Higher-Order Vagueness

Throughout this dissertation I mention the phenomenon of higher-order vagueness multiple times. In this appendix, I will give a unified description of the phenomenon. There are two senses of 'higher-order vagueness' present in this dissertation. The first is the sense in which the predicate 'vague' is vague. This sense is only used in chapter 3 as part of the argument that semantic nihilism is self-undermining. The second is the sense in which it can be indeterminate whether or not something is indeterminate. For example, it can be indeterminate whether someone is indeterminately tall or determinately tall.

A.1 Vagueness of Vague

The property <being vague> is vague; it allows for borderline cases. Here is Roy Sorensen's argument for the vagueness of 'vague.'[1] As this is Sorensen's argument about predicates I will present it as such, but it could be done with properties as well. Let n-small be a predicate that applies to a natural number x iff $x < n$ or x is small. The predicate 1-small would apply to any natural number that is less than 1 or is small. In the case of 1-small only the 'small' part of the predicate will be relevant. Since 'small' is vague, 1-small will have borderline cases, and so 1-small is vague. 10000-small will apply to natural numbers that are either less than 10000 or are small. Here the 'small' component of the predicate is irrelevant because all of the indeterminately small natural numbers are less than 10000. 10000-small essentially reduces to 'is less than 10000,' which is not vague.

[1] Sorensen, 1985.

Base Premise: 1-small is vague.

Inductive Premise: For any n, if n-small is vague, then so is $(n+1)$-small.

Conclusion: 10000-small is vague.

The point at which a predicate n-small goes from being vague to not being vague is when the 'small' component of the predicate goes from being relevant to not. This occurs when n is large enough that all of the indeterminately small natural numbers are less than n. When this happens is vague, and so the inductive premise seems to be true. Sorensen treats sorites-susceptibility as a sufficient condition for vagueness, so 'vague' is vague. It can be vague whether or not a given predicate is vague. Or, as my account would put it, it can be vague whether or not a given *property* is vague.

A.2 Indeterminate Indeterminacy

Part of the reason that the sorites argument in the previous section works is that <being small> isn't just vague, it is higher-order vague. To know when <being small> is no longer relevant in evaluating whether or not something is n-small, we need to know if n is greater than all of the indeterminately small natural numbers. It can be indeterminate whether or not n is greater than all of the indeterminately small natural numbers. This means that it can be indeterminate that it is indeterminate whether or not n is indeterminately small. If n were greater than all of the indeterminately small natural numbers, then it would be determinately not small. If it were greater than some but not all of the indeterminately small natural numbers, then it would be indeterminately small. It is indeterminate whether or not n is greater than all of the indeterminately small natural numbers, and so it is indeterminately indeterminate whether or not n is not-small. The vagueness of <being vague>, then, is derived from the higher-order vagueness of properties like <being small>.

If vague properties are not precise enough to draw a boundary between the objects that instantiate them and the ones that don't, then it is unlikely that they are precise enough to draw two boundaries, one between the objects that determinately instantiate them and the ones that indeterminately instantiate them and the other between the objects that indeterminately instantiate them and the ones that determinately do not instantiate them. If vague properties are not precise enough to draw two boundaries, then it is unlikely they are precise enough to draw four. The first of these four would be between the determinately determinate cases and the indeterminately determinate cases. The second would be between the indeterminately determinate cases and the determinately indeterminate cases. The third would be between the determinately indeterminate case and the cases where the property is indeterminately determinately not instantiated. Finally, the fourth would be between the cases where the property is indetermiantely determinately not instantiated and those where it determinately determinately not instantiated. This process can go on infinitely.

A.3 The Structure of Higher-Order Vagueness

The structure of higher-order vagueness I've just described is depicted in the following figure.

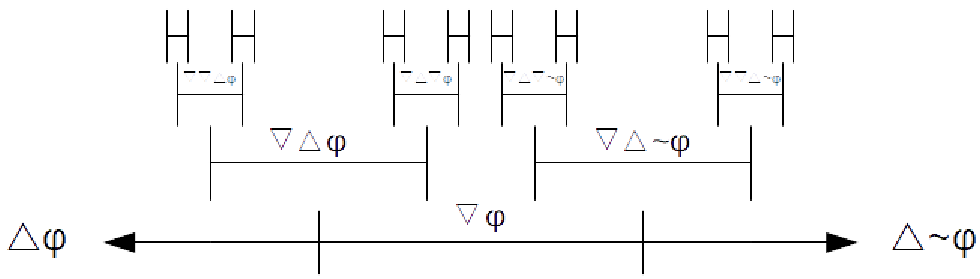

At each level of higher-order vagueness we get smaller and smaller divisions of the border area. We start off with just a division into the determinate cases and the indeterminate ones. At the second level we get a division into determinately determinate ϕ cases, indeterminately determinate ϕ cases, determinately indeterminate ϕ cases, indeterminately determinate $\sim\phi$ cases, and determinately determinate $\sim\phi$ cases.

