

Copyright by
Avram Hiller
2005

VAGUENESS AND NORMATIVITY

by

Avram Hiller

Department of Philosophy
Duke University

Date: _____

Approved: _____

David Sanford, Supervisor

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Philosophy in the Graduate School
of Duke University

2005

ABSTRACT

VAGUENESS AND NORMATIVITY

by

Avram Hiller

Department of Philosophy
Duke University

Date: _____

Approved: _____

David Sanford, Supervisor

An abstract of a dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Philosophy in the Graduate School
of Duke University

2005

ABSTRACT

Philosophers have devoted a lot of attention to vagueness in recent years, but there is still no general consensus about how to resolve the Sorites paradox. Timothy Williamson's *epistemic view*, which claims that our vague terms have unknown sharp boundaries, is the most popular and most controversial current account. No one has shown exactly what is wrong the epistemic view and no one has provided a satisfying alternative to it. These two projects – articulating what is wrong with the epistemic view, and providing a plausible alternative – are the primary goals of this dissertation.

Additionally, I survey ordinary intuitions that underlie Sorites paradoxes, and I note how these intuitions inform, and are informed by, a number of deeper philosophical debates in metaphysics, epistemology, philosophy of language, and ethics. In part, this serves as an explanation of why the Sorites paradox has remained so difficult to resolve.

The most common objection to the epistemic view – that it provides an unsatisfactory account of the connection between meaning and use – has not been successful in undermining the view. My own objection is a metaphysical, and not a semantic, objection: the epistemic view fails to provide the best explanation of what objects and properties exist. Instead, an eliminativist account of macro-level objects and properties, according to which there are no mountains and there is no property of being lavender-colored, is a better metaphysical account than one that claims that there are mountains and color properties that have sharp boundaries.

Of course, this eliminativist view is intuitively unappealing, and to show how statements in ordinary language can in some way be taken to be true, I introduce the *normative choice* account. According to this view, although non-normative facts about

linguistic behavior and about the external world do not determine a precise reference for our terms, our choices may do so. I claim that this provides all that is needed for there to be semantic normativity. First, we are still guided in our choices to some extent by psychological tendencies, and second, there are resources in semantic deliberation to respond to aberrant uses of language.

ACKNOWLEDGEMENTS

The writing of this dissertation was in many ways a long journey, involving many time periods and time zones, and those playing important roles in the completion of it are far too many to name here. The gratitude I feel to these individuals for their guidance and support cannot be adequately expressed in words. As a gesture, though, I would like to say “thank you” to the following:

David Sanford supervised the writing of this dissertation and has been a very careful and critical reader. I appreciate the ways in which he helped shape it. Even though there was no *official* North Carolina State Competition in Counterexamples, I count him as the undisputed champion. The other members of my committee – Michael Ferejohn, Owen Flanagan, and David Wong – have all provided invaluable assistance along the way. I also benefited from discussions with Carl Posy during the preliminary stages of the dissertation. Alasdair MacIntyre and Simon Blackburn gave very helpful comments on the dissertation proposal, and both have shaped my thinking in numerous ways.

My fellow graduate students at Duke provided wonderful companionship through the years. I learned a great deal from Frédéric Bouchard, Caery Evangelist, Steve Geisz, Marion Hourdequin, Boris Kukso, Paul Medeiros, Eddy Nahmias, Tom Polger, Brook Sadler, and Ana Santiago, all of whom made the writing of this dissertation much more pleasant than it otherwise would have been. I also had the very good fortune to have lived with two sharp philosophers who are also excellent listeners and good friends – Andrew Terjesen and Anthony Williams. And Robin Smith helped me to be more patient, among other things.

I thank Richard Heck for not once, but twice, welcoming me into his seminars on vagueness at Harvard. My approach to vagueness was surely shaped by these classes. Faculty in discussion groups at the Claremont Colleges and at Wake Forest University also provided welcome input. I'd like to specially thank Ralph Kennedy and George Graham for all their help and guidance. I am grateful to Stavroula Glezakos, who has been a supportive colleague and friend, and who better watch out, because I am going to win our little wager. At UNC, Ram Neta provided encouragement and illuminating discussion. In particular, our conversations helped shape Chapter Five – some of the text in it was lifted from an email to him. I also thank all those philosophers whom I have not yet met or had any causal contact with whatsoever, for providing inspiration. (I didn't have the chance to argue against causal accounts of reference in this dissertation, so I'm glad to use these acknowledgements to put in a jab.) And I am grateful to the many students I've had through the years. I learned much from Pasha Majdi and Ian Nance, among others.

The members of the 2004 Boston Red Sox convinced me of the possibility of things that I had previously thought impossible. I would also be remiss if I didn't mention certain non-humans. Emma, Bingo, Paris, and Query provided love and companionship at various stages of the writing of this work. Mountains and open spaces in North Carolina, Virginia, Tennessee, Arizona, Texas, New Mexico, Massachusetts, New Hampshire, California, Oregon, Montana, Colorado, and Utah provided a great deal of joy, peace, and perspective over the last few years. Some of the most pleasant nights of my life were spent near or under juniper trees along the Colorado Plateau.

I am fortunate to have an extraordinarily supportive and loving family. My parents, Gerald and Esther Hiller, have given an unlimited amount of support. Phyllis Fine and Susan Troen are incredibly fabulous sisters. Though it still feels weird to write my sisters' last names as something other than 'Hiller', I am also grateful to my brothers-in-law, Ken Fine and Aron Troen. I have very cool cousins, and one of them, April Hiller, formatted the dissertation properly and provided some comic relief. Finally, my niece, Melissa Fine, though not much aware of the nuances of the Sorites paradox, brought me many lasting smiles that helped me make it through.

TABLE OF CONTENTS

ABSTRACT	iv
ACKNOWLEDGMENTS	vi
LIST OF FIGURES	xii
CHAPTER 1. THE SOURCES OF SORITES	1
1.1 Goals for this Chapter.....	1
1.2 A Sorites Paradox	3
1.3 Preliminary Remarks on Realism/Anti-realism.....	4
1.4 Language as a Tool.....	6
1.5 Eliminativism, Heracliteanism, and Reductionism.....	8
1.6 Varieties of Realism, Anti-Realism, Relativism, and Skepticism	9
1.7 The Major Premise	20
1.8 Rules of Language	24
1.9 Forced Marches and Toleration.....	28
1.10 Concepts as Functions	34
1.11 Higher-Order Vagueness	38
1.12 Degree Theoretic Intuitions	48
1.13 Bivalence	51
1.14 Moderate Indeterminism, the Phenomenological Sorites, and the ‘Myth’ of the Given	55
1.15 Conclusions: Vagueness as Noticed Indeterminacy, and the Is/Ought Gap in Semantics	63

CHAPTER 2. SOME CONTEMPORARY APPROACHES TO SORITES	67
2.1 Graff's Interest-Relative Account.....	68
2.1.1 Graff's "Shifting Sands"	68
2.1.2 A Response to Graff: But the Sands Don't Really Shift.....	71
2.2 Wright's Intuitionist Account.....	83
2.2.1 Wright's "On Being in a Quandary"	83
2.2.2 A Response to Wright: On <i>Ceasing</i> to Be in a Quandary.....	88
2.3 Degree-Theoretic Approaches.....	95
2.4 Supervaluation.....	103
CHAPTER 3. THE EPISTEMIC VIEW.....	110
3.1 The Epistemic View and Bivalence.....	112
3.2 Margins for Error and Epistemicism	122
3.2.1 Williamson's Margin for Error Principle	122
3.2.2 The Error in Williamson's Margin for Error Principle	125
3.2.3 A Reconstructed Margin for Error Principle Applied to Vagueness.....	133
3.2.4 Margin for Error Principles Marginalized	137
3.3 Interlude: The Normative Semantic Problem.....	140
3.4 The Connection Between Meaning and Use	143
3.5 Interlude: Sorenson's Epistemicism	146
3.6 Vagueness and Realism	147
CHAPTER 4. VAGUENESS AND NORMATIVITY	167
4.1 Kripke's Wittgenstein, Vagueness, and the 'New Riddle of Rule-Following'	169
4.2 The Is/Ought Gap in Semantics	181

4.3	Millikan’s Teleo-Semantics.....	186
4.4	Normativity and Vagueness	195
CHAPTER 5. THE NORMATIVE CHOICE SOLUTION		201
5.1	The Normative Choice Solution	203
5.1.1	The Basic Bivalent Approach.....	203
5.1.2	Degree Theoretic Normative Choice	205
5.1.3	Supervaluationist Normative Choice.....	207
5.1.4	Three Valued Normative Choice.....	210
5.2	Problems for Normative Choice	210
5.2.1	Aberrant Choices and Semantic Deliberation.....	210
5.2.2	Choice and Determinism	214
5.2.3	Stipulations and Posits	218
5.3	Problems for Eliminativism.....	220
5.3.1	Are There Really Any Eliminativist Explanations?.....	221
5.3.2	An Argument from Williamson	223
5.3.3	Is it Really the Best Explanation?	224
5.4	Concluding Thoughts	229
BIBLIOGRAPHY		234
BIOGRAPHY		239

LIST OF FIGURES

Figure 1. The Stipulated Supervaluationist Spectrum	207
---	-----

Let us still give special consideration to the formation of concepts. Every word immediately becomes a concept, inasmuch as it is not intended to serve as a reminder of the unique and wholly individualized original experience to which it owes its birth, but must at the same time fit innumerable, more or less similar cases—which means, strictly speaking, never equal—in other words, a lot of unequal cases. Every concept originates through our equating what is unequal.

- Nietzsche

And so castles made of sand slip into the sea... eventually.

- Jimi Hendrix

CHAPTER 1. THE SOURCES OF SORITES

1.1 Goals for this Chapter

I undertake a difficult task in this chapter – to inventory the beliefs about vagueness and vague predicates that lead to Sorites paradoxes. The working title for this chapter was “Vagueness: an Intuition Inventory”;¹ I had hoped to put on the table our most basic intuitions which lead us into Sorites paradoxes, without my arguing for or against any particular view. Eventually I realized that many of the views I discuss here are more appropriately called ‘considerations’ than ‘intuitions’, and that the complication involved in specifying each consideration, as well as my desire to point towards what seems to me to be the correct approach to vagueness, led me to decide against a totally

¹ Other titles I considered were ‘Vagueness: A Compendium of Considerations’, ‘Vagueness: a Road Map’, ‘An Insider’s Guide to Vagueness’, and my sentimental favorite, ‘Vagueness on \$40 a Day’.

neutral discussion here. If one has characterized the problem of vagueness in clear, brief, explicit terms, likely one has done something wrong.

My goals in beginning my dissertation this way (rather than, say, by surveying the literature) are several. First, given that Sorites is deeply paradoxical, it is likely that we will have to abandon at least one well-entrenched belief about linguistic practice. Because this chapter provides a guide to what commitments we seem to have, it will help determine exactly where the disagreements lie, and what the costs are for adopting any given view of Sorites. I believe that a brief discussion of many views is worthwhile, as many of these views have already been widely discussed in philosophical literature. I shall mention the views, discuss them somewhat briefly, and indicate how the views relate to Sorites paradoxes; in most cases, I note the direction of analysis I give in later chapters.

I discuss how vagueness relates to wider questions of reductionism and eliminativism in metaphysics/philosophy of science; to questions of conceptual/non-conceptual content in epistemology; to questions concerning free choice in the philosophy of mind; to questions concerning the is/ought gap in ethics. I show how views on both sides of the realism/anti-realism debate contribute to Sorites paradoxes. And I hope noting these connections serves as an explanation for how difficult the problem of vagueness is, and for why, after many years, there is little consensus as to how to solve Sorites paradoxes.

Given that there is widespread dissatisfaction with the present state of literature on vagueness, I believe that it is good to begin a dissertation on vagueness somewhat independently of the literature. My own final view is different than all others currently on

offer, and I believe that beginning with a more general discussion is helpful in setting the stage for what I believe is the right approach to Sorites. In particular, there is one substantive point that I argue for in this chapter: vagueness shows that either Timothy Williamson's *epistemic view*, which claims that there are unknowable sharp boundaries to vague concepts, is correct, or that there is *radical* indeterminacy of semantics from non-semantic facts. This is a rejection of the common-sensical view that vague predicates are determinate for some cases and indeterminate for other (border-area) cases. I believe that the general commitment to a moderate form of indeterminism is undermined by taking vagueness, and higher-order vagueness in particular, seriously.

1.2 A Sorites Paradox

It's best to start with a Sorites paradox. Each common statement of the paradox (using notions of 'baldness', 'heap', or color predicates) has some advantages and disadvantages relative to others. I hope that nothing important rides on my choice of formulation – I use color concepts, though I do have qualms about this formulation. Hopefully, whatever problems in the very formulation of the paradox are discussed adequately in what follows.

Imagine a very long box of candies, numbered left to right from 1 to 1000. Candy #1 is lavender-colored, and candy #1000 is clearly not lavender; let us say that it is azure. The candies in between are ordered such that every candy is indistinguishable in color from the candies on either side, such that the series is a gradual progression from lavender to azure. I discuss issues of distinguishability below in further detail, but for the moment, let us assume that for any adjacent pair of candies, when viewed together, a

normal observer cannot detect any difference in color. However, I shall assume that for any pair that are separated by one candy, such as 567 and 569, there will be a slight noticeable difference. This assumption simplifies things, and nothing controversial trades on it. I shall call this series, following convention, a *Sorites series*. In the following statement, let L_p mean: candy p is lavender.

- | | | |
|-----|---|-----------------|
| (1) | L_1 | (Minor premise) |
| (2) | $(\forall p) (L_p \rightarrow L_{p+1})$ | (Major premise) |
| | L_{1000} | |
| (3) | L_{1000} | |

(3) is said to follow from (1) and (2) from 999 instances of universal instantiation and modus ponens. But (3) is false, because

- (4) $\sim L_{1000}$

Hence, from premises that seem to be true and from logic that appears unassailable, we have concluded something that we believe to be false. This is a paradox not merely because we have generated a contradiction. It is a paradox because we are *committed* to claims that Sorites reasoning shows to be incoherent.²

1.3 Preliminary Remarks on Realism/Anti-realism

In a 1983 paper, Hilary Putnam claims that vagueness demonstrates that there is something wrong with a view he then calls *metaphysical realism*.³ His argument there, following Dummett, is that if metaphysical realism is true, then there is a fact of the

² Richard Heck made a similar point to begin his seminar on vagueness.

³ Putnam (1983). I shall give a more complete analysis of the realism/anti-realism debate in Chapter Three, so for the moment, consider the terms ‘realist’, ‘anti-realist’, and ‘objective’ in my discussion here and in 1.6 as mere guideposts to ideas to be discussed in detail later.

matter whether each sentence is true or false. Because some vague sentences lack a truth value, metaphysical realism is false. What Putnam seems to have underestimated is the difficulty of giving a coherent anti-realist account of vagueness, and in the nearly twenty years since, very little progress has been made in that project. In fact, the prominence of Williamson's realist epistemic view in contemporary literature indicates that may be more plausible to use Sorites as a *reductio* of the major premise, and deny the claim that some vague sentences lack a truth value. Because Sorites paradoxes arise out of the combination of a realist premise with an anti-realist premise, they can potentially be used as a *reductio* of either – not just of realism, as Putnam supposed. And because Sorites paradoxes depend upon the resolution of a deep realism/anti-realism debate, views that simply introduce a logic to deal with the paradox (supervaluation, degree-theory, three-valued approaches, para-consistent, intuitionist) without resolving the underlying issues are bound to be unsatisfactory.

In particular, the major premise seems motivated by an anti-realist assumption that if one object is lavender and another is not, there must be a difference in hue perceivable by normal human observers using casual observation. Hence, vague-terms are in some way mind-dependent; that the truth conditions for terms such as 'lavender' are in some way non-trivial way *mind-dependent*. On the other hand, the minor premise seems to carry along with it realist intuitions – that objects are lavender independent of whether anyone has seen the object; that even if there were no people, there would still be colors in the world; that our perceptual mechanisms have no effect on the color of external objects. In short, that the truth conditions for terms such as 'lavender' are in some non-trivial way *mind-independent*. What is needed in order to resolve Sorites

paradoxes is either a view which combines both realist and anti-realist intuitions, or a view which takes one side in the realism/anti-realism debate and carefully explains away, using independent considerations, our motivations for maintaining the intuitions on the other side.

1.4 Language as a Tool

Though most accounts of the Sorites paradox focus on the major premise, it will be helpful to begin by discussing the *minor* premise (1) as well as (4).⁴ A common response from non-philosophers to Sorites is not to deny (1) and (4), but to claim that the terms used in (1) and (4) are *our* words, and that they shouldn't cause the trouble that Sorites causes. What considerations might support such a response?

One consideration I shall call the *language is a tool* view. I believe this is a widely held view, though of course it comes in many forms and with varying levels of justification. A basic characterization of it is as follows: we use language because it is helpful to us, and we need not consider language to be representational. If I say to a friend: "do not eat the azure candies!", the speech act has a certain functional equivalence to an act of physically preventing my friend from eating the azure candies – both acts serve the purpose of getting my friend not to eat them. (This is not to say that the acts are identical, but they are both tools to bring about the same purpose.) Because the act of physically preventing my friend from eating the azure candies itself does not represent anything, and has the same function as the speech act, our best explanation of the speech

⁴ Note that the paradox, as I state it, is a thought experiment, where I have stipulated (1) and (4) as being true; so this discussion should be interpreted as focusing on the legitimacy of making such stipulations.

act might not claim that I am using the words to represent anything. On this view, ‘lavender’ and ‘azure’ do not refer, and so the minor premise is not true.

One might add to the *language is a tool* view a pragmatic theory of truth, under which the minor premise would be true because it is in some way useful to us. How would this view solve Sorites? One possible approach is to claim that because the minor premise leads to paradox, it is not actually helpful, and so is not true. However, this would have the effect of making most or all of our statements false, because most terms are Sorites-susceptible. In response, a defender of this pragmatic theory of truth might add contextual factors as partly determining the pragmatic value, hence truth, of statements; in ordinary conversation, statements such as ‘that is lavender’ are helpful, whereas within philosophical discussions of vagueness, they are not. For various reasons, I’m not optimistic about this view.

I believe a better pragmatic theory of truth would approach Sorites by claiming that pragmatic value comes in *degrees*. Saying that Candy 1 is lavender is more useful than saying that Candy 500 is lavender; pragmatic value decreases as one proceeds down the Sorites series. Degree of pragmatic value isn’t a completely mysterious notion, and it can provide a motivation for claiming that the *major* premise in Sorites is not fully true – the antecedent in many instances will have a very slightly higher value than the consequent. The details of such an account would have to be worked out, and degree theory approaches face certain problems (as I show below), but this approach might be promising.

However, a proponent of the *language is a tool* views seems obligated to provide some explanation of *why* it is that language is a *successful* tool. It seems that the most

plausible explanation is that our words refer to actual features of the world – ‘azure candies’ refers to intersubjectively identifiable objects – and it is because of this that my speech act ‘Do not eat the azure candies!’ functions as a successful tool. So we have not dispensed with the notion of representation. The language as a tool intuition does provide some motivation for the view I call *relativism* that I discuss in 1.6, so I shall return to these issues shortly.

1.5 Eliminativism, Heracliteanism, and Reductionism

One might approach the minor premise not by claiming that there is no representation, but rather by claiming that there is no representation of macro-level objects, as our everyday language seems to presuppose. This brings the debate to questions of *eliminativism* and *reductionism*. An eliminativist ontology does not consist in macro-level ordinary objects; we simply use words like ‘azure’ because we are now unable to give a scientifically legitimate account of colors that is both correct and useful for everyday language use. In an ideal scientific language, Sorites paradoxes would not arise because terms such as ‘azure’ would not be used. However, as I attempt to show in Chapters Four and Five, there is a version of eliminativism that is not motivated by such scientific assumptions.

A related intuition that might lead one to reject the minor premise is a *Heraclitean* intuition. On this view, one cannot step into the same Sorites-series twice.⁵ We know that particles, even within solid objects, are constantly in flux, and this might be reason to

⁵ This intuition is related to the context-sensitivity view, put forth by Graff (2000) and Raffman (1994, 1995), which I discuss in more detail in Chapter Two.

deny any claims which assume that there are stable objects which persist over time. This view might deny the minor premise not on the mere grounds that azure is not a property, but on the grounds that candies do not exist. I am unsure of the prospects for a satisfactory Heraclitean ontology, though it seems that it is a kind of eliminativism.

In order to support the minor premise, one might be a *reductionist* rather than an eliminativist. On this view, our macro-level terms are reducible to micro-level terms, and so are in fact legitimate; we simply do not presently understand the reduction. On this view, the reference of 'azure' can in theory be modeled in terms of objective, precise, micro-level qualities of objects, such as their light reflectance. This view would resolve the Sorites by denying the major premise. Reductionism may face the complication that it is uncertain whether even our best scientific explanations of the world will be precise. I believe this should not be an immediate worry for the reductionist, because at issue at the present moment is Sorites as I state it above. While vagueness and imprecision at micro-levels may be paradoxical, without a full scientific theory of the world, it is difficult to anticipate what those paradoxes will look like; indeed, concerns of micro-level Sorites paradoxes might provide some guidance for scientists in attempting to conceptualize micro-phenomena, and presumably theories will be in some way constrained and motivated by an effort to avoid such paradoxes.

1.6 Varieties of Realism, Anti-Realism, Relativism, and Skepticism

I shall not discuss the three views mentioned in 1.5 much further in this chapter, except to note, first, that if one of them is correct, it may be used to resolve Sorites, and

second, that all three relate to much larger questions in metaphysics and philosophy of science that are beyond the scope of this dissertation.

The reductionist view may underlie what I shall call the *naïve epistemic* view, as it connects with a form of naïve realism. On this view, our terms refer to real precise properties and objects in the world (regardless of whether they are reducible to micro-level properties or objects). We may not *know* exactly what their precise boundaries are, but there are boundaries. This view faces the problem that it does not seem that for terms such as ‘lavender’, there is anything more that we could discover that would determine their precise reference. However, if an account of that can be given, then it is an elegant way to resolve Sorites, because it simply denies the major premise. I shall have much more to say about it in Chapter Three.

Another possible view is a *naïve vagueness-in-rebus* view, according to which our vague terms refer to real imprecise properties of the world. Contrary to a general perception that it is difficult to motivate this view, it does come from two highly plausible claims: (i) ‘lavender’ refers to lavender (a real property in the world); and (ii) lavender is imprecise (as evidenced by Sorites paradoxes). There may be a variety of *in rebus* views, because (i) may be interpreted with different levels of metaphysical commitment – one can be a minimalist about reference and another may believe in a deep metaphysical correspondence between word and property/object. It might be easier to understand the *in rebus* view for proper names – terms such as ‘Mt. Baldy’ refer to a mountain with vague boundaries (see McGee, 1997). An *in rebus* view might claim that while there is no precise boundary for Mt. Baldy, there are still areas which are clearly not part of it. So the world is divided into three groupings: areas that are clearly part of Mt. Baldy, areas that

are clearly not part of Mt. Baldy, and areas (for example, surrounding its base) that are *partially* part of it. The *in rebus* view would need a coherent account of what it is to be partially part of something, and it would also face a problem of higher-order vagueness – it is unclear where the boundary is between the partially-part and the clearly-not-part.

Contrasting with the *naïve epistemic* account of terms such as ‘azure’ is a view on which truth conditions for statements involving the term ‘azure’ relate both to facts external and internal to humans. There is one sense in which no one should deny this – ‘azure’ only represents a certain color given that it is a term of *English*, so of course truth conditions for sentences will be relative to certain facts about a speaker, such as which language the speaker is speaking. But there is a non-trivial, realist, *representation-is-relational* view, or *relational realist* view for short. On this view, whether a representation is veridical depends not only on whether the world is actually the way in which it is represented, but also upon other factors about humans, such as facts about perceptual mechanisms. The motivation behind this is what motivates *subjectivist* theories of color – color properties are ontologically inseparable from both external and internal states (see esp. Thompson, 1995). Because this view claims that there is a fact of the matter concerning the relationship between our perceptual mechanisms and external features responsible for colors, I consider it to be a realist view; one may even say that the property is an *objective* one, even though one half of the objective relationship is a perceiving subject.⁶ This view supports a *relational epistemic* view, which solves Sorites by claiming that we are unaware of the exact structure of the relation which makes terms

⁶ An interesting discussion of the compatibility of objectivist and subjectivist theories of *ethics* is in Wiggins 1996; I believe similar considerations might apply to color perception and other observational properties.

such as ‘lavender’ precise, so the major premise is false even though we do not know where the boundary is.

Note that there may also be a *vagueness-in-rebus relational realist* view, which claims that the relation is *imprecise*; this view, like the standard *in rebus* view, will have problems with higher-order vagueness. There also will be other *in rebus* views that may piggyback on the other relational views discussed below. I shall not even mention these *in rebus* views because they also suffer the same problems as the initial *in rebus* view, and are much more poorly motivated.

One way to explain the relation between our words and the world is to note that we have some interest in perceiving colors. There are several varieties of might be called *interest-relational* views. On one of them, what matters are evolutionarily-determined interests. Evolutionary considerations may be used to motivate a naïve non-relational realist view, on the grounds that evolutionary success is indicative that humans have been able to discern real features of the world. What I have in mind here, however, is an evolutionary-relational view, which combines a subjectivist account of perception with an evolutionary explanation of our perceptual capacities. This view should be considered realist because it claims that there is a fact of the matter concerning the evolutionary history of the color-relation, and as such, it may support an interest-relational epistemic view.

On other interest-relational views, the truth conditions for terms such as ‘azure’ are more explicitly mind-dependent. The evolutionary relational view may include features of mind as part of the relation, but another interest-relational view, which I shall call an *agent-relational* view, points out how our interests are generated *consciously* and

deliberately. While realist relational views would likely claim that all humans share the same relevant internal structure which functions as a relatum, the agent-relational view need not do so. This view has much in common with *anti-realism* and *relativism*, but I believe that it also may still be considered a form of *realism*, for it too claims that there are facts about the world (given that an agent's interests are part of the world) that determine the reference for terms such as 'azure'.

A good way to explicate and assess the agent-relational view is to discuss its connection to relativism. A common objection to many relativist views is that the very claims that relativists make, such as 'truth is relative to different cultures', are stated in absolutist terms, and hence are not at core *relativist*. An initial agent-relational view agrees that the relation between external features and the agent's interests is an objective relation, and that is why I consider it to be a kind of relational realism and not truly a kind of relativism. This view may support a degree-theory approach to Sorites that I mention with regard to pragmatism. One of the problems such a degree theory might have is in determining the exact pragmatic value (60%? 59.9%?) of claims like 'Candy 500 is lavender'. This view, given that it claims that there is an objective fact of the matter about our interests, and that we may not fully be aware of exactly what our interests are, is consistent with the claim that there is some precise pragmatic value of statements like 'Candy 500 is lavender' that we are unaware of. That may be a benefit of the view, because it can be used to solve Sorites, but it may also be a drawback, because some account must be given of how there is a fact of the matter despite our ignorance. This view also faces a problem in that it may be difficult to explain the motivation for being anti-realist on the first-order, but realist at higher-orders. One further distinction worth

noting at this point is whether the terms represent a relation between aspects of the world and an individual language-user (an *individual agent-relational* view), or whether it is between aspects of the world and a group of language users (a *public agent-relational* view).

There is also what I shall call a *relativist* view. On this view, first-order truth conditions are interest-relative, but unlike the basic agent-relational view, there are no other levels of explanation which are objective. This is relativism “all the way down”. Adherents to this view must provide some account of why there is as much agreement as there is, given that there are no non-relative facts of the matter.⁷ And it is unclear for the moment how the relativist view purports to resolve Sorites. On a relativist view, only humans are the ultimate arbiters for the truth of claims. Given that there is no objective fact of the matter concerning whether candies are lavender, or even objective facts of the matter concerning our interests, it is unclear what constraints there are, if any, on our language. So the relativist may be able to reject whatever parts of Sorites she deems problematic.

This kind of relativism may be unappealing, but at present I would like to say something in its defense. I believe that it is difficult to phrase the most common objection to this kind of relativism. One form of the objection I have in mind is the attempt to show that the relativist is committed to the absurd consequence that in other possible worlds in which there are no people, there are no lavender objects (for example). But there is no reason the relativist cannot say (along with the realist) that that is simply a world in

⁷ Note that the view I am calling *relativist*, by giving a particular account of how interests are agent-relative, is much more specific than pragmatism – perhaps all interest-relational views can be considered pragmatist.

which no one is there to call those things ‘lavender’ – given that whether those items in those worlds are lavender is simply relative to how we are in the *actual* world. This does not show that ‘lavender’ is not dependent upon humans, but it does show that counterfactuals are not a good way to express the dependency relation. And it is unclear if there *are* any good ways to express the dependency relation such that it causes problems for the relativist.⁸

There is another objection to relativism that I believe is a strong one. Does the kind of relativist that I describe have any grounds for claiming that the interests in question that form the second half of the relation are anything but the relativist’s own? On common relativist views today, truth is relative to the interests of a *group* of individuals. But that involves a claim that there is a commonality between oneself and at least some other individuals. It also seems to involve a claim that there is a commonality between an individual at one time and the same individual at a different time. The commonality claim might be easy to motivate if one is really an agent-relational realist, and not an ‘all the way down’ relativist, and believes that there is a fact of the matter concerning how humans are similar. On those views, the truth conditions for ‘candy #1 is lavender’ can be expressed as: ‘in relation to us, candy #1 is lavender’, and that fact is not involved in any further subjective relations—it is an objective fact about the world (which includes us). But if one is an ‘all the way down’ relativist, and believes that the relational truth-conditions for terms like ‘lavender’ are based in groups, then it should be expressed as ‘according to us, in relation to us, candy #1 is lavender’.

⁸ This is partly inspired by Simon Blackburn’s defense of quasi-realism against similar objections; cf. Blackburn 1993, p. 178.

The reason why the first one is ‘according to us’ and the second is ‘in relation to us’ is that on this version of relativism, these relations are in principle knowable to us, and ‘in relation to us’ does not necessarily imply knowability whereas ‘according to us’ does. If such a relativist is pressed about how it is known that according to us, in relation to us, candy #1 is lavender, the relativist can prefix the claim with another level of ‘according to us’. But this seems unsatisfactory – if one is making a claim about something outside of oneself or about oneself at a different time, there seems to be a requirement for an epistemological explanation of how one knows such things. But a relativist seems to lack any grounds to give such an explanation –making another claim in the form ‘according to us...’ seems unsatisfactory, because it relies on the very same evidence to prove its point. However, when one is making a claim merely about oneself, there does not seem to be such a requirement.⁹ It seems that a more feasible relativism interprets such claims as: according to me, in relation to us, candy #1 is lavender. There are in fact multiple ways to be what I shall call an *individualist relativist*. Here are just three more different ways relativists can interpret such claims: (1) according to me, candy #1 is lavender. (2) according to me, according to us, in relation to us, candy #1 is lavender. (3) according to me, according to us, candy #1 is lavender. None of these claims seems as epistemologically problematic as the non-individualist relativist epistemology, but I am unsure which of these is the best characterization of a typical

⁹ Though even that is questionable – self-knowledge is by no means simple. Still, the problem of one’s own mind pales in comparison to the problem of other minds. I believe that non-individualist relativist views need an answer to the latter, and cannot give it, and that individualist relativist views need an answer to the former, and the prospects are *better*. Note that my terms *individualist* and *non-individualist* are not used to represent the debate in philosophy of mind and language concerning individualism and anti-individualism.

relativist view, and it is not my goal here to give a full catalogue of relativisms. However, I believe that such an enterprise would be worthwhile, despite its complicated multi-meta-level structure.¹⁰

I raise these issues not to give a full discussion of relativism, but to steer the discussion towards questions of *normativity*. Individualist relativisms seem to fail to provide a necessary normative component of language. On a normative picture of language, the truth or falsity of utterances must depend upon *intersubjectively* available standards. A realist might press the point that the best explanation for the intersubjectivity of the evaluation of judgments is that judgments purport to describe the world itself, and the world either is or is not the way in which the judgment describes it to be. For the individualist relativist, though, it seems that one is not saying anything that bears on anyone else, or even oneself at a different time, and this view is only barely different from skepticism. Though it is very difficult to motivate a non-individualist relativist epistemology (given that one is actually a relativist and not a relational realist of some sort), but an individualist relativist epistemology likely turns out to be little different from skepticism.

There is another distinction worth making. One might claim, with the realist, that the very notion of truth conditions dictates that truth conditions must not be human-relative, but agree with the relativist that nothing fits such a notion, and so there is no truth. This view should be called *skepticism*. The skeptic makes no first-order assertions

¹⁰ As an aside, I believe that many issues of vagueness would be resolved if we did dig deeper into such meta-level details, especially where higher-order vagueness is concerned. At times, it has struck me that the Sorites paradox could not hinge upon complications that are four meta-levels up, and very far removed from first-order ordinary language. However, ultimately, I do believe that such discussion is required to resolve Sorites. I shall discuss this type of issue further in 1.12.

that purport to be true, and therefore does not accept the minor premise to Sorites (or the major premise either, for that matter). This is a powerful way to reject Sorites paradoxes, though it comes at the cost of preventing the skeptic from making any *assertions*. I have found myself explaining a Sorites paradox to a non-philosopher, and saying something like, ‘if we know anything about the outside world, then we know that this is lavender’, and hearing the reply, ‘maybe we don’t know anything about the world outside of us, then’. A stalwart view, and I believe that in discussions of vagueness, it merits more consideration. I am unsure whether this (or any) skeptical argument, insofar as it is put forth as an argument, is coherent. Another form of skepticism is *Pyrrhonism*, in which one remains silent and provides no arguments. This is an interesting option, but a discussion of it is outside the scope of this dissertation.

Let me recapitulate the views I have discussed, and how they might resolve

Sorites:

1. *Eliminativism*: macro-level terms such as ‘lavender’ do not refer, because they are not part of a satisfactory scientific language. Hence the minor premise of Sorites is not true.
2. *Reductionism*: the referential structure of macro-level terms such as ‘lavender’ is reducible to micro-level properties, such as wavelength reflectance, which are precise. Hence the reference of macro-level terms is precise and the major premise is false, though we do not presently know which candy is the last lavender one.
3. *Naïve epistemic* view: terms such as ‘lavender’ refer to precise objective (human-independent) properties in the world. Hence the reference of terms is precise and the major premise is false, and we may never know which candy is the last lavender one.
4. *Naïve in rebus* view: terms such as ‘lavender’ refer to imprecise objective properties in the world. Hence the reference of terms is imprecise – there is no candy that is fully lavender and is next to a candy that is fully not

lavender. However, this view will need some account of higher-order vagueness.

5. *Relational realist* view: terms such as ‘lavender’ represent objective relations between human features and external properties. Hence the reference of terms is precise and the major premise is false, though we may never know which candy is the last lavender one.
6. *Evolutionary relational* view: terms such as ‘lavender’ represent objective relations between our evolutionary history and aspects of the world. Hence the reference of terms is precise and the major premise is false. Evolutionary biology will shed light upon which candy is the last lavender one, though we may never know which candy is the last one.
7. *Agent-relational* view: terms such as ‘lavender’ represent objective relations between one’s conscious interests and aspects of the world. This relation might not be understood by normal speakers, but it is precise, so the major premise is false. Psychology or socio-linguistics would shed light upon which candy is the last lavender one, though we may never know which candy is the last one.
8. *Relativism*: terms such as ‘lavender’ represent relations between our/one’s conscious interests and aspects of the world. This relation may itself only be understood subjectively, and truth conditions will never lie outside the scope of an individual. Hence there are many options for us to avoid Sorites; either the minor or major premise might be denied.
9. *Skepticism*: truth conditions must be non-relative, but there are no non-relative truth conditions, and so there is no truth. Hence both the minor premise and the major premise are false.
10. *Pyrrhonian skepticism*: [There is no paradox]

Though some versions of 1 depend upon realist assumptions, I shall show in the dissertation that there is another kind of eliminativism which is more closely aligned with anti-realism. 2 and 3 uncontroversially should be considered realist, and I argue that 4 through 6 also should be. 4 and 5 might be considered anti-realist; 6 and 7 definitely should. 8 might also be considered anti-realist. This classification relies on the categorization of realist views as ultimately relying only on objective facts, and anti-

realist views being at some level subjectivist. This warrants further inquiry, especially given the result that certain views are both realist and anti-realist, but the views that are both anti-realist and realist may be especially promising because they may respect both the anti-realist intuitions underlying the major premise and the realist intuitions underlying the minor premise. I discuss realism and anti-realism more explicitly in Chapter Three, and discuss more fully the prospects of these combined realist/anti-realist views. I also formulate a means of distinguishing between realism and anti-realism.

I should also note that all these views seem to come from a naturalistic perspective. The one exception may be the individualist relativist view, which relies on an irreducible notion of human thought as ultimately underwriting normativity, and is thus not easily accommodated within a naturalistic worldview. Because this is the direction I myself take, and I'm skeptical of current naturalistic approaches in philosophy of mind, I believe that Sorites paradoxes are an insurmountable obstacle for naturalism in philosophy of language. I hope to show that there is an is/ought fallacy committed by naturalistic semantics, and the consequence is that we need to turn to a non-naturalistic view of semantics. This is a bold claim, and requires a large project; hopefully, the other chapters in this dissertation will show that such a project is, first, required, if we wish to maintain any notion of representational semantics and resolve Sorites, and second, plausible in its own terms.

1.7 The Major Premise

On what grounds might one be inclined to accept the major premise in the first place? It seems to represent the widespread thought that if the color of one candy is

indiscriminable (or very nearly so) from the color of another candy, it can't be the case that one is lavender and the other is not. Real examples of heaps demonstrate the basic intuition well. Grains of sand are really tiny. Very small increments, such as the removal of one grain of sand, should not lead to different judgments about the status of two similar objects. I shall call this the *no ontological difference without a perceptual difference* intuition, or the *OD*→*PD* intuition, for short. This leaves open the question (for the moment) of whether the claim that ontological differences require discriminable differences is *necessary*, such that our inability to discriminate between the color two items partly *constitutes* the fact that they are the same color, or *contingent*, such that our inability to discriminate between the two is merely an *indicator* that they are the same color. The latter view seems in line with the *in rebus* view, and the former in line with one of the anti-realist views and relativism. The naïve epistemic view, on the other hand, flat-out denies the *OD*→*PD* intuition.

One might argue that a universal generalization of a material conditional misrepresents our beliefs about the status of two adjacent pieces in a Sorites series. Intuitions for (2) might come from clear cases – it seems like people would agree that if candy 2 is lavender, then so is candy 3. Both are clearly lavender. But whether even this supports the use of the material conditional is dubious – in natural language, if one says “if 2 is lavender, then so is 3”, one is not making a claim that is intended to be logically equivalent to a specification of the truth tables for the material conditional; one is likely reporting a fact about the similarity between 2 and 3.

The use of the classical material conditional as a representation of what underlies the major premise has a further problem, because for instances of the universal

generalization for candies which are in the *border-area*¹¹ between lavender and not-lavender, it is hard to say what our intuitive response concerning the truth of the material conditional would be. Both antecedent and consequent (seemingly) represent indeterminate cases, and we should not assume at the outset that on a logic that purports to represent our ordinary linguistic intuitions, ‘ $P \rightarrow Q$ ’ is true when both P and Q seem to be indeterminate.¹² The use of a counterfactual conditional would not help much, because it is hard to understand the following counterfactual in a way that makes it true: If candy c were lavender, then candy c+1 would also be lavender. And as stated, it is false – candy c+1 would not change color if candy c did. Furthermore, perhaps c’s color is an essential property of it, which might make the antecedent always false, hence the counterfactual necessarily true. A better characterization of our intuitions behind the major premise is something like the following conditional: if this candy (c) were lavender, then any candy which would differ from it as little as c+1 actually differs from c, would also be lavender. This faces a problem concerning the notion of ‘color difference’, which is not easily quantifiable, rendering claims such as “x and y are more different in color than are w and z” problematic. Still, such a notion might be manufactured. A generalization of a proper major premise might be something like the following: For any two candies, x and x’, and the color difference between them (ΔC), if x were lavender, then any candy differing

¹¹ I shall use the notion of *border-area* neutrally – I do not mean, merely by using the term, to argue that the border area is actually indeterminate. I simply use it to denote the vague range of candies, say, from 400 to 700, where common sense does not dictate an answer, and I do not wish to incorporate within the notion any particular account of higher-order vagueness that indicates whether the border of the border-area is vague or not.