Note that some of the labeled areas on the figure could be written differently. Indeterminately determinate cases are also indeterminately indeterminate. The area labeled $\triangledown\triangle\phi$ could instead be labeled $\triangledown\triangledown\phi$. This is because indeterminacy stands in between two options. If it is indeterminate that ϕ, then it is indeterminate that $\sim\phi$. If it is indeterminate that $\triangle\phi$, then it is indeterminate that $\triangledown\phi$.

Not everyone agrees that higher-order vagueness has the structure described here. Suzanne Bobzien[2] argues for what she calls Columnar Higher-Order Vagueness. It has the following structure:

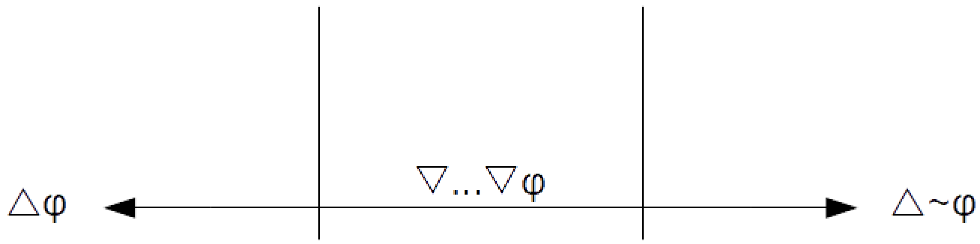

According to Bobzien, we get three columns: the infinitely indeterminate ϕ cases, the completely determinate ϕ cases, and the completely determinate $\sim\phi$ cases. I briefly argue against Bobzien's picture of higher-order vagueness in chapter 5.

[2]Bobzien, 2015.

Appendix B

Important Results from the Logic of States of Affairs

B.1 Semantics Review

The following are the obtaining and unobtaining conditions for the standard set of logical operators.

Negation Obtaining Condition #3: A negative state of affairs obtains iff its object state of affairs unobtains or it is metaphysically impossible for the object state of affairs to obtain.

Negation Unobtaining Condition #3: A negative state of affairs unobtains iff its object state of affairs obtains or it is metaphysically impossible for the object state of affairs to unobtain.

Conjunction Obtaining Condition #2: A conjunctive state of affairs obtains iff both of its object states of affairs obtain.

Conjunction Unobtaining Condition #2: A conjunctive state of affairs obtains iff at least one of of its object states of affairs unobtains or it's metaphysically impossible for both to obtain.

Disjunction Obtaining Condition #1: A disjunctive state of affairs obtains iff at least one of its object states of affairs obtains.

Disjunction Unobtaining Condition #1: A disjunctive state of affairs unobtains iff both of its object states of affairs unobtain.

Material Conditional Obtaining Condition #3: A material conditional state of affairs obtains iff either its antecedent state of affairs unobtains or its consequent state of affairs obtains or it is metaphysically impossible for the antecendent state of affairs to obtain when the consequent state of affairs doesn't.

Material Conditional Unobtaining Condition #3: A material conditional state of affairs unobtains iff its antecedent state of affairs obtains and its consequent state of affairs unobtains.

Material Biconditional Obtaining Condition #3: A material biconditional state of affairs obtains iff its object states of affairs both obtain or both unobtain or it's metaphysically impossible for its object states of affairs to have different obtaining-values.

Material Biconditional Unobtaining Condition #3: A material biconditional state of affairs unobtains iff one of its object states of affairs obtains and the other unobtains.

Universal Generalization Obtaining Condition #2: A universally quantified state of affairs obtains iff every instantiation of the quantifier obtains.

Universal Generalization Unobtaining Condition #2: A universally quantified state of affairs unobtains iff at least one intantiation of the quantifier unobtains or it is metaphysically impossible for every instantiation to obtain.

Existential Generalization Obtaining Condition: An existentially quantified state of affairs obtains iff at least one instantiation of the quantifier obtains.