¹² There are non-classical logics that weigh in on this issue, but it is not my purpose here to survey the variety of alternative logics.

from x by ΔC would also be lavender. It seems that the most basic intuition underlying the major premise is: ΔC is not enough to make a difference between two objects being different colors. This is simply the $OD \rightarrow PD$ intuition.

Crispin Wright and Stephen Read note (Read and Wright 1985, Wright 1987) that a paradox can also be generated from the following:

$$(2^*) \quad \sim(\exists p) (L_p \ \& \ \sim L_{p+1})$$

Wright demonstrates a paradox can be derived from (1) and (2*), even under intuitionist logic. (2*) seems to have a great deal of intuitive appeal. It can be read as: there is no pair of candies between which lies the boundary between lavender and not-lavender. I shall call this intuition the *No Sharp Boundary (NSB)* intuition, as Wright calls a paradox using (2*) the *No Sharp Boundaries* paradox.¹³ Because of the possibility of the No Sharp Boundaries paradox, a paradoxical result may be obtained even if (2) is not something anyone would assent to, and so the issue concerning the lack of intuitive appeal of the *material* conditional is insignificant.¹⁴

Why exactly is (2*) appealing? Its negation entails under classical logic the Problematic Existential:

$$(PE) \quad (\exists p) (\sim T_p \ \& \ T_{p+1})$$

¹³ Note that the NSB intuition (as I use it) just says that there is no sharp boundary between lavender and not lavender, and not that there are no sharp boundaries at all. Below, I shall discuss the connection between the $OD \rightarrow PD$ intuition and the NSB intuition.

¹⁴ As Wright points out in 1987, p. 228, fn11, (2) and (2*) are not logically equivalent under intuitionist logic unless L is effectively decidable, which is just what is being called into question by vagueness.

It should be noted that Wright's own solution, in "On Being in Quandary" (Wright 2001), is to accept the negation of (2*), which can be proved within intuitionist logic by viewing the NSB paradox as a *reductio*, but deny Double Negation Elimination, which is not intuitionistically valid, and is needed in order to obtain (PE) from the denial of (2*). If one can provide a motivation for using intuitionist logic, then there may be hope for resolving Sorites. I discuss Wright's view further in Chapter Two. In the remainder of this chapter, when I state higher-order paradoxes, I shall use the material conditional for simplicity's sake; nothing will trade on this use until my consideration of Wright's intuitionist view in Chapter Two, where I use Wright's own terminology.

1.8 Rules of Language

Underlying our intuitions concerning the boundaries of vague predicates might be considerations about the nature of language which are characterized by Wright in "Language Mastery and the Sorites Paradox" (Wright 1975) – that language use is guided by rules. However, the notion of a rule of language is highly problematic, and I'd like to make some preliminary remarks here as a way to provide some clarification. (Chapters Four and Five contains a more complete discussion.) Wright criticizes what he calls the *governing* view, which notes how observational predicates are acquired by normal observation.¹⁵ (Wright, following Dummett, uses the notion of an observational predicate in a technical way, but for our purposes, it conforms to a common-sensical understanding.) On the governing view, (A) the application of observational predicates is

¹⁵ I'm not sure whether he names it that because it is the view held by many in Oxford at the time, or because the view concerns how rules of language *govern* language. My guess is the latter.

determined/governed by some set of introspectible rules that govern our language use, and (B) these rules are such that their precision is constrained by human limitations. As such, no rule should have it that for two patches which are color-indistinguishable (using casual observation), one is accurately described as ‘lavender’ and the other is not, because our perceptual apparatus and our memory of detail are not refined enough to make fine-grained distinctions. Given this, rules of language cannot be such that a term applies to one member of a Sorites series while not applying to an adjacent member.

This picture provides a theoretical background for the NSB intuition. Wright claimed in 1975 that the Sorites was a *reductio* of the governing view. The governing view does seem plausible to begin with. But rejecting the governing view solely on the grounds that it leads to Sorites seems unsatisfactory (and the numerous papers written by Wright on vagueness since 1975 testify to that). There are two related reasons for this: first, showing that the governing view leads to paradox is not enough to temper our reasons for believing in it in the first place; second, we still need a positive account of language use which does not lead to paradox. Any argument that a common-sensical view leads to Sorites is not sufficient to get us to abandon the view, given that we have yet to discover any appealing alternative that does not lead to Sorites. So to get us to abandon the governing view, one must give independent grounds for rejecting considerations underlying it.

One of Wright’s purposes in “Language Mastery” is to support a Wittgensteinian view that rejects the possibility of the *introspectibility* of rules of language, as in part (A) of the governing view. Questions of what *introspectible* means are bound to involve deep, controversial issues in philosophy of mind, and should be approached cautiously.

My own emphasis is on the *governance* of rules, independent of whether they are introspectible; I shall argue later that (A) is problematic because rules, introspectible or not, cannot provide proper governance. Importantly, though, questions of introspectibility relate to my earlier discussion of whether or not terms such as ‘lavender’ represent a relation that, ultimately, is knowable to humans.

If rules are ultimately relative to conscious considerations, then there are good reasons for claiming that they will not be precise. This arises from more than just the limitations of our memory and perceptual acuity. Given that there is a continuous possible range of colors, no matter how finitely precise our vision is, we will never be able to specify the precise point within the potentially infinite range of colors where the boundary is (cf. Changizi 1999). Sorites paradoxes are convenient because they rely upon discrete cases, but in principle, if terms such as ‘lavender’ are precise, their precision must be an intensional precision – grouping not just all actual objects but all *possible* objects into those which fall under the concept and those which do not, i.e., the *extension* and the *anti-extension*. This puts a very strong requirement upon views that claim that color terms are precise.

An important caution in discussions of rules of language is to note what directives to behavior are issued by rules of language. There seem to be many possibilities for what norms, for behavior in general, do: they may *prohibit*, *discourage*, *permit*, *encourage*, or *require* certain behaviors. There are norms of dress and style, for example, that employ each of these categories. Four points here: (1) Permitting is equivalent to not prohibiting. (2) Discouragement and encouragement might come in degrees. (3) It might be claimed that prohibition is equivalent to 100% discouragement, and requirement is equivalent to

100% encouragement. However, such equivalencies should not simply be assumed prior to an examination of the terms. (4) The relationship between discouragement and encouragement is also unclear at the outset. Perhaps a norm neither encourages nor discourages certain behavior. However, such cases might be better considered cases where the norm simply does not apply.

A further point is that even within the realm of requirement, there is a distinction to be made between whether rules *indicate* or *dictate*. I use *dictate* to denote a traditional philosophical assumption that when rules issue requirements, they *govern* language use, whereas I shall argue in Chapter Four that the most rules can do is *guide* language use. At issue is how much normative force a rule has: whether rules are absolute governors as opposed to friendly guides. Guidance is conceptually different from encouragement in that the content of guidance one receives from a friend might be in the form of a requiring ‘you must...’, but the friend’s lack of authority makes this mere guidance and not governance. For analogous reasons, governing is not the same as requiring, because a governor can issue encouragement, in the form, ‘it is a good idea for you to....’ In sum, whether rules *dictate* or *indicate* is independent of the content of the rule. I shall in this chapter use the notation *indicate/dictate* in order to be neutral on the governing/guiding issue.

Rules are also conditional. It is not the case that rules of language indicate/dictate that candy #1 ought to be called ‘lavender’. One should not simply blurt out colors of items; there are contextual guidelines about when one ought to call something ‘lavender’. A better characterization is that rules of language indicate/dictate that if one is expressing a judgment of whether candy #1 is lavender, then one ought to say that it is lavender.

However, this is not quite right: If I wish to deceive someone, then I should say that it is not lavender. Another attempt: Rules of language indicate/dictate that if one wishes to truthfully express a judgment of the color of candy #1, one ought to call it ‘lavender’. This isn’t yet quite it: I can truthfully judge it as being not yellow. A better effort: Rules of language indicate/dictate that if one wishes to truthfully express a judgment of whether candy #1 is lavender (i.e., of its *lavender-status*), then one ought to call it ‘lavender’.

1.9 Forced Marches and Toleration

There is more that can be said about the intuitive justification for the NSB intuition. PE claims that there exists some patch that forms the boundary between lavender and not-lavender. The NSB intuition, combined with the view that there is some boundary (I shall call this the *classification* intuition – more below), entail that there is an imprecise boundary. But consideration of actual Sorites series shows that to be incoherent: *if there is indeed a boundary, it must be located at some particular place*, unless we can provide a revisionist account of either ‘imprecise’ or ‘boundary’. One effort to do so would be to claim that there is a boundary *region* between lavender and not-lavender; however, this causes two difficulties which I shall discuss later in this chapter: it rejects bivalence, and it leads to higher-order vagueness. It will be helpful now to examine the NSB intuition in more detail, and determine why one might be inclined to hold it even after it is demonstrated that the notion of an imprecise boundary is incoherent.

One way to focus the discussion is to give examples of what Horgan (1994) calls *forced-march Sorites*, thought experiments where imagined individuals (I shall call them

marchers) are brought through Sorites series and asked questions about the members of the series. I don't think that too much can be learned from consideration of forced marches, especially given how many possible forced marches there are, but it will help shed light on higher-order vagueness and the NSB intuition. There are several cautions to keep in mind as we proceed. We must be aware of the contextual differences between forced-march occasions and natural language situations. To say that one ought to respond in a certain way in a forced-march is not to say that one ought to do so in a natural language situation. Boundaries given in forced marches may not be indicative of any real boundary, or any supposition by the marcher that there is a boundary. A common objection to supervenience is that it only analyzes ways in which vague terms can be made precise, but might not account for the very fact that *any* precisification may violate our intuitions about language or the rules of language. Eventually some account must be given of exactly what forced precisifications given by marches tell us about vague terms, if they even indicate anything. We also need to keep in mind that what a marcher says may or may not be indicative of what the marcher believes, as forced-marches highlight an important distinction between two kinds of propositional attitudes. For some cases, marchers might express judgments about the color of members of a Sorites series, but only consider these judgments to be *provisional*, given the forced nature of the context. For other cases, the marcher might make a judgment and take it to be *normative*, insofar as the marcher believes that other people ought to also make the same judgment, and that the marcher herself ought to make the same judgment even when not in the forced situation. Additionally, forced marches bring to the surface the distinction of an object either (1) appearing lavender to a marcher, or (2) being judged by a marcher as being

lavender; marchers may be asked to report on either (1) or (2). The forced marches I discuss employ the latter.

Let us begin by considering one (call it *march A*) in which a marcher is forced to give a ‘yes’ or ‘no’ answer to the question of whether each candy is lavender. In such a march, presumably, there will be some boundary, as given by where the marcher stopped saying ‘yes’; let us suppose it is at 400.¹⁶ I shall call 400 the *1st-order boundary point*.

We can then subject the marcher to another line of questioning (*march A**), also in which the marcher must answer ‘yes’ or ‘no’, where we ask for each candy: “Is it admissible for another marcher to have called it (in A) ‘not-lavender?’” If the respondent answers ‘yes’, then we ask: “Is it admissible for another person to have called it ‘lavender?’”¹⁷

Presumably, this will leave us with three ranges: a range of candies in the middle, say 301 to 500, where the subject says ‘yes’ to both questions in A*, a range of candies on the left side where the marcher says ‘no’ to the first question, and a range of candies on the right side where the marcher says ‘yes’ to the first and ‘no’ to the second. A common way to characterize this marcher’s response is to say that candies 1 to 300 are *definitely lavender*, candies 301 to 500 as the border area, and candies 501 to 1000 are *definitely not lavender*; I shall call 300 the *2nd-order boundary point*.

¹⁶ I shall assume that for all the marches to be discussed, marchers give answers which are *neat*, meaning, for example, that for any candy that the marcher calls ‘lavender’, the marcher also calls all candies to the left of that candy ‘lavender’.

¹⁷ I shall also assume that marchers must describe their own responses from a previous march as admissible.

Another marched-Sorites¹⁸ (*march B*) is when the marcher is not forced to give a ‘yes’ or ‘no’ answer to the first set of ‘is this candy lavender?’ questions. This will give us a grouping of the series into three parts – the portion to which the marcher said ‘yes’, the portion to which the marcher said ‘no’, and the portion to which the marcher said something other than ‘yes’ or ‘no’.¹⁹ One note is that it is not contradictory for the groupings given by a marcher in B to be different than those given by the same marcher in A* – it is no contradiction to both (1) claim that candy 305 is lavender, when not forced to give a ‘yes’ or ‘no’ answer, and (2) claim that it is admissible for another person to say that 305 is not lavender when that other marcher *has* been forced to give a ‘yes’ or ‘no’ answer to the question. This consistency is one reason why it has been difficult for philosophers to settle upon one explanation of the meaning of ‘definitely’ – ‘definitely lavender’ could denote a second-order range where it is inadmissible to deny (in a first-order context) that something is lavender, or it could be one range of a first-order trivalent structure. It cannot be assumed that these two are co-extensional. One last march worth mentioning here (call it *march B**) is analogous to march A* – we ask whether one would accept alternative answers to those one has given in march B. There are many, many, other possibilities for ‘definitely’ that I shall not discuss, corresponding to the many other types of marched-Sorites.

¹⁸ I call this a *marched Sorites* to distinguish it from a *forced march Sorites*; a marched Sorites is when one is marched through a series, though is not forced to give one of a limited range of responses. I believe my terminology here is non-standard, though it is helpful to make such a distinction.

¹⁹ It is possible that we might analyze this kind of marched-Sorites as giving us more than three regions if we presume that the marcher gives answers for the middle ones in groups of, say, ‘probably’ and ‘probably not’, but I shall focus on the three regions listed.

How seriously should we take the results of these forced-marched Sorites? Our aim as philosophers is to get a better understanding of vague predicates and of marchers' intuitions about them, and certainly these marches tell us something. They indicate that people are *tolerant*²⁰ in some marches of other marchers' answers when the others are forced to give a 'yes' or 'no' answer. Presumably, marchers even have some tolerance (in march B*) about others' answers when the others are not forced to give a 'yes' or 'no' answer – this relates to higher-order vagueness, which I shall discuss shortly. There are two lessons I believe that we should learn from these marches. First, the likely results of marches A* and B indicate that we have intuitions that we have no confidence in any particular sharp boundary point between lavender and not lavender. In march A*, our toleration of alternative answers indicates that we are not confident that our own answers in march A are correct. In march B, trivalent responses indicate that we would not choose to commit to any sharp boundary. These considerations bolster the NSB intuition. However, it should be noted that the considerations given above might be accounted for epistemically – it might be the case that we are tolerant because we believe that no one can know where the sharp boundary is, and not because we believe that there is no sharp boundary.

A second lesson is that these marches remind us that there is *so much* relevant information that is relevant to the concept of lavender – every marcher's response to the four marches above, plus all other possible ones, and more. According to agent-relational views, reference is partly determined by speakers' dispositions. This might bolster an

²⁰ Wright uses the notion of *tolerance* to describe concepts; here, I use it to describe people.

epistemic view that claims that there is a boundary which we are not aware of; if all marches give us a small piece of information about a concept, then it seems that there is much that we are ignorant about.

I stated that something else must underlie the NSB intuition, and I believe that something else does. Many ordinary language users do have meta-theoretic intuitions that there is no fact of the matter where the boundary is between lavender and not-lavender. It seems to me that philosophers underestimate the degree of skepticism that has pervaded common-sense intuitions about language. In fact, I would guess that in marches B or B*, where marchers are not forced to give ‘yes’ or ‘no’ answers, many actual individuals would refuse to give assent/denial to even paradigm cases of lavender/not lavender, and not merely out of impatience with the process – Sorites brings to the surface what I shall call *indeterminist* intuitions that many individuals have. (This relates to much of what I discuss in 1.4 concerning the interplay of higher- and lower-order linguistic intuitions.) A possible sociological explanation of the prevalence of these intuitions would likely involve relativist sentiments from academia seeping into mass-media entertainment and then seeping into popular culture. This, of course, goes to show that we should not rely on these sorts of unreflected intuitions as determinant of our philosophical theories. Still, analysis of language must rely on these intuitions as being partly determinant of the norms of language, so it seems there is some feedback problem at play. How much of a problem it is depends partly upon how clearly we can make a distinction between meta-linguistic intuitions and first-order intuitions. At the very least, these considerations do show how deeply entrenched our intuitions are concerning the lack of a sharp boundary between lavender and not-lavender. One of my goals in the remainder of this chapter is to

argue that the indeterminist intuition, and not the NSB intuition, is the most basic intuition underlying Sorites.

1.10 Concepts as Functions

At this point, it is worth going beyond what lessons we learn from marched Sorites to how *philosophers* should view border-area cases from a reflected perspective, which I shall call *sub specie philosophiae*.²¹ The work that I hope this notion performs is fourfold. First, it indicates that as philosophers, we are not under duress (as is a marcher in a forced march) to give just a ‘yes’ or ‘no’ answer to questions about borderline cases. We can reflect on the matter, and perhaps ultimately choose to withhold judgment, make only provisional claims, or say something else entirely. The second bit of work the *sub specie philosophiae* notion does is to claim that insofar as we wish to make normative claims, we do so because we are confident that there is one response that we consider best; we do so believing that those who disagree with our claim are in error. This is because of what I shall call the *caution principle*, which states that one should not make a judgment that one is not confident in unless something about the situation forces one to make a judgment, in which case the judgment should be made *provisionally* and not *normatively*. I have no argument in favor of the caution principle, but it seems to be a very basic rule of language that we already accept. Third, *sub specie philosophiae* allows a clear way to express the most common objection to the epistemic view. One can say: if, *sub specie philosophiae*, no determinate precise reference for terms like ‘lavender’ is

²¹ This notion is similar to a notion of reflective equilibrium. But equilibrium implies a kind of balance, and I do not want to assume in advance that such balance will occur. If Hegelian *synthesis* is a highly optimistic view of outcomes of dialectic, and reflective equilibrium is a moderate optimism, *sub specie philosophiae* is a moderately pessimistic view.

seen, then there is no determinate precise reference. Fourth, the notion of *sub specie philosophiae* will permit discussion of (allegedly) full-knowledge conditions while avoiding thorny issues concerning omniscient beings or the view *sub specie aeternitatis*, both of which I believe are problematic notions from the outset.

The range of possible claims we can make *sub specie philosophiae* about borderline cases need not be assumed to be bivalent. Presumably, we should tolerate alternative responses in forced-march situations. But should we tolerate alternative responses when no answer is forced? There is good reason to claim that we should not be tolerant: we have deeply entrenched Fregean intuitions that *concepts are functions*, even if we deny Frege's claim that concepts are bivalent. What I hope to show is an even stronger Fregean claim: that concepts are *total* functions – for every element of the domain, there is one correct thing to say about it. (This argument, however, comes prior to consideration of higher-order vagueness, and in 1.11, I show how higher-order vagueness affects it.)

For each candy, we, *sub specie philosophiae*, determine as best as we can whether it is lavender. For candies on the far left and right sides of the Sorites series, I shall presume for the moment that we are confident that rules indicate/dictate that candies are properly called 'lavender' (left ones) and 'not-lavender' (right ones). For what I shall call *unconfident cases*, we are not confident what the rules indicate/dictate. Unconfident cases are not defined the same as border-area cases – the difference is that it might be possible that there are candies that we are confident that the rules indicate/dictate both that we should not call them 'lavender' and that we should not call them 'not-lavender'. These

would be confident border-area cases. So it is worth asking whether there are *any* unconfident cases.

On the assumption that there are, we should not permit any individual not under duress to call unconfident cases either ‘lavender’ or ‘not-lavender’, unless such judgments are made provisionally. The reason for this is that if we, *sub specie philosophiae*, are unconfident about a case, and another individual makes a confident judgment, we need to determine whether the individual knows something that we do not know which justifies the individual’s confidence. If we determine that there is something, then we hadn’t actually been viewing things *sub specie philosophiae*. But if there is nothing else, then we should claim that the individual has made an error.

There is something inherently problematic in claiming that a case is an unconfident case. For, by the caution principle, we should only make such a claim if we are confident that we are unconfident. But this is contradictory – if we are confident that we are unconfident, then we have some confidence about the case. The reason is that the view *sub specie philosophiae* combines first-order considerations with higher-order considerations when it delivers what we to say about cases. Specifically, I have in mind that for what we deem to be border-area candies, we should say that individuals in forced marched situations may either say that the candies are or are not lavender, but one who is not forced should not give any non-provisional ‘yes’ or ‘no’ answer.

If we reject Williamson’s epistemic view, which claims that there *is* information unavailable to even *sub specie philosophiae* that helps determine the precise reference of ‘lavender’, then we ought to believe that the verdict *sub specie philosophiae* is indeed defined as giving *the* best verdict. If we have determined that the rules that we have been

considering do not indicate/dictate whether a candy is lavender or not, then that is fact about language that we have discovered. If the only things that determine facts about language are the rules of language, then we have simply discovered a higher-order rule. If there is something outside of the rules of language, then some account must be given of it, and this is what I attempt to do in Chapter Four.

There is one way for us to say that we are unconfident in our assessment, *sub specie philosophiae*, of a border-area candy, where the lack of confidence is due to an epistemic and not ontological issue, without being committed to Williamson's epistemic view. Perhaps claims of unconfidence are provisional, given that we, *sub specie philosophiae*, are presently engaged in theorizing about vague predicates, and perhaps we should wait until there is some resolution of this very debate before making any claims about these cases. Possibly, rather than make any claims of unconfidence about border-area candies, we should now be silent. (This, however, will run into problems of higher-order vagueness, given that it then seems unclear when we should begin to be silent.) But given that we are engaged in the project of trying to resolve Sorites paradoxes, it is because we believe that there is *something* that we are not presently aware of that in some way clarifies the referential structure of vague terms. We know that this process is theoretical – it (probably) does not involve actual empirical tests beyond what we already know. What this response indicates is that we are confident that there will be some form of resolution of allegedly unconfident cases. This is in opposition to Wright's current view, which seems to *rest* (in a somewhat literal sense) on the assumption that these cases are *quandaries*. It strikes me that the very reason that Sorites paradoxes have achieved prominence in philosophical literature is that we are unsatisfied with the status of us not

knowing what to say about supposedly borderline cases, and that we need to find the best analysis of these cases. My optimism is not based on contingent considerations; I shall argue that the truth about supposedly unconfident cases simply is what, *sub specie philosophiae*, our best theory tells us we should say. If it is true, it would be confirmation of Frege's notion that concepts are total functions – there is one best analysis *sub specie philosophiae* for borderline cases.²² On the epistemic view, borderline cases are in fact either lavender or not lavender. On the methodology I advocate, the range for the function need not be bivalent – we might say, *sub specie philosophiae*, that one should say that one should neither call it lavender nor not-lavender. My view, then, is similar to the epistemic view in that I agree that there is something that we are not presently aware of which can help us understand the referential function of terms like 'lavender', but it is dissimilar in its not accepting bivalence at the outset (though, because of considerations in 1.12, I also accept bivalence) and in its optimism that the 'best thing to say about borderline cases' is in fact knowable, though it is not presently known, and is achievable through philosophical argument.

1.11 Higher-Order Vagueness

Which cases are the border-area cases? Is there a precise boundary between clear cases and border-area cases? At what point do our confident, normative judgments about candy color turn into unconfident, provisional judgments? We have come this far and have still not yet explored what is most troubling about vagueness: *higher-order*

²² This might be further complicated if one is pessimistic about the prospects of interpersonal philosophical argumentation and theorizing. So it might be better to talk of *sub specie philosophi*, which is the perspective of an individual philosopher. This is a higher-order instance of the discussion of individual relativism; and it relates to Blackburn's (1998) analysis of relativism, which I endorse.

vagueness (HOV). What is HOV? Just as *vagueness* is difficult to characterize, so is HOV. Most basically, the intuition that there is no sharp boundary between lavender and not-lavender seems to entail that there is a ‘border-area’ between lavender and not-lavender. But this would seem to entail that there are sharp boundaries between lavender and the border-area and between the border-area and not-lavender. HOV is the worry that these boundaries are themselves not precise. Rather than solve the problem of vagueness, the common-sensical three-region view, stemming from the NSB intuition, seems to be even more problematic than the original Sorites.

As I note in 1.9, a common way to express the non-border areas is to call them the *definitely lavender region* and the *definitely not-lavender region*, or $def(L)$ and $def(\sim L)$, for short. A second order Sorites can be formulated as follows:

- (1) $def(L_1)$ (Minor premise)
- (2) $(\forall p) (def(L_p) \rightarrow def(L_{p+1}))$ (Major premise)
-
- (3) $def(L_{1000})$

(3) is said to follow from (1) and (2) from 999 instances of universal instantiation and modus ponens. But (3) is false, because

- (4) $\sim def(L_{1000})$ ²³

What motivates the major premise of this argument? I claim, following Mark Sainsbury (1991), that we believe not only that there is no sharp boundary between lavender and not-lavender, but that there are no sharp boundaries in the whole range. I shall call this the *boundarylessness* intuition. It seems likely that underlying the boundarylessness

²³ For my purposes at the moment, I am not using *def* as a sentential operator, as it is in other accounts. If it were a sentential operator, the paradox as given would be invalid, because it would involve quantifying across a sentential operator.

intuition is the $OD \rightarrow PD$ intuition, which goes some way in explaining why there are no sharp boundaries; however, I shall argue that it is not the $OD \rightarrow PD$ intuition but rather the *indeterminist* intuition.

But first, it is important to ask how high HOV goes. Is there another region, say, of $\text{def}(\text{def}(L))$ candies? If so, is there a precise boundary between $\text{def}(\text{def}(L))$ and $\text{not-}\text{def}(\text{def}(L))$? While several writers have made comments noting the limits of HOV, no one has taken the limitations seriously enough. Fine, in an offhanded concluding comment, remarks (1975, p. 297): “our intuitions seem to run out after the second or third orders of vagueness.” However, Fine sees it as a virtue of his supervaluation account that it can handle a potential infinity of orders of vagueness. I see this as a weakness: if we, *sub specie philosophiae*, can only find grounds for a few orders of vagueness, then any logical system for vague predicates should characterize that very fact, and be limited by it. The fact that there is a limit on our intuitions of the orders of vagueness probably reveals something about the semantic structure and use of terms. Furthermore, just as the lack of acuity of our perceptual and cognitive systems indicates that there is no sharp boundary between lavender and not lavender, it should also indicate that there cannot be a highly detailed hierarchy at many orders.

How, then, should we characterize the intuitive limits on HOV? More terminology might help. It is common to define a notion of *admissibility* as relating to definiteness in the same way that possibility relates to necessity in modal logic. Namely, $\text{adm}(L) = \sim\text{def}(\sim L)$, and $\text{def}(L) = \sim\text{adm}(\sim L)$. There are two problems with the intuitive backing of these identities that are not commonly noted; both stem from the fact that no one has given a settled account of exactly what facts about language use *definiteness* and *admissibility*

correspond to. First, $\text{def}(L)$ might function independently of the notion of admissibility, as in march B, in which a region of *definitely lavender* is given prior to any consideration of admissibility. Second, it might be the case that questions of admissibility permit a trivalent structure but questions of definitivity do not. For example, if, rather than asking whether it is admissible to call each candy ‘not-lavender’ and permitting only a ‘yes’ or ‘no’ answer (as is done in A*), we permit marchers to give answers other than ‘yes’ or ‘no’, then the region in which a response other than ‘yes’ or ‘no’ is given is *not* a region in which the marcher claims it is admissible to call it ‘not-lavender’, but it *is* a region that the marcher would seem to claim is not definitely lavender. The def region seems better identified only with the region that is *not* admissibly not-lavender, and the not-def region with everything else.

This might seem to be simply a choice of nomenclature for def and adm. It is plausible to suppose that if we permit a trivalent response to questions about admissibility, then we should also claim that definitivity is trivalent, and not say that the region in which an answer other than ‘yes’ or ‘no’ is given to an admissibility question is *not-definitely*; rather, we should consider that region to be border-definite, just as it is border-admissible. This supposition seems reasonable, and is likely made by authors using notions of def and adm. However, I hope to show that higher orders of vagueness give intuitive support for claiming that definitivity is bivalent whereas admissibility is trivalent. The following argument does not purport to show that we should change our nomenclature to characterize admissibility as being trivalent and definitivity as being bivalent; the fact that our technical vocabulary does not match exactly with common-sense notions is not reason enough to change the technical vocabulary. Rather, I hope to

show to show that identifying *definitely* with *not-admissibly-not* obscures an important point concerning the limits of HOV.

Let us suppose that in a march, call it A^{**} , we ask whether it is admissible to say that it is admissible to say that the candy is not-lavender (presuming that in A^* , the marcher claimed that some candies on the left side, say, those less than 300, are not admissibly called ‘not-lavender’).²⁴ I shall follow convention and represent n iterations of ‘it is admissible that’ as adm^n , and n iterations of ‘definitely’ as def^n . Presumably, A^{**} will give us a 2nd order boundary point, say, 200, which is the highest-numbered candy which is $\sim adm^2(\sim L)$. Then there will be an A^{***} , which asks whether candies are $adm^3(\sim L)$. Presumably, this will give a 3rd order boundary point, say 100. It should be obvious where this is heading – that either there will be a final boundary point (given the uncontroversial assumptions of *neatness* and of one’s holding one’s own previous answers as admissible, as noted above in 1.9), such that for some m , and all $n > m$, the n^{th} order boundary point will be some number, say 50, or there will be no candies at all that are $\sim adm^n(\sim L)$. In Chapter Two I formalize this argument to expose a deep flaw in supervaluationist logic (and state a similar argument against the epistemic view, if *adm* is construed epistemically).

This is an important result. It shows that either there is a sharp boundary somewhere, or the entire range is not fully determined by the facts of the matter. Taking

²⁴ There could be a different A^{**} in which the question is ‘Is it admissible to say that it is not admissible to say that it is lavender.’ However, for reasons not worth explaining, A^{**} as I have it in the text relates to A^* as A^* relates to A , because both A^{**} and A^* are of the form: ‘is an answer other than the one given by you in the previous march admissible?’ For higher-order B marches, the possibilities are quite numerous, because the question asked could be the above question, or it could be: ‘is a *denial* of the answer...admissible?’ This introduces a tremendous amount of complication for possible B marches; as we shall see, the A marches are complicated enough to demonstrate my point.

HOV seriously and maintaining the boundarylessness intuition involves denying the common-sensical view that terms like ‘lavender’ are determinately true when applied to candies such as 1, but are not determinately true or false in the border-area. If there were some cases that are determinately true there will be a last such case. But if there is HOV, then there cannot be a last such case. And if one denies the boundarylessness intuition, then it seems reasonable to suppose that there is a boundary between lavender and not-lavender, just as the epistemic view says. The view that there are several ranges of determinacy, but with a sharp boundary somewhere, seems more poorly motivated than the initial bivalent epistemic view (especially in light of the considerations to be given in 1.13).

I shall now focus on the philosophical question of what could possibly indicate/dictate appropriate answers to marches such as A*** and beyond. Do we really have intuitions about whether it is admissible to say that it is admissible to say that it is admissible to say that candy 1 is not lavender? We certainly have intuitions that it is definitely lavender. But that is not the same as saying that it is $\text{def}^3(L)$. It seems that one can appeal to perception in order to claim that candy 1 is $\text{def}^1(L)$, but perceptual evidence itself would not validate $\text{def}^2(L)$. Perhaps one could say something to the effect of “one does not understand the meaning of ‘lavender’ if one does not say that candy 1 is definitely lavender”. That might validate $\text{def}^2(L)$. But can the same be said for $\text{def}^3(L)$? My claim, which I shall discuss in slightly different terms in Chapters Three and Four, is that there are *no* facts which can demonstrate $\text{def}^3(L)$; in short, this means that $\sim\text{def}^4(L)$. However, that is not to say that it is $\text{adm}^4(\sim L)$. This comes out of a more explicit explanation of what *def* really means. My supposition is that if one were given an option

to give an answer other than ‘yes’ or ‘no’ to a march which asks whether candies are $\text{adm}^3(\sim L)$ – call this march the A^{***B} march – one ought to give an answer other than ‘yes’ or ‘no’, even for paradigm candies such as 1.

It is worth noting several points here. First $\sim\text{def}^4(L_1)$ does not entail $\sim\text{def}^1(L_1)$. It would only do so if our logic for vague predicates contains an analogue of the S4 axiom from modal logic, which claims $\text{def}(P) \rightarrow \text{def}(\text{def}(P))$. My claim will be that if we are uncertain, *sub specie philosophiae*, whether $\text{def}^n(P)$, then that means that we ought to claim $\sim\text{def}^{n+1}(P)$, because if we are uncertain about P *sub specie philosophiae*, then we can be certain that P is not definitely the case. However, although definity is ultimately bivalent and lack of definity at higher levels does not entail lack of definity at lower levels, it is the case that admissibility is trivalent and that if we believe, *sub specie philosophiae*, that $\text{adm}^4(P)$, then we ought to believe that $\text{adm}^1(P)$. This is because the intuitive notion of admissibility does seem to be transitive across levels – if we claim that it is admissible to say that it is admissible to say that P , then we ought to say that it is admissible that P . This is the reason why, at higher orders, we should refrain from giving a ‘yes’ or ‘no’ answer to questions of admissibility. What I am trying to characterize are intuitive, and not technical notions, of adm and def . Doing so, I believe, does not prove that our logic must coincide with our intuitive notions, but it does reveal an important fact about Sorites series – that our intuitions about admissibility run out at a certain level.

What explains the fact that our intuitions run out at approximately the third or fourth level rather than higher or lower? (This is a question that no one seems to have ever asked.) I am unsure exactly why, but I’d like to give a preliminary explanation here. There are two possibilities. The first is that our perception is acute enough to group color

series into several ranges, which could be mapped onto the number of ranges there would be in a Sorites series in which there are three or so levels of definitely. I believe that this in itself does not account for our intuitions running out at the third (or so) level – if our perception were more (or less) acute, then it still might be the case that our intuitions would run out after that many iterations. It is also worth considering whether we have the same number of ranges for Sorites series which cover a smaller range of the color continuum, such as from azure to aqua. My suspicion is that our intuitions may run out slightly sooner, but that there would still be two or three orders.

The second possibility is that the number of ranges coincides with our interests – it makes the concept the most useful to have several ranges of analysis. Why is it that concepts are most useful if they have several ranges? My answer is that it is partly because of perceptual reasons, and partly because of theoretical ones. It is useful that speakers are able to claim that others' uses are inadmissible, and of course the range of these cases are determined perceptually; in fact, it seems that questions of whether it is admissible to call a candy 'lavender' should be determinable by casual observation, just as original questions of whether a candy is lavender. However, admissibility, even at this first level, is tied in with a more theoretic analysis of language use. I am unsure how to argue for this point. In talking with non-philosophers, I've heard a wide range of response to questions of admissibility, and it became apparent that many of the responses were connected with the individuals' theoretic proclivities. My belief is that we do make fairly unreflected judgments about admissibility, but these judgments are always threatened to be undermined by theoretical doubts. Sometimes, I tell myself that an individual who claims that candy 1000 is lavender must not be speaking English; this may be the kind of

theoretic intuition that tacitly underlies inadmissibility claims of first-order uses. But it should be noted that ‘is speaking English’ is itself Sorites-susceptible. Would it be admissible for an individual to claim that another individual is speaking English in saying that candy 1000 is lavender? I’m unsure. It does seem that it is highly dependent on theoretical beliefs about the nature of language. And I’m unsure what could ground any judgment at even higher levels of admissibility. It strikes me that in the absence of an agreed upon theory of language, there are, at present, no good grounds for making any inadmissibility claims at, say, the fourth level. At the very least, there are no *further* reasons we can give if we attempt to claim that $\sim\text{adm}^4(\sim L_1)$ that we have not already given for why we believe, say, that $\sim\text{adm}^3(\sim L_1)$. This discussion reveals that we need a more developed theory of language that accounts both for evaluations of whether objects are lavender, and also of a stratified admissibility structure, and this is a project that must be undertaken to solve Sorites. So, much of the remainder of this dissertation will be an analysis of the prospects of any semantic theory that attempts to give such an account.

One important note, however, is that if we believe that $\sim\text{def}^4(L_1)$, then a common form of higher-order Sorites will be resolved. Sorites paradoxes may be solved, as in supervaluation accounts, by claiming that the major premise is false, and by appeasing intuitions about the denial of the major premise by claiming that there is no sharp boundary between *definitely* lavender and *definitely* not-lavender. This gives rise to a higher-order paradox, in which it is alleged that there is a sharp boundary between definitely lavender and not definitely lavender. As long as there is always yet another order of ‘definitely’ to turn to, one can resolve any higher-order paradoxes. However, the consideration above that Sorites series are discrete and finite demonstrates that orders of

vagueness *will* run out. But this does not mean that there will be an irresolvable Sorites.

The argument

- | | | |
|-------|---|-----------------|
| (1) | $\text{def}^4(L_1)$ | (Minor premise) |
| (2) | $(\forall p) (\text{def}^4(L_p) \rightarrow \text{def}^4(L_{p+1}))$ | (Major premise) |
| ----- | | |
| (3) | $\text{def}^4(L_{1000})$ | |

can be shown to be unsound, not because the major premise is false, but because the *minor* premise is false, on the assumption that it is at the third order of vagueness where our intuitions run out. I hope to show in later chapters that this is a promising route to pursue in resolving Sorites.

However, this move, which relies on uncertainty at higher-orders, still has not avoided all possible forms of higher-order Sorites. The following paradox might still be given, if $U(\text{adm}^3(\sim L_1))$ denotes the supposition that we are uncertain whether candy 1 is adm^3 not-lavender.

- | | | |
|-------|---|-----------------|
| (1) | $U(\text{adm}^3(\sim L_1))$ | (Minor premise) |
| (2) | $(\forall p) (U(\text{adm}^3(\sim L_p)) \rightarrow U(\text{adm}^3(\sim L_{p+1})))$ | (Major premise) |
| ----- | | |
| (3) | $U(\text{adm}^3(\sim L_{1000}))$ | |

But (3) is false, because we should be certain that candy 1000 *is admissibly*³ called ‘not-lavender’ – it is not lavender, after all. This, I claim, is the deepest form of the Sorites, whereby, if we can resolve it, we can finally lay the problem of vagueness to rest. I shall call it the *paradox of the uncertainty of higher-order admissibility*, or *PUHA*. For the moment, I leave it open if the uncertainty should be accounted for metaphysically or epistemically. I propose, though I will not go into it in much further detail, that it is solved by formulating reasons to re-analyze $\text{adm}(P)$ as $\sim \text{def}(\sim P)$. As such we will be

certain that $\text{adm}^3(\sim L_1)$. But I maintain the claim that admissibility is transitive across levels, so I will admit that $\text{adm}^1(\sim L_1)$. And I hope that the analysis I shall give of def and adm will show why $\text{adm}^1(\sim L_1)$ is not an unacceptable result, despite its seeming counterintuitiveness. This has the effect of demonstrating that there is, ultimately, no higher-order vagueness.

1.12 Degree Theoretic Intuitions

It should be noted that certainty comes in degrees, and if so, perhaps PUHA can be solved using a degree-theoretic logic. It may, but first I would like to discuss here what might motivate using a degree-theoretic logic to analyze PUHA or any other Sorites paradox. One of the benefits of degree-theory accounts is that they are able to explain what Fine calls *penumbral connection*. Two important types of penumbral claims are what I shall call *ordinal claims* and *similarity claims*. For candies in the border-range, we still can say that candies on the left are *more* lavender than those to the right; there is something about the concept of lavender that gives an ordered array by lavender-status. (In Chapter Two I discuss the prospects for an *ordinal degree theory*; standard degree theories are *cardinal* degree theories because they claim that a real number corresponds to each candy.) Degree theory accounts for ordinal claims quite well because the degrees to which two candies are lavender sets a mode of comparison between them. Secondly, even for candies in the border range, we say that they are *similar* in color. Degree theory can give a proper account of similarity by claiming that two candies in the lavender series are similar in color insofar as their lavender-degree values are close.