Existential Generalization Unobtaining Condition: An existentially quantified state of affairs unobtains iff every instantiation of the quantifier unobtains.

These do not exhaust the logical relations that can hold between states of affairs nor do they exhaust the options for the given logical operators. However, I argue in chapter 5 that these are the best options for capturing our ordinary understandings of the operators.

B.2 Excluded Middle

$$\phi \vee \sim\phi \tag{B.1}$$

Excluded middle is not a tautology in the logic of states of affairs. If ϕ is indeterminate and it is possible for ϕ to obtain, then $\sim\phi$ is indeterminate. Since it is not the case that the disjunction obtaining condition is met and it's not the case that the disjunction unobtaining condition is met, (B.1) is indeterminate.

However, since it is impossible for $\phi \vee \sim\phi$ to unobtain, $\sim(\phi \vee \sim\phi)$ unobtains by the negation unobtaining condition.

B.3 Non-Contradiction

$$\sim(\phi \wedge \sim\phi) \tag{B.2}$$

The principle of non-contradiction is a tautology in the logic of states of affairs. Since it is impossible for ϕ to both obtain and unobtain at the same time, it is impossible for both conjuncts in (B.2) to obtain at the same time. So, by the conjunction unobtaining conditions, $\phi \wedge \sim\phi$ unobtains and (B.2) obtains.

B.4 $\phi \rightarrow \phi$

$$\phi \rightarrow \phi \tag{B.3}$$

(B.3) is a tautology in the logic of states of affairs. Since it is impossible for ϕ to both obtain and unobtain at the same time, it is impossible for the antecedent of (B.3) to obtain when the consequent unobtains. So, by the material conditional obtaining condition, (B.3) obtains.

B.5 De Morgan's

One important result from the logic of states of affairs is that $\sim(\phi \wedge \psi)$ is not equivalent to $\sim\phi \vee \sim\psi$. Consider the following example.

$$\sim(\phi \wedge \sim\phi) \nvDash \sim\phi \vee \sim\sim\phi \tag{B.4}$$

Suppose that ϕ is indeterminate. Since a state of affairs cannot both obtain and unobtain at the same time, $\phi \wedge \sim\phi$ unobtains. This entails that $\sim(\phi \wedge \sim\phi)$ always obtains. However, the disjunction on the right of (B.4) will be indeterminate when ϕ is indeterminate. So, the consequence relation does not hold between the left and right side of (B.4). Since $\sim(\phi \wedge \sim\phi)$ is a tautology, the consequence relation does hold from right to left.

As for the other form of De Morgan's, the equivalence using excluded middle and non-contradiction holds.

$$\sim(\phi \vee \sim\phi) \dashv\models \sim\phi \wedge \sim\sim\phi \tag{B.5}$$

The left side unobtains because $\phi \vee \sim\phi$ couldn't possibly unobtain. The right side also unobtains because it is impossible for a state of affairs to obtain and unobtain at the same time. Both sides of the equivalence are contradictions, and so the equivalence holds.

If the metaphysical impossibility conditions in the non-truth-functional operators are not enacted, then the equivalences between states of affairs of the forms in De Morgan's Law hold.

B.6 Contraposition

$$\phi \to \psi \dashv\models \sim\psi \to \sim\phi \tag{B.6}$$

The above equivalence does not hold in the logic of states of affairs. Consider the following instantiation.

$$\sim(\phi \vee \sim\phi) \to \sim\psi \models \psi \to (\phi \vee \sim\phi) \tag{B.7}$$

Suppose that ϕ and ψ are both indeterminate. We have already seen that $\sim(\phi \vee \sim\phi)$ unobtains and so the left side of (B.7) obtains by the material conditional obtaining condition. The right side, on the other hand, is indeterminate. Both the antecedent and the consequent of the right side are indeterminate. So, the consequence relation only indeterminately holds between the left and right side of (B.7).

B.7 Biconditional Exchange

$$\phi \longleftrightarrow \psi \dashv\vDash (\phi \rightarrow \psi) \land (\psi \rightarrow \phi) \tag{B.8}$$

(B.8) holds in the logic of states of affairs. Suppose that $\phi \longleftrightarrow \psi$ obtains, then both ϕ and ψ obtain, both unobtain, or it is impossible for ϕ and ψ to have different obtaining-values. If they both obtain, then both conjuncts on the right side of (B.8) obtain. If ϕ and ψ both unobtain, then both conjuncts obtain. Finally, if it is impossible for ϕ and ψ to have different obtaining-values, then it is impossible for ϕ to obtain when ψ doesn't and *vice versa*. So, by the material conditional obtaining condition, both conjuncts obtain.

Going the other direction, if $\phi \rightarrow \psi$ obtains, then either ϕ unobtains or ψ obtains or it's impossible for ϕ to obtain when ψ unobtains. If the right conjunct obtains, then either ψ unobtains or ϕ obtains or it's impossible for ψ to obtain when ϕ unobtains. Combining these together we get the following, either both ϕ and ψ obtain, both unobtain, or it is impossible for them to differ in obtaining-values. These are just the obtaining-conditions for biconditionals.

B.8 Double Negation

$$\sim\sim\phi \dashv\vDash \phi \tag{B.9}$$

(B.12) does not hold in the logic of states of affairs. Consider the following instantiation of (B.12).

$$\sim\sim(\phi \lor \sim\phi) \vDash \phi \lor \sim\phi \tag{B.10}$$

Suppose that ϕ is indeterminate. It follows from disjunction conditions #1 that $\phi \lor \sim\phi$ is indeterminate. However, since $\phi \lor \sim\phi$ cannot unobtain, it follows that $\sim(\phi \lor \sim\phi)$ unobtains and $\sim\sim(\phi \lor \sim\phi)$ obtains.

In the other direction the consequence relation holds.

$$\phi \vDash \sim\sim\phi \tag{B.11}$$

If ϕ obtains, then $\sim\phi$ unobtains. If $\sim\phi$ unobtains, then $\sim\sim\phi$ obtains.

B.9 Instantiation

Universal instantiation and existential instantiation are both valid in the logic of states of affairs.

$$\forall xFx \vDash Fa \tag{B.12}$$

If $\forall xFx$ obtains, then every member of the domain of quantification is such that it is F. Therefore, a must be F. It is not possible in this situation for a to be an indeterminately existing object. If it was, then $\forall xFx$ would be indeterminate by virtue of the fact that a could only indeterminately instantiate F.

Similar reasoning applies to existential instantiation. If $\exists xFx$ obtains, then there is at least one member of the domain, a, such that Fa obtains. Going the other direction existential generalization holds. If Fa obtains, then there is at least one member of the domain that is F, and so the domain has the property of having at least one member be F.

B.10 Existential and Universal Quantifiers (Incomplete)

One of the equivalences between universal and existential quantifiers does not hold in the logic of states of affairs, but the others do.

$$\exists xFx \Dashv\vDash \sim\forall x\sim Fx \tag{B.13}$$

If there exists an object that is F, then there is some object in the domain that is F, and so it is not the case that every member of the domain is not F, $\sim\forall x\sim Fx$.

If the right side of the equivalence obtains, then either $\forall x\sim Fx$ unobtains or it is impossible for $\forall x\sim Fx$ to obtain. If it unobtains, then there is some object such that it is F. If it is impossible for $\forall x\sim Fx$ to obtain, then this would entail that $\forall x\sim Fx$ unobtains by the universal generalization unobtaining conditions. To reiterate, a universal generalization unobtains when it is impossible for all of its instantiations to obtain. Since $\forall x{\sim}Fx$ unobtains, any instantiation of it must either unobtain or be indeterminate. It is possible that every instantiation is indeterminate, and so there is no object of the domain that is determinately F. As such, $\exists xFx$ is indeterminate. There is only indeterminately a consequence relation going from right to left in (B.13).

$$\forall xFx \not\dashv \vDash {\sim}\exists x{\sim}Fx \tag{B.14}$$

First, the proof from the left side to the right will start with the premise $\forall xFx$. Assume for a contradiction that $\exists x{\sim}Fx$. Since existential generalization holds in the logic of states of affairs, ${\sim}Fa$. Universal instantiation also holds, and so Fa. This is a contradiction, so if $\forall xFx$ obtains, then so does ${\sim}\exists x{\sim}Fx$.

${\sim}\exists x{\sim}Fx$ does not entail $\forall xFx$. Assuming that ${\sim}\exists x{\sim}Fx$, $\exists x{\sim}Fx$ either unobtains or it is impossible for it to obtain. Suppose that it unobtains. By the existential generalization unobtaining condition, every instantiation of $\exists x{\sim}Fx$ unobtains. For any x, it is either determinately not the case that Fa or it is metaphysically impossible that Fa. It could then be that Fa is indeterminate. So, ${\sim}\exists x{\sim}Fx$ does not entail that every object in the domain, x, is such that Fx.