As I argue above, one of the intuitions leading to Sorites is the belief that for observational concepts, there is no ontological difference without a perceptual difference. I believe that degree theory in fact respects this intuition better than any other account of vagueness on the market (except perhaps for the context-sensitivity theory such as given by Raffman and Graff, which has other problems of its own as I shall discuss in Chapter Two). Sanford, who himself advocates a degree-theory approach to vagueness, remarks (1976, p. 201): “There is admittedly something ironic about responding to the imprecision of natural language by adopting a semantics which allows infinitely precise discriminations of truth-value.” What follows is a partial vindication of degree-theoretic views from this worry.

Let’s say for the moment that candies 1 to 250 have a lavender-value of 1, candies 750 to 1000 have lavender-value of 0, and each candy c in the border area has lavender value of $1 - ((c - 250) \times .002)$. Thus the lavender-value of 251 is .998; 400 is .7; 401 is .698; and 749 is .002. Perhaps underlying Sanford’s worry is the objection that degree theory involves even more ontological differences without perceptual differences. However, I believe objection fails. For even if candies 400 and 401 are not *pairwise-indiscriminable* (which means that one cannot discriminate between them when just looking at the two of them), there is indeed *some* detectable perceptual difference. Specifically, 400 will be pairwise-indiscriminable from some patch, say 399, that 401 is not pairwise-indiscriminable from. (If 400 and 401 are indeed pairwise-indiscriminable from all and only the same ones as each other, then their degree-values should be the same. But Sorites paradoxes are premised on pairwise-indiscriminability not being transitive, so I shall assume that 400 and 401 are instances of that.) So, on degree theory, there is a

minute degree-value difference between 400 and 401, but that simply represents the minute *ontological* difference between the color of candy 400 and the color of candy 401, which is due to a minute perceptual difference. Importantly, the same can be said for candies 250 and 251. The difference between 1 and .998 (i.e., .002) is no more substantial than the difference between .7 and .698, so degree does not presuppose that there are any ontological differences that do not correspond to perceptual differences.

This is not to say that degree theory is correct; it still seems to violate the *boundarylessness* intuition. For even though the range of degree theory is a continuum from $[1,0]$, that corresponds to the potential continuum of colors of possible candies, there still is a precise boundary between 100% lavender and less than 100%, and we do not know where this boundary is. I believe that this is an important result, because it separates the OD→PD intuition from the boundarylessness intuition, and indicates that the boundarylessness intuition is motivated in a different way. What I believe does underlie the boundarylessness intuition, and Sorites in general, are intuitions about *indeterminacy*. It is not determinate (I hold, in line with common sense intuitions) exactly which candies are 100% lavender. But this is not an issue specifically dependent on boundary issues. For it is also not determinate, for example, what candy is 70%, either. It is completely unclear how the facts that we might be aware of, such as facts about light reflectance and facts about usage, can determine *semantic* normative facts, regardless of whether the semantic facts come in degree form or are bivalent. Sorites makes us notice the fact that reference is indeterminate. The epistemic view claims that the non-semantic facts do determine the semantic fact; what I say here isn't meant as a refutation of it; my objections to it will be the focus of Chapter Three. What I hope the discussion here shows

is that the intuitions that lead most people to reject the epistemic view are intuitions that vagueness is, at its basis, due to indeterminacy.

Let me finish this section by returning to the point that began the section. If we can give a degree-theoretic account of *certainty*, I claim that we can solve Sorites; we do not need a full degree-theoretic account of color properties. However, philosophy of mind (and psychology) teaches that we might be not good judges of our own mental states. Is there any *objective* (I cannot say *mind-independent*, of course) standard that indicates how certain we are about something? For reasons I state in the next two sections, I believe that there is not.

1.13 Bivalence

My argument that, ultimately, there is no higher-order vagueness, comes from several considerations. First, I reject the epistemic view (for reasons given in Chapter Three). Second, I argue that *sub specie philosophiae*, there is one best thing to say about each case (this is the argument given in 1.10). Third, I shall argue that the best thing to say, *sub specie philosophiae*, simply is the correct answer. But these considerations leave open the possibility that first-order structure, at the end of the day, will be *multivalent*. But in fact, in contexts in which language is used representationally, first-order structure is indeed *bivalent* – this is a concession to Williamson’s view. Though it will be important to consider the role contextual factors play in determining what first-order structure is most desirable in particular situations – this is a concession to those who focus on context as underlying vagueness.

There seem to be strong grounds for claiming, as Frege does, that valuations of sentences are bivalent. A simple argument in favor of bivalence comes from four claims: (1) sentences express propositions; (2) a proposition says that the world is a certain way; (3) there is one way that the world is; and (4) either it is the way that the proposition says it is or it is not. (3) may be controversial, but I cannot conceive of how it could be false. Even if one is a relativist or pluralist about our descriptions of the world, it seems that the world itself simply *is*. And to say “the world *is*” is the same as saying “the world is a certain way”. Hence I regard (3) as following from the claim that *something* exists. Secondly, the very notion of a proposition, as derived from Frege and the early Wittgenstein, is designed such that propositions have the same logical structure as the world itself, which is to say that (2) is true, by the very meaning of ‘proposition’. Propositions are simply things that say that the world is a certain way. (4) seems undeniable as well, though I do not have an argument for it; if we have a strong enough notion of a proposition, then it will follow. It seems that if one is inclined to reject bivalence, then one must actually be committed to something stronger – a rejection of the referential picture of language that (1) encapsulates. A skeptic may deny that there is representation, but the kind of relativist view I advocate claims that despite the indeterminacy which renders naturalistic semantics untenable, we need not abandon normative representational semantics altogether.

If concepts are bivalent, it would exclude a degree theory and other theories that admit three truth values. It would also deny that there is a range for which it is indeterminate whether a sentence is true or false. Perhaps, as noted above, multivalent theories can be reconstrued as analyzing levels of propositional attitudes – it’s not that we

fully believe that candy 400 is 70% lavender; rather, we have a certain propositional attitude by which we assent, with 70% confidence, that it is lavender (see Sainsbury, 1986). There seem to be two problems with such a degree theory. First, it does not seem at all apparent that we can precisely quantify our propositional attitudes as such. Why is it 70% and not 71%? Can we really make such distinctions? This might suggest an ordinal rather than a cardinal degree theory, as I shall discuss in Chapter Two. Second, even if it does explain our levels of belief about border-area candies, it still doesn't explain what the ontological status is of these candies. We can analyze some speakers' beliefs as such, but what should we, *sub specie philosophiae*, say about these cases? It would seem odd, *sub specie philosophiae*, for us to say that we are 70% confident in candy's 400 being lavender, and that we have nothing more to say about its status. What could explain us having such an attitude? It seems that it would require some account of our lack of confidence. Such an account might be considered a *degree-theoretic epistemic theory*; but it seems like it would be more poorly motivated than the original epistemic theory.

Still, the bivalence of propositions seems to not map on well to the apparent non-bivalence of the normative structure of ordinary linguistic behavior. For in practice, we do not assess speech acts in a bivalent scale, just as we do not assess styles of dress bivalently. For example, when learning language, parents can give more subtle cues encouraging or discouraging children's linguistic behavior which do not simply give full encouragement to true statements and full discouragement to false statements. Furthermore, expressing truth is not the only goal of ordinary language use, and, as noted in 1.8, norms of truth are only one part of norms of language – there also are norms of relevance, norms of appropriateness, and norms of perspicuity, to name just three.

While something like a degree of belief theory might correctly describe our degree of willingness to assent to certain claims, it does not get to the heart of the ontological issues underlying vagueness. Considerations in the preceding paragraph might lead one to believe that concepts are not bivalent, but it strikes me that these considerations are poor motivation to hold that view. For while language is a tool, it is plausible to say, as above, that one reason why some language is a successful tool is that it refers to objects in the world. *Sub specie philosophiae*, when we are analyzing whether border-area candies are lavender, we bracket off these other norms and focus solely on questions of truth, and in doing so, assume that part of what one does in uttering a sentence such as “That candy is lavender” is express a proposition, even if one is also doing many other things.

There may be another motivation, though, for believing that there is a disconnect between norms of ordinary language and norms of truth. Let’s consider for the moment a sense-data theory of perception, otherwise known as the ‘Myth of the Given’. On this theory, our sensory experience is not laden with conceptual content. Our minds then (in some way) turn our brute sensations into propositional knowledge. Of course, much of analytic epistemology in the second half of the 20th Century, culminating (perhaps) with John McDowell’s *Mind and World* (1994), has rejected this view. On McDowell’s picture, our sensory input is already in the logical form of a Fregean ‘thought’.²⁵ If so, and if (2), (3) and (4) above are true, then sensory representations are either veridical or

²⁵ In conversation, I discussed with McDowell an objection that if concepts are implicated in sensory experience, one should be able to say exactly what those concepts are, but it does not seem that we can say exactly what concepts are implicated in particular sensory experiences. McDowell responded by saying that the input is already in the logical form of a Fregean ‘thought’. I found his response unsatisfactory.

not. And if language is an expression of the propositions implicated in our representations, then language is bivalent.

It is interesting that the argument McDowell gives in *Mind and World* is a kind of transcendental argument – justification of our beliefs is possible only if our representations themselves are concept-laden. What vagueness might call into question is the very claim that our attribution of our concepts to the world is justified – for our concepts, being those of creatures with imprecise perceptual capacities, may not fit the world how it really is. The world seems much more precise than ordinary, vague, human concepts seem to allow, even if the world itself is not fully precise. To me, this is an argument favoring a return to something like the “Myth” of the Given. I shall say return to this issue momentarily, though a full examination of it would take me well beyond the scope of this dissertation.

1.14 Moderate Indeterminism, the Phenomenological Sorites, and the ‘Myth’ of the Given

What exactly are we doing when analyzing terms like ‘lavender’ *sub specie philosophiae*? ‘Lavender’ is not a purely observational term. Merely observing that an object looks lavender is not enough to prove that it is – it may only appear lavender due to unusual lighting conditions. Whether I believe, *sub specie philosophiae*, that a candy is lavender is different from whether it appears to me as being lavender.

I have discussed how the problems of vagueness are really problems of indeterminacy, and I shall advocate a solution to Sorites that involves a fairly radical form of indeterminacy. But a more common-sensical view is that vague predicates are determinate only in the border-area, but are not indeterminate for end cases. Preliminary

considerations of HOV cast doubt on that *moderate indeterminist* view, but I need to say more about why it seems right in the first place. One way might be to claim that the referential structure of concepts relates to paradigm cases. For instance, there are paradigm instances of lavender, and whether an object is lavender depends on its similarity to paradigm cases. (This may motivate a degree theory account.)

I believe that considerations of paradigm cases do not motivate a moderate indeterminist view. A question here is whether similarity of color is a brute perceptual fact, or whether it is primarily something for scientists to explain. It may be very difficult to give a full scientifically respectable account of color similarity, as color perception is multi-dimensional. But if such an account can be given, then likely we will have, in the process, successfully reduced the very notion of ‘lavender’ to something else that is presumably precise. And if that’s the case, then we could indeed solve the whole Sorites paradox, and so reference would be seen to be determinate even for ‘border-area’ cases. On the other hand, similarity to paradigm cases can be considered a brute perceptual fact, where perceivers judge whether a candy is similar in color to paradigm cases. But if the similarity relationship is based upon a brute perceptual fact, then the paradigm account does not give us any further explanation of what we’re doing when we call something ‘lavender’ than what we had already known, as questions of whether the facts of the matter determine the relevant similarity relations between two candies are no easier to answer than questions of whether the facts of the matter determine whether the candies are lavender. I conclude from this that the thought that the referential structure of terms like ‘lavender’ relates to paradigm cases might be right, but introducing the notion of

similarity to paradigms in no way justifies the view that Sorites series are determinate at the ends but not in the middle, which is precisely what HOV calls into question.

A different motivation for the common-sensical moderate view is best brought out by discussing the solution to Sorites paradoxes given by John Burgess (1990). On Burgess's view, an object is lavender if it would appear lavender to most normal observers under normal conditions. This seems to me to be correct, and it may provide a way to resolve Sorites paradoxes without entering into the deep metaphysical/epistemological debates I note in 1.6 and elsewhere. Or, it could be used as a basic framework for explicating a relational realist view of vague predicates, and might be used for explicating a form of the relational epistemic view – there is so much information we could obtain how who normal observers are and what normal conditions are, as well as what normal observers would judge, that we could never know exactly where the boundary is; such an epistemic theorist would claim that there is a precise fact of the matter about who is a normal observer, and what normal conditions are.

But it might also be grounds for a common-sensical three-value view – there are some conditions that are definitely normal, some observers who are definitely normal, and some amounts of observers that are definitely 'most', even though all three are vague. If that's the case, then we would have a three-valued structure. This is Burgess's own view. The complete details of his account aren't important here – I'm just using his account of color terms to give some motivation for the view that they are determinate on the edges and indeterminate in the border-area. Burgess attempts to solve Sorites paradoxes by claiming that while it is vague what 'most' means, and also it is vague what normal conditions and normal observers are – and so color terms are vague – these terms

are not ineradicably vague at all higher-orders. For example, ‘most’ is first-order vague – it can apply to any number between 50% of speakers and all speakers. But ‘most’ is not higher-order vague – there are clear boundaries to ‘most’. Burgess gives similar accounts for the notions of normal speakers and normal observers, showing that they are first-order, but not higher-order, vague. So on Burgess’s view, there is a limited amount of higher-order vagueness while there will remain cases that are clearly lavender; he argues that this solves Sorites paradoxes. If his account is correct, it is an account of how there is indeterminacy for border-area cases while the outer-areas are determinate.

This is an ingenious approach to Sorites, but it has a problem. It relies on it being determinate whether objects appear lavender to individuals. This relates to the discussion noted in the previous two sections. Consider this *phenomenological Sorites*:

- (1) Candy 1 appears to me to be lavender
- (2) For any c , if candy c appears to me to be lavender, then so does $c+1$
-
- (3) Therefore, Candy 1000 appears to me to be lavender.

But this (3) is false, because candy 1000 does not appear to me to be lavender. Of course, at issue here will be the status of the major premise. It seems that whether a candy appears lavender to me is just a factual judgment about me, and not a normative semantic one involving the meaning of vague predicates. If that’s the case, it seems like the paradox can be used as a simple *reductio* of the major premise. It cannot be true, even if we are unsure which candy forms the boundary. This is revealing in a number of ways. First, it shows that we presume that there is a fact of the matter about whether candies appear lavender to an individual, whereas, by the indeterminist intuition I discuss above, we do not presume that there is a fact of the matter about whether a candy is lavender.

Second, given that we believe that there is a fact of the matter about whether we perceive something to be lavender, it is unproblematic to view the Sorites as a *reductio* of the major premise of the phenomenological Sorites.

Problems arise, though, if we think seriously about where the last candy is that appears lavender to me. It is not simply a matter of it being impossible for two candies to look the same without one appearing different than the other – that is trivially true.

Phenomenological matters are quite complex. First of all, even within one single candy, it is likely that the color does not appear uniform – different parts of it will reflect light differently, and different edges may appear different upon examination. Additionally, there are different categories of indistinguishability: pairwise indistinguishability, in which two candies look to be the same when examined next to each other with no other candies to measure against, and presumably with a neutral background. My belief is that pairwise indistinguishability is not best captured by the claim that observers would judge that they appear to be the same color, because even for a single candy, if an observer attends to it carefully, it will likely appear differently at different moments; rather, I believe pairwise indistinguishability would be better captured by a claim that observers, if they did not know which of the pair was lower numbered (i.e., closer to lavender), would have an even distribution of claiming that one patch is more lavender than the other if they were forced to make such a judgment. Next, there may also be indistinguishability within a series; test subjects are allowed to compare patches to other patches in the series to determine whether there are other patches which, upon investigation, can show that only one of two patches is identical to a third patch. There is also a level of indistinguishability *sub specie philosophiae*, in which we are presumably

allowed scientific experiments, an extended amount of time to give an answer, plus an extended amount of time to theorize about color properties.

I include the foregoing not to prove anything in particular, but rather to show that the major premise in the phenomenological Sorites does not seem easily motivated by phenomenological considerations. But is it false? Is there a candy that will appear lavender to me but the one next to it will not? It strikes me that the phenomenological paradox is not caused by slippery-slope considerations; rather, like with the original Sorites, it is caused by considerations of indeterminacy. I don't have a particularly good argument for this, but my intuitions are as follows. I perceive a manifold, and in my field of vision, there is an area of a certain color. While my perceiving it as such does not come prior to my understanding a language, it does come prior to consciously ascribing to it any particular color concept. Saying that it 'appears lavender' is just a label one can put on a perceptual experience, just as calling an object 'lavender' is just a label on that object. Why does it seem like there's a fact of the matter on the left side of the continuum that candies there appear red, whereas for border-area cases, there may be no fact of the matter? Because saying that something 'appears lavender' is just a label that we ourselves *actively* put on it; there is no objective fact of the matter about the proper label for our subjective experiences. (Given that I believe that objective and subjective are not in principle incompatible, I don't intend the previous sentence to be analytically true.) If we are *deciding* to call something 'lavender' on the left edge, it makes the objective legitimacy of the naming of the whole range problematic. There's no fact of the matter that we're choosing; we're simply choosing labels. My conclusion is that labels for

phenomenological experiences will be chosen, and will not easily be determinable by any naturalistic theory of mind and language.

I am denying that there is an objectively determinable category of ‘appears lavender’. So a third potential motivation for the moderate indeterminist view is that radical indeterminism seems to undermine much of what we hold dear about language use. My view seems to be a return to something like the ‘Myth’ of the given. And it seems to fall into the very worry which motivates McDowell to claim that perceptual experiences already have conceptual content – that if perceptual experiences lack conceptual content, then in explaining our beliefs, we cannot give justifications of the beliefs. Certainly, if sense experience is not concept-laden, some account must be given of where concepts come from. My belief is that it is not entirely *active* choice – there is some mental structure which automatically attaches concepts to sense perception most of the time. However, there still seems to be some room for freedom of thought, too. I believe that nothing short of a grand solution to questions of freedom and determinism in philosophy of mind is needed to resolve these issues. Needless to say, a grand epistemological theory and a grand theory of mind are beyond the scope of this dissertation.

What are the costs of denying that there are any relevant similarities between two individual lavender experiences? It seems that there are good evolutionary grounds for believing that there will be relevant similarities. And if there are relevant similarities, that’s just what concepts are supposed to capture – relevant similarities are in the world itself, and perceptual similarities are merely a *reflection* of that, and our concepts are simply what we use to represent similarities. This would be a direct realist view.

Alternatively, humans evolved with an ability to distinguish between colors; there are facts of the matter about our biology that make our experiences similar. This is an evolutionary relational realist position, but not a direct realist position. Either of these two, I believe, is a fairly strong argument that there *must be* color categories.

However, I believe that a view in philosophy of mind that rejects the primacy of concepts, and still accounts for evolutionary success (and other kinds of success as well), is not so implausible. In particular, it might be argued that we have a faculty that judges relative similarity of aspects of our sensory input and our memories, and this is, in principle, separable from any linguistic faculty of ours. (This claim might be bolstered by an appeal to other animals.) If one eats a certain berry that caused illness, one can notice similarities between that berry and other berries, and one will know to avoid berries that look similar to the first one. A story about this faculty can be told, I believe, without relying upon any notion of color concepts. The similarity is a brute perceptual similarity, and the color concept, if there is one, is only added later. I hope this kind of account may make indeterminacy of language less unappealing. And I shall dedicate much of Chapter Five to showing how the indeterminist view of language I advocate still gives us what we need in order to have normative evaluation of language – guidance and governance.

It is interesting that McDowell himself also rejects fully ('bald') naturalized semantics; perhaps there is more in common in our views. Space would prevent me from giving a full analysis of McDowell's view and a more developed epistemological theory, though at present I do not have a full epistemological view worked out. My own general view is that we will need a successful account of similarity before we can assume that the best way to describe our phenomenology is that there are conceptual categories of

perceptual experience. I do believe in a certain holism of the mental, which might be a position similar to McDowell's. However, I disagree with McDowell that the best way to describe this is that reasons take over the whole space. I believe that a *lack* of full reason pervades the whole space.