B.11 *Modus Ponens*

P1: $\phi \rightarrow \psi$

P2: ϕ

C: ψ

Suppose for a *reductio* that C is not a consequence of P1 and P2. There are two ways that the consequence relation could fail to hold between the premises and conclusion of this argument. Suppose that P1 and P2 both obtain and that C unobtains. Since ϕ obtains and ψ unobtains, P1 unobtains by the material conditional unobtaining condition. Suppose that P1 and P2 obtain but that C is indeterminate. The only way that P1 could obtain under these conditions is for it to be impossible for ϕ to obtain when ψ doesn't. If it's impossible for ϕ to obtain when ψ doesn't, then the assumption that P2 obtains would entail that ψ also obtains.

B.12 General Points

In general, inferences that are valid in classical logic will still be truth-preserving in the logic of states of affairs when we restrict ourselves to states of affairs that either obtain or unobtain. For those inferences that are not preserved when we add indeterminacy, we can amend them and get a similar inference that is valid.

$$\sim(\phi \wedge \psi) \not\models \sim\phi \vee \sim\psi \tag{B.15}$$

De Morgan's is not valid in the logic of states of affairs. However, an inference like the one from left to right in (B.15) is valid.

$$\{\sim(\phi \wedge \psi), \Diamond(\phi \wedge \psi)\} \models \sim\phi \vee \sim\psi \tag{B.16}$$

If it is possible for ϕ and ψ to both obtain at the same time, then the modal part of the conjunction unobtaining conditions #2/3 will not be met. So, the truth of $\sim(\phi \wedge \psi)$ will entail that at least one of ϕ or ψ unobtains.

The inferences that are invalidated by the move to non-truth-functionality can be amended in ways like (B.16) to give arguments for which the consequence relation holds. This will need to be done with the operators that have modal components to their obtaining and unobtaining conditions. Here is a list of how to compensate for the modal intricacies of the operators.

$\sim\phi$: Add $\Diamond\sim\phi$ or $\Diamond(\phi)$ as a premise.

$\phi \wedge \psi$: Add $\Diamond(\phi \wedge \psi)$ as a premise.

$\phi \to \psi$: Add $\Diamond(\phi \wedge (\sim\psi \vee \bigtriangledown\psi))$ as a premise.

$\phi \longleftrightarrow \psi$: Add $\Diamond((\phi \wedge (\sim\psi \vee \bigtriangledown\psi)) \vee (\psi \wedge (\sim\phi \vee \bigtriangledown\phi)))$ as a premise.

Here's another example of this method at work. Contraposition is invalid in the logic of states of affairs, but an altered inference is valid.

$$\{\sim\phi \rightarrow \sim\psi, \Diamond(\sim\phi)\} \vDash \psi \rightarrow \phi \tag{B.17}$$

The initial problem with contraposition was the possibility that ϕ was an instance of excluded middle, and so the antecedent of $\sim\phi \rightarrow \sim\psi$ was always false. Since excluded middle can be indeterminate, the conseqent of $\psi \rightarrow \phi$ could be indeterminate. The issue arose because of the semantics for negation. Following the list above, we compensate for the modal conditions on negation by adding the premise that $\sim\phi$ is possible. ϕ's being impossible will not affect the obtaining-value of $\sim\phi$, and so it will follow the truth-functional semantics for negation.

Bibliography

Barnes, Elizabeth (2010). "Ontic Vagueness: A Guide for the Perplexed". In: *Noûs* 44.4, pp. 601–627.

Bobzien, Susanne (2015). "I—Columnar Higher-Order Vagueness, or Vagueness is Higher-Order Vagueness". In: *Aristotelian Society Supplementary Volume*. Vol. 89. 1. Wiley Online Library, pp. 61–87.

Braun, David and Theodore Sider (2007). "Vague, so untrue". In: *Noûs* 41.2, pp. 133–156.

Cameron, Ross P (2007). "Turtles all the way down: Regress, priority and fundamentality". In: *The Philosophical Quarterly* 58.230, pp. 1–14.

Edgington, Dorothy (1992). "Validity, uncertainty and vagueness". In: *Analysis* 52.4, pp. 193–204.

Eklund, Matti (2005). "What vagueness consists in". In: *Philosophical Studies* 125.1, pp. 27–60.