Though I agree with Sartre that choice ultimately underwrites normativity, I believe that we need not have an existentialist view that choices are totally free – we seem to *tend* to choose in certain ways. And if we have a tendency to do choose, it casts doubt upon whether the choice are truly free. Hence it is not a pure existentialist freedom of thought. But neither is it a determinist argument that we are unfree. It is motivated by what I believe is a common sensical view that claims that if we are giving an account of language and the mind, we need an evolutionary story, a biological story, a historical-linguistic story, and a developmental story, but we also need a story that accounts for conscious choices that are not presently fully understood within naturalistic philosophy of mind. There is biology, but there is also choice. I shall return to these issues in Chapter Five, though a full analysis of the seemingly contradictory notion of our 'tendency' to choose certain things is beyond the scope of this dissertation.

1.15 Conclusions: Vagueness as Noticed Indeterminacy, and the Is/Ought Gap in Semantics

I hope I have placed in view the range of possible views concerning realism/anti-realism/relativism. Achieving consensus on how to solve Sorites will require consensus on which of these views is correct. I also hope that I have shown that vagueness is tied closely to other large unresolved debates in contemporary philosophy, and this serves as an explanation for why Sorites has been so hard to resolve.

I have argued that the common-sensical moderate indeterminist view of vagueness, that vague predicates are determinate at the ends of the continuum, but are indeterminate in the middle, is cast into serious doubt by considerations of higher-order vagueness. Intuitions concerning the lack of sharp boundaries are not what ultimately underlie Sorites paradoxes. Mark Sainsbury's view (1991) is that vagueness is *boundaryless classification*. We have intuitions that there are no boundaries of color concepts, and we have intuitions that we use color concepts to classify objects in the world into color categories, but I believe that because of Sorites, no coherent account of boundaryless classification can be given. However, I have argued that the focus on boundarylessness misconstrues the real source of vagueness: it is not that we want to classify without boundaries; rather, we want to classify without there being any determinate way to classify. On my view, vagueness is *noticed indeterminacy* – we simply do not pay attention to problems of indeterminacy for non-vague cases, but, as I believe higher-order vagueness shows, the problems are still there.

I find it somewhat surprising that the kind of radical indeterminist view that I advocate has not been more fully discussed in the literature. It is partly motivated by Quinean and Wittgensteinian concerns. Furthermore, I believe that there is an important connection between problems in naturalized semantics and problems in naturalized ethics. Famously, it is difficult to derive ethical 'ought' statements from 'is' statements. I believe that the problem in semantics is similar – we can know as many non-normative facts about light reflectance, about our perceptual mechanisms, about our dispositions to use words, but from these facts alone we cannot deduce normative facts about how words like 'lavender' *ought* to be used. I believe this kind of consideration can be used to reject

fully naturalized kinds of semantics. However, the analogy with ethics may provide us with some direction – we have not abandoned ethical theorizing despite the is/ought gap, and in Chapters Four and Five, I apply lessons learned in ethical normativity to resolve the issues concerning semantic normativity. The solution to Sorites that I present in Chapters Four and Five, the *normative choice* solution, is motivated by an analogous view about the role choice plays in determining our ethical norms.

In Chapter Two, I discuss current views on offer in the literature. I explain more specifically how the considerations above apply to supervaluation and degree theory. I have not yet given any arguments against Graff's context-sensitivity view, and I have not adequately discussed Wright's most recent view. I also have not given an argument against Williamson's epistemic view. In fact, much of this chapter can be seen as an argument favoring Williamson's view over other views currently on offer. But in Chapter Three, I formulate an argument against it that relates to wider questions of realism and anti-realism. In Chapter Four, I discuss the notion of rules, in light of Wittgenstein's rule-following considerations, to more fully support my own view that rules of language do not determine the normative structure of language. I shall also discuss rules of ethics, and show how facts of the matter do not determine ethical norms, and discuss the lessons learned from ethics. In Chapter Five, I formulate what I shall call the *normative choice* solution to Sorites, and I show how it is internally coherent, and succeeds where other indeterminist views fail. I hope to show how choice still provides the two things that we need – guidance and governance – for there to be normative evaluation and justification of our practices.

I should also note what I believe the dissertation leaves out, and could be considered future projects. First, a further examination of the issues in epistemology and philosophy of mind concerning whether perception is concept-laden would be desirable. Second, lack of space will prevent me from detailing a full meta-ethical picture that shows how there are ethical norms despite indeterminacy; I have developed such an account and I believe that it can be used to validate semantic norms. Third, it is possible that the connection between linguistic and ethical norms is more than analogical – it may be that epistemic norms are themselves norms only insofar as they are pragmatically useful, in which case norms of language are simply one kind of behavioral norm that are subject to ethical analysis. Fourth, the importance of the notion of choice demands a more complete discussion within debates of freedom, in regard to contemporary analytic philosophy of mind, existentialism, and Kantian philosophy. I hope I say enough in Chapter Five to show that my solution is a plausible route, but I do not believe that I give a complete enough account of the consciously active aspects of our mentality and aspects which are beyond our conscious choice. Such an endeavor would be a life project, not a dissertation project.

It is the business of philosophy, not to resolve a contradiction by means of a mathematical or logico-mathematical discovery, but to make it possible for us to get a clear view of the state of mathematics that troubles us: the state of affairs before the contradiction is resolved. (And this does not mean that one is sidestepping a difficulty.)

- Wittgenstein, *Philosophical Investigations*, §125

CHAPTER 2. SOME CONTEMPORARY APPROACHES TO SORITES

In this chapter, I examine several of the solutions to Sorites on offer. I have already laid the groundwork for my analysis of each in Chapter One, and in two cases – degree theory and supervaluation – I believe that Williamson (1994) has already analyzed them quite well. My goal in discussing these two views will be to bring out aspects of these views that are important for development of either my own view or my criticisms, in the next chapter, of the epistemic view. In two other cases – Delia Graff's view and Crispin Wright's most recent article – the views are new and have not yet received attention in the literature, and I will discuss them in more detail. I discuss all four of these views because I believe that each has something right about it, and I attempt to bring out what I believe is right. (Because Timothy Williamson's *epistemic view* deserves extended analysis, I reserve all of Chapter Three for it.)

2.1 Graff's Interest-Relative Account

2.1.1 Graff's "Shifting Sands"

The view expressed by Delia Graff in "Shifting Sands" (2000) is an odd one, and it resembles my own in several ways. In this sub-section, I explicate her view, and in the next I shall argue that it is not an adequate account of vague predicates.

On Graff's view, the precise reference of vague predicates is interest-dependent – it falls under the category I describe in Chapter One as an *interest-relational* view. On this view, what 'red' refers to is determined by the interests of the person using the term, and by the context in which the term is used. Thus, vague adjectives are relational. Graff explains (22):

- That car is expensive is to be analyzed as meaning:
- That car costs a lot, which in turn is to be analyzed as meaning:
- That car costs significantly more than is typical.

Graff goes on to argue that the expression 'significantly more than is typical' shows that truth conditions for 'expensive' relate to some norm of car cost; and furthermore, whether a car is *significantly* more costly than the norm depends on the interests of some individuals. Hence, 'expensive' is relational.

According to Graff, predicates like 'tall' and 'expensive' are in fact precise, because interests are precise, even though they often do not appear to be. This positions Graff to resolve Sorites paradoxes because it both shows why we should reject the major premise, and why we were originally inclined to accept it in the first place. To show that interests are precise, she discusses the question of how much coffee will satisfy her interests. She writes:

I take it ... that when I desire some coffee it is just a brute fact that there is a least amount of coffee of which it is true to say that it will satisfy my desire. Anything less will not do. I know that many philosophers will protest; they will say, "But how could it be that your desire for coffee is like *that*?" I say, given that a teaspoon of coffee is not enough to satisfy my desire, how could it *not* be like that! Moreover, I have an explanation for why my desire seems that it is not like that. There should be nothing surprising or doubtful in the suggestion that we have inexact knowledge (in Williamson's sense) of the satisfaction conditions of our desires. I have never bothered to figure out exactly how the enjoyment I get from coffee maps on to the different amounts of it I might drink. (28)

Hence Graff's view is a kind of epistemic view. On her view, bivalence holds (either she will be satisfied by the amount of coffee or not), and there will be unknown boundaries of our terms. One of the important aspects of her account is that she provides an explanation of what it is we are ignorant of – something that Williamson does not provide. She provides an account of the connection between use and meaning – we have some knowledge of our interests, which determine the meaning of the term, and that explains why we use the terms well enough. (I shall grant to Graff that if 'enough coffee to satisfy my desire' is precise, then so are whatever interests that are related to the norms for other vague adjectives.)

It might be argued that Graff's account is better suited for a degree-theoretic logic, given that interests may be satisfied to different degrees. Graff expressly states that the initial aspects of her theory are to be consistent with non-classical logics. But rather than accept the degree-theoretic logic, on Graff's account, there is a notion of 'fully satisfied', such that those cases that satisfy her interests in the range between 0% and 100% (exclusive) do not count as satisfying her interests fully.

Her account might even be extended to give not only an explanation of why there is first-order vagueness, and why it is not a problem, but also an account of why there is

higher-order vagueness, and why it is not a problem either. For there is a potential range of values which might be the precise boundary for how much coffee will satisfy her interests.

Because on Graff's account, the extension of a predicate is relative to a language-user in a context, an important factor is *salient* to the individual in the context. *Salience* becomes a prominent notion in Graff's account, though I will argue in the next subsection that the notion introduces more problems than it resolves, and that an interest-relative account is best left without such a notion. On Graff's view, if we are looking at two adjacent paintings in a Sorites series on the red-orange continuum, differences between other paintings won't be salient at that moment. Graff writes:

Whatever standard is in use for a vague expression, anything that is *saliently* similar, in the relevant respect, to something that meets the standard itself meets the standard; anything saliently similar to something that fails to meet the standard itself fails to meet the standard. Put another way, if two things are saliently similar, then it cannot be that one is in the extension of a vague predicate, or in its anti-extension, while the other is not. If two things are similar in the relevant respect, but *not* saliently so, then it may be that one is in the extension, or in the anti-extension, of the predicate while the other is not. One reason for requiring that the similarity be a *salient* one is to block the absurd conclusions that would otherwise follow by Sorites reasoning, since any two dissimilar things can be connected by a similarity-chain. (14)

If there are two members of a Sorites series that are saliently similar (to some language-user in some context), then it cannot be the case that one is in the extension of the relevant predicate and the other is not. Though her initial use of 'salient' related to perceptual salience, Graff makes an ingenious maneuver using this notion: she defines *saliently similar* as being equivalent to *the same for present purposes*, which in turn means, as she explains: "I will say that two things are the same (in a certain respect) for

present purposes when the cost of discriminating between them (in that respect) outweighs the benefits.” (25)

So, on Graff’s account, if we are looking at two consecutive members of a Sorites series, then the two will be saliently similar, and so the precise boundary will not be between those two. However, Graff continues:

Even if there is some cost associated with any discrimination we might make between similar heights, ages, etc., the cost of making the discrimination somewhere does not outweigh the benefits of making it somewhere. The boundary between those differences that are significant and those that are not will try to locate itself, so to speak, at a place where there is little or no resistance. (28)

The boundary will shift depending upon what part of a Sorites series a language-user is looking at – hence her title, “Shifting Sands...”. She concludes:

As an adherent of classical logic and bivalence, I believe that Sorites sentences are false and that “sharp boundaries” claims are true. On any Sorites series for a vague expression, I believe that somewhere in the series (not where we’re looking) there is an object that possesses the property expressed by an utterance involving a vague expression right next to an object that lacks that property. I refuse, however, to call the proposed boundary between the property possessor and the property lacker a *sharp* boundary since as I have stressed, this is but a metaphor, and I have as much right to the metaphor as does the proponent of gaps or degrees. I would cash out the metaphor in the following way: the boundary between the possessors and the lackers in a Sorites series is not sharp in the sense that we can never bring it into focus; any attempt to bring it into focus causes it to shift somewhere else. (33)

2.1.2 A Response to Graff: But the Sands Don’t Really Shift...

I have several criticisms of Graff’s account. I argue that it is difficult to show that truth conditions for sentences shift in the manner Graff indicates, and that language is more public than her account has it. But one can retreat to accommodate these criticisms

and still have an interest-relative account. But in the end, I argue that the very notion of an *interest* is more problematic than Graff deems it.

First, it is unclear exactly what Graff means by ‘costs’, ‘benefit’, and ‘discrimination’ when she writes that “the cost of making the discrimination somewhere does not outweigh the benefits of making it somewhere” (28). Presumably, the benefit to which she is referring is the benefit of not being caught in a Sorites paradox, but it is difficult to measure what kind of benefit that is. It is also unclear what the costs are – there are practical costs in measuring, for example, precisely how much redness is in a painting, but Graff uses the notion of ‘discrimination’ ambiguously. In one sense, it means an activity to have the difference become perceptually salient, and in the other sense, it just means attributing a predicate to one thing and not another, hence discriminating between the two things. If that’s all it takes, then there seems to be little cost in saying that the boundary is somewhere.

But then, it seems like one can simply say, while looking at two adjacent members of a red-orange Sorites series, that one is red and the other is not, even though for practical purposes, one would have a great deal of difficulty discriminating between the two items. The reason why Graff rejects this is that (1) terms are interest-relative, and that (2) if two adjacent items are salient, it won’t be in our interests to discriminate between them, because the costs will outweigh the benefits. But I am unsure why (2) is true, and why Graff even needs it to be true for her account to succeed. What is needed is a more developed account of interests, and of costs and benefits, than Graff supplies.

To further this response to Graff, consider the following passage:

Imagine an eccentric art collector who reserves one room for her paintings that contain just red pigments, and reserves another room for her paintings that contain just orange pigments. One day she is presented as a gift a painted color spectrum ranging from primary red on one end, to orange on the other. She resolves to cut the canvas in half. Now if she cuts without thinking, perhaps in a state of mad excitement because she is so eccentric, she will most likely cut in just the right place—by which I mean that once the halves are re-framed and hung, she will still be able to truly proclaim that her paintings containing just red pigments are in one room, and that her paintings containing just orange pigments are in another. Although the right-hand edge of the painting in the red room is extremely similar to the left-hand edge of the painting in the orange room, their similarity is not *salient*, so the boundary between red and orange may occur between them. If the decision about where to cut is labored, in contrast, the collector will likely find herself unable to locate the boundary between the red and the orange, the pigments on either side of any proposed cut being too similar—and when the decision is labored, *saliently* similar—for one to go in the red room and the other in the orange. (15-16)

This is truly odd. Why is it that only if one is mad can one cut right at the boundary?

Can't one who is reasonably-minded simply say, 'I need to cut the painting somewhere, so I should simply cut it here'? If one's has an interest in mind in preserving a red/orange split, it seems that salience should play little role in determining what is red and what is orange. Graff includes the notion of salience so that she deny that there is a salient sharp boundary, but it seems that she does not give enough reason to support the claim the cost of discriminating between two consecutive items that one is looking at is always greater than the cost of not doing so (on the assumption that the cost of not discriminating *somewhere* is greater than the cost of not doing so). Furthermore, it seems that even if one's own interests do shift based on momentary context – for example, the precise amount of coffee required to satisfy me one minute might be slightly greater than the previous minute – it seems that for her own account to work, Graff needs some story of

how our shifting interests correspond to our shifting *perceptual* input, rather than with, for example, our ever-changing internal states.

Perhaps Graff may be willing to abandon the eccentric paint collector example because of these considerations; but there are other problems even if the notion of salience does not depend on perceptual input as in that example. Another reason why Graff's example may strike one as odd is that it seems that truth conditions for a sentence should not continuously vary based upon the continued experiences of one particular individual. So Graff needs more of a metaphysical account of what exactly is shifting. Perhaps one who wants to have a coherent practice should never have such a red/orange policy about pigments in the room, because it is never knowingly satisfiable, given the problems with borderline cases. This supports my own inclination that if the attribution of the term 'red' to an object can only occur if the extension of 'red' varies as such, then there really is no property of redness, and that we should move towards eliminativism for the use of the term.

This is simply an unargued rejection of Graff's interest-relative account; if we *do* accept that terms are interest-relative, and that interests do change with context (though perhaps not quite in the same way that Graff claims), then it would entail that the truth-conditions for the application of vague terms do change quickly. So, I should give some more motivation for my rejection of the interest-relative account. This consideration comes not from metaphysics, but from philosophy of language: We have a well-entrenched intuition that whether an object is red or not should not depend upon something as contingent as the state of mind of one particular individual. I believe this connects to the notion of *sub specie philosophiae* that I introduced in Chapter One. It is

from this perspective that the linguistic intuitions of all language users are compiled so as to determine the single referential function for all terms. The very notion of *sub specie philosophiae* doesn't exclude vagueness, for it may claim that the single referential function includes borderline cases for which there is no correct answer (apart from the response that there is no correct answer) to whether it is red or not. This would account for the *stability* of vague terms – we have intuitions that we must appeal to a single standard of correct use of a term, rather than have standards shift as such.

The notion of *sub specie philosophiae* is based upon a limiting assumption that we are all speaking one common language. I do not argue for this assumption, and I am unsure whether it is ultimately tenable. The assumption is used because some account is needed to explain the success of communication. However, if one who believes that truth conditions vary by individual-in-a-momentary-context can provide some account of how communication is still feasible despite shifting, then the motivation for claiming that referential structure for a term of a language is determined from *sub specie philosophiae* might be undermined, and the intuitions noted in the last sentence of the previous paragraph may be overcome. I shall leave this as an open question, though I cast more doubt upon it below.

One further reason for considering the view *sub specie philosophiae* is to make sense of *disputes*. Perhaps for a given case, one person says that some item is red, and another says that it isn't. On Graff's account, it could be that they are both right, given that the context for each individual might be slightly different. Her account may not allow for an explanation of how two individuals might resolve a dispute. This kind of objection is a common response to relativist views, though I am unsure how much force it

has in general if some account of how individuals communicate using different idiolects can be given. Specifically for Graff's form of relativism, perhaps there is some level of overlap between the interests and contexts of two individuals. She could then give some account of what happens in disagreements, in the following form:

Person A: This is red-for-my-context

Person B: This is red-for-my-context

Both A and B: Our contexts have a great deal of overlap. So, given that this is red-for-my-context, and that our contexts have overlap, this is probably also red-for-your-context.

Such an account is not easy to give, and would require much further development. The *sub specie philosophiae* perspective simplifies matters, but because of the possibility that a relativist such as Graff can give some such account, it is not a non-starter.

A further reason why the sands might not shift so quickly is that it seems possible for the whole spectrum in a Sorites series to be in view at one time; the whole range is perceptually salient. Given that this seems empirically plausible, then it might be a simple reason to reject Graff's claim that on an interest-relative view, the sharp boundary is never where one is looking, and so her whole introduction of the technical notion of salience, defined in terms of costs and benefits, does not do the work that she wants it to. Another case is where a Sorites paradox can be given where *none* of the members of the Sorites series are perceptually salient to anyone. As point of fact, almost all Sorites series are this way, given that they are, in their most common form, written as thought experiments in philosophy papers. In what way does the reference change with changing interests in such instances of Sorites?

Graff notes (28-29) that her account of positive vague adjectives, such as ‘expensive’, ‘tall’, and ‘old’, as meaning “having significantly more ___ than is typical” does not apply to another kind of adjectives – negative vague adjectives, such as ‘cheap’, ‘short’, and ‘young’ – these all should be analyzed as “having significantly less ___ than is typical. I am unsure whether ‘young’ is positive or negative – it may mean “having significantly more youth than is typical”. Semantics is a tough business – there are many many semantically relevant features of each of these terms. However, I will grant Graff the point that these terms might be interest-relative in a way that semantics can ultimately make sense of.

Still, there are other classes of adjectives that Graff does not give an account for. How should ‘red’, for example, be analyzed? Should it be analyzed as being ‘significantly redder than items typically are’? But it is difficult to see what could be substituted for ‘redder’ within the analysis that would give a non-trivial interest-relative account of ‘red’. The best possibility, I believe, would be ‘significantly closer to paradigm red than is typical’. Perhaps this succeeds, but there are other vague adjectives where it is difficult to how such an account would work. One example is ‘typical’. Could it be analyzed as ‘being significantly more typical (or normal, ordinary, or whatever) than is typical’? In Graff’s defense, perhaps some such analysis might work – perhaps it can be analyzed as ‘having fewer eccentricities than one typically does’. Of course, what ‘one typically does’ means in this context needs spelling out – it could mean: median; mean; median/mean over all items in all possible worlds; median/mean compared to what is expected. But it might be the case that just what ‘typical’ means in this context is relative

to the interests of an individual or group of individuals using a term, and so the spirit of Graff's account might still be preserved.

However, as Graff admits, she does not account for vague *nouns*. She says that an analysis of nouns would have to be done on a case-by-case basis. She writes:

There is a good explanation for why the treatment of vague adjectives runs smoothly, while extension of the proposal to vague nouns seems strained. The reason is that it is a semantic feature of adjectives that they are associated with some dimension of variation—one needed for the formation of comparatives such as 'taller', 'older' and 'more expensive'. But it is not a semantic feature of nouns that they are associated with a dimension of variation. (29)

However, it is much more difficult to provide an interest-relative view of nouns. But there are many similarities between the use and semantic structure of vague adjectives and vague nouns. These similarities may thus cast doubt upon her primary claim that vague adjectives are interest-relative if one believes that nouns are not. This can be seen especially with proper nouns. Consider Mt. Baldy, the highest of the San Gabriel Mountains. First of all, it seems that if you hadn't heard of Mt. Baldy before reading the previous sentence, you do now, because reference to it has been passed by a causal chain, *a la* Kripke's account. It is hard to tell how an interest-relative view could be given for 'Mt. Baldy'. It would probably involve rejecting Kripke's causal theory; certainly it would involve rejecting a form of externalism which claims that meanings are entirely outside the head. It seems that for objects such as mountains, we defer to the view *sub specie philosophiae* to determine reference.¹ Furthermore, it is also highly questionable how the issue of salience could come in.

¹ One note is that one way to characterize Graff's account is to consider it that she is providing *sub specie philosophiae*, a formula in the meta-language for determining the truth conditions of each vague

Notably, the same kinds of examples can be given for adjectives. Consider someone who has never drunk wine, but who has conversed with wine-drinkers and thus acquired use of the term ‘oaky’. It seems this person may be able to understand a Sorites paradox on that term, and it is unclear how truth conditions for the person’s use of the term ‘oaky’ depend on the individual’s interests. Now, one might dispute that there can be Sorites paradoxes simply on paper without the referent there. But these paper-bound paradoxes do seem to make sense, despite the context not being perceptually salient. This is evidence that both the notion of salience isn’t required for an interest-relative view of vagueness, and also that interests are not necessarily changeable in any fast-shifting manner.

All these objections above are demands that a much more detailed semantic account is needed to make any interest-relative view work, and also to show that Graff’s particular account seems to be less plausible than other interest-relative accounts that depend less on a notion of salience. Still, these objections leave open the possibility that some interest-relative account might work.² However, there is a further objection looming concerning the very notion of ‘interest’. Graff argues, using the coffee example above, that there must be some precise amount of coffee that satisfies her present interests. She deduces this by using Sorites-style reasoning: one teaspoon of coffee is not enough, and (let’s say) two full cups of coffee are enough, and so there must be some

object-language sentence in each context – they are all relative to the individual’s interest at that time. So there will be something in common between all individuals, though the particular reference for each will differ. But it is unclear whether this is enough to satisfy what we want out of a common language.

² Graff does support her account of interests by appealing to the semantic theory of linguist Chris Kennedy (30 - 33). But this simply provides a technical model for the semantics of vague expressions, and does not provide real clarification of the notion of interest as demanded in the objections above.

exact amount at which point it is enough coffee. It is a *reductio* on the hidden premise that for any amount of coffee c , if c is not enough coffee to satisfy my present interest/desire, then $c+\epsilon$ is not enough, either, for some small ϵ . Why is it that this kind of argument seems more acceptable as a *reductio* on the major premise than other Sorites arguments? Examination of this question proves to be very revealing.

I am unsure exactly what motivates Graff's argument here, but I speculate that the difference between Sorites paradoxes for 'satisfies my desire for coffee' and 'is expensive' is that it one might think that one's desires are indubitably real and non-relational³ – they are mental states – whereas something is expensive only insofar as it relates to the interests of an individual. Presumably, Graff does not want there to be a higher-order account of what it is to satisfy one's desire such that whether something satisfies someone's desire is relative to some other desire. The 'brute fact' (28) that there is a least amount of coffee that will satisfy her desire is not a relational, or conventional, fact.

However, it will be difficult to base a semantic account on desires. First of all, there are independent grounds within philosophy of mind to go eliminativist concerning desires; I will not follow this line of criticism here. Second, for vague terms like 'expensive', 'red', 'heap', what matters is not desires (as a particular mental state) but one's *present interest*, an expression that Graff uses frequently in her paper. So in order to give an account of the relational aspect of vague predicates, we will need an account of *present interest*. This needs much further explanation. Likely, what is in one's present

³ Though it might appear that 'desire', as Graff uses it, is relational (I desire coffee), it is rather that she is in a mental state such that it would be satisfied by a certain amount of coffee; but 'coffee' should not be taken to be part of a desiring relation; it should be understood as I have a desire-of-a-certain sort.

interests is also dependent upon long-term considerations about one's interests in general. Is this amount of coffee enough to satisfy my present interest? Let's imagine that on the one hand, I need a pick-me-up, and this desire is motivating me to drink coffee. On the other, I have a GI condition such that I need to limit the amount of coffee I drink. Is it in my present interest to drink any coffee? I am unsure. But what can possibly provide an answer for me? A scientific account of how much pleasure drinking the coffee will give me, minus how much pain the GI problems will cause? What might such an account look like?

The interest-Sorites can potentially be used to call into question the very notion of an interest, rather than to prove that interests are precise. On Graff's view, for some n and just about *any* non-zero ϵ , n liters of coffee will not satisfy my present interests but $n+\epsilon$ liters will. This seems to strain the very notion of 'satisfy my present interest'; Graff is correct that we are not fully aware of our own interests, but I cannot fathom that any amount of examination of an individual can lead to such a precise demarcation. I approach 'satisfy my present interest' just like 'looks red' in the first chapter – both seem to be potentially precise, because they seem to be dependent upon real mental states (or something reducible to mental states). But does this intuition hold up when painstaking questions of precision arise? It is hard to imagine that when $\epsilon = 1$ nanoliter, there is anything about me that makes it such that I will be satisfied with $n + \epsilon$ liters of coffee, but not n . This, to me, is an argument for the vagueness of mental state terms, and with it, an argument for the vagueness of 'satisfaction of interest'.⁴ In clear cases, we say of

⁴ See my "Are Minds Vague?" (Hiller 2000).

ourselves ‘I am satisfied’; but this does not show that ‘am satisfied’ is a precise, or even coherent, concept.

I am unsure how persuasive this is; perhaps actually drinking a cup of coffee will help (in more ways than one). What actually happens when one drinks a full cup, and goes (allegedly) from being less than fully satisfied to fully satisfied? Is there a real fact of the matter that there is a point at which one becomes fully satisfied? Even if one did have full self-awareness, it is unclear whether one would be self-aware of being 100% satisfied. In other words, a full description of our mental states may not include the notion of full satisfaction; hence there are good grounds for being eliminativist about the term, even if one is not an eliminativist about mental concepts in general.

The preceding discussion was interpreting ‘interest’ as a name for a mental state; but the notion of ‘interest’ is a value-laden, normative, ethical term. (I am unsure how Graff herself intends it.) In the next chapter, I argue that because of the is/ought gap, one cannot reduce normative truth conditions of a sentence to something non-normative; but Graff’s account is notable because it does not attempt to do that, insofar as her own intention with ‘interest’ is to denote something more than simply a mental state. I am not sure that this aspect of her view makes it any more favorable, however, unless the implicit pragmatism can be worked out in more detail.

Furthermore, even at the level of normative terms, arising out of moral skepticism there are independent reasons for denying the existence of interests, and this would lend support to the eliminativist approach to the interest-level Sorites. Secondly, ‘interest’ is a normative term, and so we will need some account of ethical normativity, which is no easy problem to confront. While it might in fact be a feature, rather than a flaw, of

Graff's account that it ties semantic normativity to ethical normativity, it leaves it in need of much more explanation.

2.2 Wright's Intuitionist Account

2.2.1 Wright's "On Being in a Quandary"

Crispin Wright has written much on the topic of vagueness since 1976, and is a significant reason why the field of vagueness has become such an interesting and exciting one. Wright relates Sorites paradoxes to the realism/anti-realism debate, and recently to relativism, and much of his work on vagueness informs, and is informed by, his work on meta-ethics and other questions of truth and objectivity. Though much of his work deserves comment, I shall restrict my attention here to "On Being in a Quandary: Relativism, Vagueness, Logical Revisionism" (2001). In it, he defends the use of intuitionist logic to resolve the Sorites paradox. In this subsection, I shall explicate his view, and in the next, I shall argue against it as a solution to Sorites.

Wright differentiates between two different kinds of relativism. There is what he calls *indexical relativism*, according to which truth conditions of statements in a discourse are relative to some individual or group. An example of indexical relativism is Gilbert Harman's moral relativism, as expressed in Harman/Thomson 1996. On this view, Wright claims, there are no real disputes, and it thus mischaracterizes our ordinary understanding of disputes within a discourse (51-52). Though I believe Wright's characterization of this view is too quick (it is a version of the brief anti-relativist argument I give above against Graff), his purpose is not to give a full argument against it but instead to clear taxonomical space for his own favored view, which he terms "True

Relativism”. On this view, there may be disputes about a common claim such that two parties have conflicting views but are such that neither is in error (53). (As we shall see, though this is how Wright initially characterizes the view, in the end, he does not want to say quite that.) I myself am unsure if “True Relativism” is a good name for this view – it is unclear what is taken to be *relative* to what on this view; it seems it is more like a pluralism than relativism. But I shall follow Wright’s own usage in what follows.

Wright also ventures to resolve a problem concerning intuitionist mathematics’ rejection of classical logic, and hopes to provide an intuitionist solution to Sorites, following Putnam (1983). Sorites may be seen as a *reductio* of the major premise

$$(1) \quad (\forall p) (L_p \rightarrow L_{p+1})$$

It can also be seen as a *reductio* of the following premise, which, for reasons I discuss in Chapter One, is a better candidate for the inference step in Sorites paradoxes.

$$(2) \quad \sim(\exists p) (L_p \ \& \ \sim L_{p+1})$$

The benefit of intuitionist logic is that it does not permit one to use a *reductio* on either (1) or (2) to the “unpalatable existential”:

$$(3) \quad (\exists p) (L_p \ \& \ \sim L_{p+1})$$

Instead, using intuitionist logic, a Sorites-based *reductio* of (2) will only give us:

$$(4) \quad \sim\sim(\exists p) (L_p \ \& \ \sim L_{p+1})$$

which is not logically equivalent in intuitionist logic to (3) because double negation elimination (DNE) is not intuitionistically valid. The goal for Wright, then, is to find independent motivation for disallowing DNE, and to provide some understanding of how (4) should be interpreted.

Wright's view is a form of epistemic view, because he argues that borderline cases are cases in which we do not know the truth-status. But Wright is sure to differentiate his view from the epistemic views of Williamson and Sorensen, who maintain bivalence, and with it, an undesirable (for Wright, and for myself too) form of realism. I believe that Wright is successful in the goal of separating the epistemic view from bivalence, but is not successful in providing an adequate alternative resolution to Sorites paradoxes.

Wright introduces a principle of *evidential constraint* (59):

(EC) $P \rightarrow$ it is feasible to know that P

Though Wright does not believe that (EC) holds for all propositions P, he does accept it for propositions such as that expressed by 'this is red' because, on his anti-realist account, the nature of the concept of redness is such that if an item were red, it is because it looks a certain way, and hence if an item is red, in principle it should be feasible that we should know that it is such. Wright uses the (EC) as an alternative to Williamson's kind of epistemic view – Williamson claims that it may be impossible to discover whether some borderline cases are red, when they are. At issue for me in what follows is not whether (EC) is correct. Though I believe that even for terms like 'red', (EC) fails, I will not criticize that aspect of Wright's account. What is important is whether Wright succeeds in providing a coherent account of vagueness that accepts (EC).

On Wright's view, borderline cases are those in which we are in a *quandary*. He specifies quandaries as follows (92):

A proposition P presents a quandary for a thinker T just when the following conditions are met:

- (i) T does not know whether or not P
- (ii) T does not know any way of knowing whether or not P
- (iii) T does not know that there is any way of knowing whether or not P
- (iv) T does not know that it is (metaphysically) possible to know whether or not P⁵

Wright further notes that (92):

The satisfaction of each of these conditions would be entailed by

- (v) T knows that it is impossible to know whether or not P

but that condition is excluded by Quandary as we intend it – a quandary is uncertain through and through.

Wright doesn't explain it this way, but (v) and (EC) lead to contradiction⁶:

If it is impossible to know that P for some P (as entailed by (v)), then, from (EC) and contraposition, $\sim P$. One substitution instance of (EC), as Wright himself notes (p. 59, fn.

15), is:

$\sim P \rightarrow$ it is feasible to know that $\sim P$

If it is impossible to know that $\sim P$ (as from (v)), then, from this instance of EC and contraposition, $\sim\sim P$. Even if we cannot deduce P from this using intuitionist logic, we can conclude $\sim\sim P$, and $(\sim P \ \& \ \sim\sim P)$ is still a contradiction. Hence, either (v) or (EC) should be rejected, even using intuitionist logic, and given that we have resolved to keep (EC), we reject (v).

⁵ David Sanford writes (personal correspondence): “Does this definition have more clauses than necessary? It seems to me that (iv) entails (iii) which it turn entails (ii). Moreover, (ii) conjoined with ‘whenever one knows something, there necessarily is some way one knows it’ entails (i). Why not use just clause (iv)?” I think Sanford may be right about this. I am unsure whether anything rides on the independence of the four clauses (apart from theoretic conciseness).

⁶ In what follows I use ‘it is impossible to know’ to mean the same as ‘it is not feasible to know’.

Let's imagine a borderline situation for the proposition P. One person, A, says P, another person, B, says \sim P. What should we (*sub specie philosophiae*) say about this situation? It seems like, if *we* believe that it is a real borderline case, then both A and B have committed a cognitive error. However, on Wright's view, though we cannot claim that A and B are *not* cognitively blameless, if we meet the conditions of being in a quandary, we are also not in a position to say that either has committed a cognitive error. All we can prove that it is not the case that A and B do not have a cognitive shortcoming (Wright 60). This is where the rejection of DNE comes in. The double negation is simply an expression of our being unable positively say that they are wrong, because we are in a thorough quandary about the status of borderline cases. That lack of knowledge is not inconsistent with one who asserts P, and is also not inconsistent with one who asserts \sim P. Insofar as Wright's view is relativist, it is because of this.

This might seem like it violates the notion of a concept as a function, as Frege understands it, and as I discuss in Chapter One. If a concept is a function, then it cannot be the case that we should permit A to say that an object falls under the concept and permit B to claim that the object does not; there must be one correct answer. However, given that Wright's view is at core an epistemic view, he is not arguing that there are more than one correct answer to whether a certain item is red; his view entails just that we withhold judgment on anyone who makes a positive judgment. We are not in fact permitting the answer either A or B gives; rather, the answers are not inadmissible, and this is not the same as affirming that it is admissible. I shall return to this issue in the next subsection.

So, for quandary cases, we do not make any assertions about the case that does not begin with a negation, where the negation is understood intuitionistically. We do not make any positive assertions about the case. This view is close to what I refer to in Chapter One as Pyrrhonism. However, it differs in from Pyrrhonism in that we do make some negative assertions; we deny the major premise.

2.2.2 A Response to Wright: On *Ceasing to Be in a Quandary*

I believe that as far as this all goes, Wright is correct. The logic is impeccable. But there are two major categories of problems for this account. The first is how it treats higher-order vagueness; Wright does not give any account of it in the paper.⁷ Is there a first candy at which point we are in a quandary? That would seem odd; at least, there doesn't appear to be. (And we would have a hard time explaining how examining one candy puts one in a quandary but examining the one next to it does not.) It may be unclear to us what we know. Is it a quandary what the bounds of the quandary cases are? Such a quandary quandary case would be one in which (substituting 'I do not know whether or not P' for P in (i)):

(Qi*) I do not know whether or not I do not know whether or not P.

Given that there are four conditions for being in a quandary, there are fifteen other statements that one would hold when one is in a quandary about whether one is in a quandary.

⁷ Elsewhere (1992a), Wright argues against the existence of higher-order vagueness, and this may be the reason why he does not give an account of it here. However, I believe that Heck's (1993) rebuttal of Wright's argument successful restores the coherence of higher-order vagueness.

On Wright's view, is there any difference between (Qi*) and (i)? From what it says, it might be the case that I know that P, but I do not know that I know it. This would involve a rejection of the KK principle. Would Wright endorse such a rejection? He can accept (i), because it might be the case that I do not know P even though P is true. (What he cannot accept, under (EC), is that I *cannot* know that P even though P is true.) Given the anti-realist motivation behind (EC), it does not seem like one can know that P is true without knowing that one knows it. This is not a knock-down argument against this possible treatment of higher-order vagueness; it is recommended by the motivating consideration that the rejection of the KK principle is normally an externalist move, and Wright's account of 'red' using EC is extremely internalist. Still I have not proven that there is a logical inconsistency in (Qi*).

However, even if one could provide some such account for the higher-order vagueness of quandaries, it will still raise the specter that there will be one last item in the Sorites series which does not put us in some kind of (possibly higher-order) quandary, whereby the next one does. It seems that our confidence in our knowledge is gradual; what decreases in the center area is our confidence in our assessment, though not necessarily our knowledge. This criticism will reappear in analyzing Williamson's view in the next chapter. Another possibility would be for us to be at some level of quandary even for Candy 1; this is a radical move, and seems to go against the spirit of (EC) – if we are in a quandary about everything, then (EC) has no application.

Putting aside the concern over higher-order vagueness, there is another problem with Wright's account. As above, on Wright's view, withholding judgment on a claim is not the same as saying the claim is admissible. But do these two states lead to the same

practice? If so, is it feasible to claim that there is a difference between them? In what follows, I will defend Wright's account by showing that there is indeed a difference in these two states; however, this comes at the expense of demonstrating that Wright's account does little (or nothing) as far as advancing our understanding of Sorites.

Assuming that there is no higher-order vagueness, there seem to be three possible responses that we can give, *sub specie philosophiae*, for any member of a Sorites series: That it is L, that it is not-L, or that we are in quandary about its status. Let me continue an analogy that Wright uses concerning an actual quandary – Goldbach's conjecture. It seems true that we all are in a quandary concerning Goldbach's conjecture. What should we do if someone claims to know the truth about it? Given that we are in a quandary, we cannot say that the judgment is inadmissible, though we should not also claim that it is admissible. Here's why this view of Wright's is correct, though I am unsure whether Wright himself would accept my characterization of the problem this way⁸:

We can imagine a number of possibilities where someone claims to know that Goldbach's conjecture is, say, true. At this stage, we should withhold judgment on the claim, as dictated by Wright's view. But we should not claim that the person's assertion is admissible, either. What we should do is ask the person why she believes that it is true. One possible response is that the person simply claims that it seems true to her, and can

⁸ The reason I am unsure whether Wright would accept my characterization is that I have said definitively that once we have determined that there is no reasoning behind the individual's assertion, we can claim that he is wrong. But this is not entailed by anything that Wright says. The difference is that I employ the notion of *sub specie philosophiae* to do some work: if we claim to be at the best state of information as possible, and we claim to be in a quandary, then if another person gives an affirmative answer, we should analyze the person's reasons. If the person adds little or nothing to our understanding, then we should say that she is wrong. If the person does add something to our understanding, then we should change our own view, and cease to be in a quandary. This dialectic arises out of a debate in ethics concerning relativism; I incorporate Blackburn's analysis of it here.

provide no further reason for her belief in it.⁹ Here, we can indeed suppose that there is an error in the person's *reasoning*, though not necessarily with the person's *assertion* (by itself) that Goldbach's conjecture is true. For we know what *might* count as reason to believe in Goldbach's conjecture – a proof of it that meets certain accepted procedures. For even if we do not know how to get at the truth, we have a framework of knowing what procedures should and should not grant one the standing of knowing the truth. The foregoing is not a violation of (ii) of the definition of a quandary,

(ii) T does not know any way of knowing whether or not P

because it is just a case where

(K) There are certain ways that we know do not produce knowledge whether or not P.