Evans, Gareth (1978). "Can There Be Vague Objects?" In: *Analysis* 38.4, p. 208.

Fara, Delia Graff (2000). "Shifting sands: An interest-relative theory of vagueness". In: *Philosophical topics* 28.1, pp. 45–81.

Fine, Kit (1975). "Vagueness, truth and logic". In: *Synthese* 30.3-4, pp. 265–300.

Garrett, Brian (1991). "Vague identity and vague objects". In: *Noûs* 25.3, pp. 341–351.

Gómez-Torrente, Mario (1997). "Two problems for an epistemicist view of vagueness". In: *Philosophical issues* 8, pp. 237–245.

IFAB (2019). *Laws of the Game 2019/20*. International Football Association Board.

Johnsen, Bruce (1989). "Is vague identity incoherent?" In: *Analysis* 49.3, pp. 103–112.

Kaplan, David (1979). "On the logic of demonstratives". In: *Journal of philosophical logic* 8.1, pp. 81–98.

Keefe, Rosanna (1998). "Vagueness by numbers". In: *Mind* 107.427, pp. 565–580.

— (2000). *Theories of vagueness*. Cambridge University Press.

Lowe, EJ (1994). "Vague identity and quantum indeterminacy". In: *Analysis* 54.2, pp. 110–114.

Ludwig, Kirk and Greg Ray (2002). "Vagueness and the sorites paradox". In: *Noûs* 36, pp. 419–461.

Lukasiewicz, Jan (1930). "Philosophical remarks on many-valued systems of propositional logic". In: *Jan Lukasiewicz Selected Works*.

MacFarlane, John (2016). "I—Vagueness as Indecision". In: *Aristotelian Society Supplementary Volume*. Vol. 90. 1. Oxford University Press, pp. 255–283.

Parsons, Terence and Peter Woodruff (1995). "Worldly indeterminacy of identity". In: *Proceedings of the Aristotelian Society*. JSTOR, pp. 171–191.

Plantinga, Alvin (1978). *The nature of necessity*. Oxford University Press on Demand.

Puryear, Stephen (2012). "Frege on vagueness and ordinary language". In: *The Philosophical Quarterly* 63.250, pp. 120–140.

Quine, WV (1981). "What price bivalence?" In: *The Journal of Philosophy* 78.2, pp. 90–95.

Raffman, Diana (1994). "Vagueness without paradox". In: *The Philosophical Review* 103.1, pp. 41–74.

— (1996). "Vagueness and context-relativity". In: *Philosophical Studies* 81.2, pp. 175–192.

— (2013). *Unruly words: A study of vague language*. Oxford University Press.

Sainsbury, RM (1989). "What is a vague object?" In: *Analysis* 49.3, pp. 99–103.

Sider, Theodore (1993). "Van Inwagen and the possibility of gunk". In: *Analysis* 53.4, pp. 285–289.

Smith, Nicholas JJ (2008). *Vagueness and degrees of truth*. Oxford University Press.

Sorensen, Roy (2001). *Vagueness and Contradiction*. Clarendon Press.

— (2005). "Précis of *Vagueness and Contradiction*". In: *Philosophy and Phenomenological Research* 71.3, pp. 678–685.

Sorensen, Roy A (1985). "An Argument for the Vagueness of'Vague'". In: *Analysis* 45.3, pp. 134–137.

Stanley, Jason (2003). "Context, interest relativity and the sorites". In: *Analysis* 63.280, pp. 269–281.

Thomasson, Amie (2007). *Ordinary objects*. Oxford University Press.

Tye, Michael (1990). "Vague objects". In: *Mind* 99.396, pp. 535–557.

— (1994). "Sorites paradoxes and the semantics of vagueness". In: *Philosophical perspectives* 8, pp. 189–206.

Unger, Peter (1979). "There are no ordinary things". In: *Synthese* 41.2, pp. 117–154.

Van Inwagen, Peter (1988). "How to reason about vague objects". In: *Philosophical Topics* 16.1, pp. 255–284.

White City Shopping Ctr., LP v. PR Rests., LLC (2006).

Williamson, Timothy (1997). "Imagination, stipulation and vagueness". In: *Philosophical Issues* 8, pp. 215–228.

— (2002a). *Knowledge and its Limits*. Oxford University Press on Demand.

— (2002b). *Vagueness*. Routledge.

Wright, Crispin (1975). "On the coherence of vague predicates". In: *Synthese* 30.3-4, pp. 325–365.