(K) is not precluded by the statement that we are in a quandary, and there is no reason to suppose that the intuitive reasons why we should say that borderline cases fit conditions (i) through (iv) cast any doubt on (K).

On this defense of Wright's view, we withhold judgment from the assertion, but we do not say that the assertion is admissible; rather, we press forward to determine the individual's grounds for believing in the questionable proposition. What if our claimant to knowledge claims that she knows that Goldbach's conjecture is true because she has a proof of its truth using standard methods of proof utilized by mathematicians? How should we react there? We should analyze the proof. We can come to agree that it is true,

⁹ My argument might be susceptible to cases of people like Ramanujan, who was able to intuit the truth of formulas impeccably. I am unsure what to say about these kinds of cases. Though Ramanujan was then often able to produce a proof of the formula, Sanford suggests (personal correspondence) that it seems farfetched for someone to claim that a formula is true, and then not produce any reason for it. In such cases, we have reason to doubt that the individual has understood the formula.

and, hence cease to be in a quandary. Given that, for all we know, Goldbach's conjecture is true, we should not claim that an *assertion* of its truth is inadmissible. Quandaries leave open the possibility that one might come to know the truth, even if, at the given moment, one does not know, and one does not even know how to find out whether or not it is true.

I use this analogy to then ask: what should we say when faced with A and B above – two individuals who claim to know the truth about a borderline candy? We should look at their reasons. If A says “just look at it; it's red”, we should not count that as satisfactory, for a reason analogous to the reason why we should not accept the flippant response that Goldbach's conjecture just seems true. On the hypothesis that we are in a quandary, we have already examined the candy, and we have perhaps surveyed other individuals' responses to the candy, and we have not deemed that we know its status; one more person's assertion about it will not change our beliefs about its status, and we should claim that the person has used faulty reasoning and should abandon her belief that, say, the candy is lavender, and retreat (at the very least) to the view that it seems lavender to her, but that she does not know whether it is lavender.

What of the situation analogous to the claim that one has a proof of Goldbach's conjecture – a claim that one has a proof about the truth-status of borderline cases? Right now, it might be the case that Wright is correct, and that borderline cases are quandaries – we do not know whether they are true, and we do not know whether it is possible to know. (We do not know how to resolve Sorites paradoxes.) But that's not to say that we should not approve of any claims for the truth-status of borderline cases; *one might give a philosophical theory of vagueness that does answer the question*. Wright's account simply places us at the point where we should *begin* philosophical/semantic analysis of

vagueness. We should not simply rest (as I say in Chapter One) with the thought that we are in a quandary, if we believe we are in a quandary.

Given that one rejects (v), one should not give up on the resources of philosophical/semantic analysis to answer problems of vagueness. We do not now know the status of borderline cases, but this does not entail that we should not endeavor to find out. So, I believe, Wright's view is, at best, a claim about the current state of vagueness scholarship. While he might be right, it is a disappointing view. I am unsure what it ultimately adds to current scholarship, except as a means to make room for an epistemic conception of vagueness that is not a realist one. It provides me with extra motivation (not that I needed any) to try to resolve Sorites and get us out of this crazy quandary.

There is a further point to be made. Let's consider another example of a quandary. I am at a party, and I hear a rumor that there four floobs in the room. Is this true? I don't know. Frankly, I don't even know what a floob is. But I hope to find out. But for the moment, it seems as if I am in a quandary. I don't know the truth value of any sentence "that's a floob", and I don't know how to find out, etc. But I don't even know if the person who told me the rumor is just pulling my leg – that there is no such thing as a floob, and the person just wanted to play a party game with me by putting me on a 'wild floob hunt'. If this is so, it is not the case that there are no floobs in the room; rather, the word floob is semantically defective – it has no meaning or truth conditions. As such, the sentence "That is a floob" turns out not to express a proposition at all. Hence it is not the case that I am in a quandary, as Wright defines it. This is because Wright writes of an individual T being in a quandary concerning a *proposition* P. I am merely uncertain about the *sentence* "that's a floob".

This is more than simply a technical point about the wording of Wright's account of a quandary; because what we need for Wright's account to even apply to vague sentences is to assume that for sentences such as "That is red", the sentence is not semantically defective and that it does indeed express a proposition. This is no idle criticism because on some other accounts of vagueness, such as supervaluation, sentences such as 'C is red' do not express propositions.

In his paper, Wright makes reference to his notion of a "minimally truth-apt discourse", as developed in his *Truth and Objectivity* (1992b). In defense of truth in ethics, Wright argues in favor of a minimalist conception of truth and meaning. This is in a Wittgensteinian spirit, though he does not accept the form of "quietism" that he claims McDowell attributes to Wittgenstein by which *all* our first-order practices are beyond metaphysical dispute. Setting aside the subtleties in Wright's account in (1992b), it still seems as if even if we were to replace 'proposition' with 'sentence' in Wright's account, it would become immediately less appealing. Because it puts us at the place where we must begin doing philosophy of language, which is what I have argued above. The quandary definition would be Wright saying "I don't know the truth value of the sentence 'That is red', and I don't know any way of knowing it, and I don't know if it's possible to know it, etc." That might be true, and it might be true for everyone right now. But that's just the place where we're at when we are first presented with the Sorites paradox. We're not sure what's to do. But it's also where the study of semantics (and metaphysics and epistemology) begins. So while Wright's account, as modified with the replacement of sentence for proposition, may be commendable, it is merely a confession that we need to do semantics. We should not rest with it.

2.3 Degree-Theoretic Approaches

Approaches using degree theory have been suggested by Sanford (1975, 1976), Machina (1976), Edgington (1997), Peacocke (1981a), and others. I believe that the prospects for this kind of solution have been given proper treatment in Williamson (1994), but I shall make a few remarks here that go beyond what Williamson writes. In Chapter One I argue that degree-theory approaches to vagueness accomplish what no other proposed solution does: they respect what I call the OD→PD intuition – the intuition that there cannot be a relevant ontological difference between two items without a corresponding perceptual difference. I believe that this rebuts a common objection to degree-theory approaches, such as from Tye (1994, p. 191).

Surely it is absurd to suppose that there is some single hair addition that divides the bald from the borderline bald, that changes the degree of truth from 1 to a precise degree less than 1. Indeed this seems to me no less absurd than supposing that some single hair addition transforms the bald into the non-bald (as would be the case were ‘bald’ a precise predicate).

First, the ‘surely it is absurd’ form of argument seems misplaced; one owes a view that one opposes some reason why one believes that it is wrong. For there seem to be two different objections to degree theory approaches: that they entail that there are sharp boundaries when there cannot be any, and that they assign precise values where there should not be. If the intuitive prohibition against sharp boundaries arises from the OD→PD intuition¹⁰, then, as I argue in Chapter One, degree theory approaches

¹⁰ The intuition that there cannot be an ontological difference between two items – in particular, it can’t be the case that one is red, and the other is not – without there being a perceptual difference between the two.

successfully meet this objection; at the very least, contra Tye, they meet the objection better than approaches that suppose that ‘bald’ is a precise predicate.

Another way to phrase this response to Tye is to examine a powerful but slightly misleading argument given by Horgan (1994, 173-174), of a forced-march Sorites:

For the Sorites series of baldness statements $B(0)$, $B(1)$, ... $B(10^7)$, the argument goes as follows.

- (A1) Consider the true statement $B(0)$, together with its right neighbor $B(1)$. What are the possibilities concerning the semantic status of $B(1)$? Allow as many different possible kinds of semantic status for $B(1)$ as you like – e.g. (1) true; (2) false; (3) neither true nor false; (4) neither true, nor false, nor neither true nor false; (5) indefinite whether true or false; (6) indefinite whether true, false, or indefinite whether true or false; ...; etc. No matter how many such possibilities there might be (even infinitely many), either $B(1)$ has the *same* semantic status of $B(0)$ itself – viz., truth – or else $B(1)$ *differs* from $B(0)$ in semantic status. But if $B(0)$ and $B(1)$ differ in semantic status, then there is a sharp semantic boundary between them – which is incompatible with the robustness of genuine vagueness. Hence $B(1)$ is true.
- (A2) Consider the true statement $B(1)$, together with its right neighbor $B(2)$. [Etc....] Hence $B(2)$ is true.
- (A 10^7) [...] Hence $B(10^7)$ is true.

Horgan employs a technical notion of *robust vagueness*: “Let an expression E be *robustly vague* if there is nothing in our actual semantic norms that sanctions any single candidate-precisification of E as correct, over and above various other candidate-precisifications” (1994, p. 162). I call Horgan’s forced-march argument misleading because it is not because of the robustness of vagueness, *per se*, that there cannot be a sharp semantic boundary between $B(0)$ and $B(1)$. Robust vagueness, on its own, does not entail that there cannot be a sharp boundary between two consecutive items – perhaps the two items have several possible valuations (this qualifies as robust vagueness), but they have slightly

different possible values (that qualifies as a semantic difference, or, in Horgan's severe-sounding phrase, a *sharp semantic boundary*). What degree-theory approaches do is soften the boundary, on the grounds that there really is a very slight semantic difference, even between two items that look identical when viewed pairwise (though, by stipulation, do not look identical to all and only the same other patches when viewed pairwise – see my 1.13). This seems quite reasonable on its own.

However, it seems that the real trouble for degree-theory approaches is the very problem of robust vagueness, prior to consideration of the forced-march argument – it seems impossible to motivate any particular assignment of sharp values to members of a Sorites series. It should be noted that this is the same problem that faces *all* proposed solutions to Sorites currently on offer¹¹ – the benefit of degree-theoretic accounts is that they have the resources to resolve the forced-march argument, whereas others face that difficulty in addition to the one presently under consideration. To preserve something like degree-theory, some account would be needed to describe how the degree-theoretic semantic model relates to our use of vague terms.

Interestingly, one such proposal could be adapted from Graff's view – degree of truth corresponds to degree of satisfaction of interest for the individual to whom the predicate is related to. Graff herself, an adherent to bivalence, does not accept this, but it seems like she has supplied a framework for such a theory. Let's assume that there is some real comparative relationship, along the dimension of relative satisfaction, between mental states that one would be in if one were to drink varying amounts of coffee

¹¹ Except perhaps for Horgan's own "transvaluationist" view (1998), which I do not discuss here.

(varying between, say, one teaspoon and one 16oz cup). What a degree-theorist could do is map customer satisfaction to values between 0 and 1. The problem for this view may be first brought out by considering that at some point, the satisfaction-value will become greater than 0. But which point is it? Graff argues in her own context that it must be at some point, and that it is reasonable to assume that because we are not fully self-aware, there can be some ignorance on this issue. The thing to note here is that though this problem is brought out by looking at a the case in which the truth value becomes greater than 0%, the same problem can be brought out by asking where the point is where the truth value becomes something greater than 42%. I contend that this is a problem no more or less worrisome than the case of the point at which the value goes to something greater than 0%. But worrisome it is... What could determine when I'm 42% percent satisfied, rather than 41.8? I agree with Graff that we are not fully aware of our interests, but the question is, what could we become more aware of that could guide us in the answer? This echoes the concern from 2.1.

Another objection is to ask what the *meaning* of '42% true' is. I understand truth in a way as expressed in Chapter One: a sentence is true if the world is the way it says the world is. What would it mean to say that the world is 42% the way that the sentence says it is? I simply cannot understand what this could be. What the degree theorist would need is a theory of truth whereby degree of truth applies to degree of some standard, such as degree of assertibility, degree of belief, or degree of justification. Another possibility is to consider that the motivation to reject bivalence is better explained as a failure of a sentence to express a proposition, rather than a proposition being something less than true and more than false. Following this, we may say that a sentence can express a proposition

to a degree of 42%, and this might correspond to a propositional attitude towards the sentence (See Sainsbury, 1986). This may be a possible avenue to pursue, but I'd want to see the whole theory fleshed out some more. (For reasons not worth giving here, Sainsbury's account fails; see Williamson 1994, Chapter 4.)

Another problem for degree-theoretic accounts is that there are seemingly too many of them: Sanford's differs from Machina's and they both differ from Edgington's – assignments for connectives differ on these accounts. Of course, one would need to look at the details of each account to determine which seems best. However, it is my suspicion that there may be different accounts that are equally good; but it cannot be the case that more than one model is the correct one for our terms. This urges that rather than have a particular degree theoretic semantic model, we instead need a solution to Sorites paradoxes that respects the motivating intuition that some vague sentences seem truer than others but captures this intuition without a specific formal approach.

There are arguments that Williamson makes against degree theory that I believe are not successful. In particular, he argues as follows:

Suppose that p is true to the same degree as q . Thus the first and second conjuncts of $p \& q$ match the first and second conjuncts of $p \& p$ respectively in degree of truth. By generalized truth-functionality, it follows that $p \& q$ is true to the same degree as $p \& p$. Since $p \& p$ is true to the same degree as p , $p \& q$ is true to the same degree as p . Now imagine someone drifting off to sleep. The sentences 'He is awake' and 'He is asleep' are vague. According to the degree theorist, as the former falls in degree of truth, the latter rises. At some point they have the same degree of truth, an intermediate one. By what has just been argued, the conjunction 'He is awake and he is asleep' also has that intermediate degree of truth. But how can that be? Waking and sleep by definition exclude each other. 'He is awake and he is asleep' has no chance at all of being true. (1994, p. 136)

In response to this argument, we should note that not all degree theorists maintain truth-functionality. Sanford (1975, p. 30): “The appropriate value of a conjunction is thus not determined by just the values of the conjuncts. It makes a difference whether or not the conjuncts are logically related in certain ways.” And it seems that there are reasons to not be committed to truth-functionality. In particular, we can say that truth-functionality fails only in certain instances – when, for example, the truth values of two component sentences are dependent. Hence, we should have a logic that accepts truth functionality except for some cases. It is unclear what the costs of bracketing off these kinds of cases are. And neither does it make the approach seem overly ad-hoc: there are good general grounds for claiming that the semantics for tautologies and contradictions (which form the two prominent classes of exceptions to truth-functionality to Sanford’s account) should be given different logical treatment than empirical statements. Approaches that bracket off the exceptions to truth-functionality might make things more difficult in determining the truth-value of complex sentences, but that’s not any reason to abandon it.

One further possibility that might avoid some of the aforementioned problems for degree theory that is mentioned in Chapter One is that there could be an *ordinal degree-theory* rather than a cardinal one – that, rather than truth values being given to sentences, we could simply place them in order. One variant of this view would be that some line is placed such that all the sentences in the order beyond that line are true; but this approach would seem to fall prey to Tye’s criticism above. The other variant would be to claim that there is no such thing as 100% truth; there is only comparison of truth value. I believe such an approach, though radical, is worthy of discussion. If we can still make inferences using such a method, and be justified in believing certain statements (if, for example, the

statement is ordered closer to the true side than its negation), then we might have a theory of truth which provides all that is needed out of truth. For example, we can assign a value to $A \rightarrow B$ such that it is true if B is at least as true as A, and false otherwise.

This view respects the judgment that I make earlier, in the discussion of a degree theory based in Graff's view, that placing any values on any member of a theory would be a stipulation by some individual, and not fully determined by the stipulation-independent facts of the case. Such an account would maintain penumbral connection, one of the primary benefits of using a degree theory (though one which supervaluation shares).

Williamson objects that sentences like “‘It is wet’ is truer than ‘It is cold’” will be vague in many cases, and hence there is a problem for the ordinal degree-theory. Sometimes, say, on cold, dry days, the sentence will have a determinate truth value (false). But on other days, such as cold, rainy, ones, it will be difficult to place the truth value of it. It will not do to have a ‘coldness’ truth scale that is separate from the ‘wetness’ truth scale. Having separate truth scales for each predicate would eliminate the use for a truth predicate altogether – we could instead just talk about where each sentence lies on a scale relative to other sentences of its kind; and this seems to have limited use.

One possible way for an ordinal degree theory to work would be to claim that cases where it is difficult to compare the truth value of coldness to wetness are simply cases with equal truth value. Williamson does not seem to consider this possibility; perhaps he does not because this view faces the following problem. Let's assume that based upon the evidence we have:

(E) “It is cold at time T” is equal in truth-value with “It is rainy at T”.

Now, let's assume that after ϵ seconds, it gets a fraction of a degree colder, such that:

(F) "It is cold at time T" is less true than "It is cold at T+ ϵ ".¹²

But now, how should we compare "It is cold at T+ ϵ " with "It is rainy at T"? It seems like the same reasons that would motivate us to claim (E) would lead us to claim that

(E*) "It is cold at time T+ ϵ " is equal in truth-value with "It is rainy at T".

But, (E), (F), and (E*) together violate the transitivity of truth value.¹³

One response would be to claim that "It is cold at time T+ ϵ " really has a slightly higher truth value than "It is rainy at T", if we wish to hold steadfastly to (E). However, it will be arbitrary whether to hold steadfastly to (E) and reject (E*), or instead to hold steadfastly to (E*) and claim, instead of (E):

(EX) "It is cold at time T" has less truth-value than "It is rainy at T".

Given that this choice is arbitrary, for the ordinal degree theory to work, it must claim that is some fact about the relative truth values that is beyond our ken. But given this, we might be better off simply having an epistemic theory.

One further possibility, suggested by Sanford (1993), is that there may be a supervaluation account by which each level of valuation is a degree-theoretic one, rather than a bivalent one, as in Fine's version of supervaluation. On this view, there is no single value attributable to each member of the Sorites series. However, it is unclear,

¹² Sanford comments (personal correspondence) that not all cases in which it gets a fraction of a degree colder does it make (F) true. For example, it getting slightly colder on the South Pole in winter does not change the truth value of the sentence "It is cold now". The cases I wish to consider for the sake of this argument are only those in which (F) is true.

¹³ There is a comparable problem in the literature in ethics; the problem of marginal differences, as given by Ruth Chang 1997. This argument concerns the relative values of different items. I find this problem especially compelling because it can be interpreted not just a problem of values, but a problem for preferences. Chang's own solution, that there is a fourth comparative term (apart from $<$, $>$, and $=$), *parity*, is unsatisfying to me. I do not have room to discuss this issue further.

then, exactly what the semantic status of each member of the series would be. I shall discuss this view further in the next subsection, as it is better considered a supervenient view than a degree-theory view.

2.4 Supervenience

Like for degree theory, I believe that Supervenient treatments of Sorites have been given proper enough treatment in Williamson (1994), but I shall make a few remarks here that go beyond what Williamson writes. I hope to show in this section that there are internal technical problems with the view, and secondly, that it still does not capture our intuitions concerning the nature of vagueness.

Keefe (2000) defends supervenience, and I shall respond to some of her arguments in what follows. It is worth noting that, like Williamson and his epistemic view, Keefe claims that

I believe that, unfortunately, there is no straightforward argument for the correctness of this view... The case for it must be made by showing the success with which it fulfils the tasks facing any theory of vagueness... I claim that on this cost-benefit method of assessment, it does vastly better than its rivals. (2000, p. 152)

Though it is frank, I don't find this kind of justification particularly inspiring. If all formal solutions to Sorites paradox have some faults, then why accept any formal solution at all? One of the problems with formal solutions is that they often are stated at an abstract enough level that one could make a slight alteration in the formal machinery and not be able to adjudicate which of the two solutions is better. This is an objection I level against degree theory above, and a similar objection can be made to supervenience as well. What is needed most is an account of the intuitions underlying the formalism; only after they are on the table, then can we decide upon a formal structure. Keefe herself

expresses sentiments similar to these (see her Chapter 2), and does not give a formal account of the supervaluationist view she believes is best. Because she is no friend of formalism, she sees this as a benefit of her account. She argues (59-61) that there may indeed be more than one formal solution consistent with her account, but that that is no flaw. Still, a supervaluationist view without the formalism seems no better than a supervaluationist view with the formalism – we need to give some account of what it is about our linguistic practice that makes one or more of the formal supervaluationist possibilities feasible, and do so we must be presented with the formal account.

Williamson argues (1994, Chapter Five) that though supervaluation preserves LEM, several other important features of classical logic – conditional proof, argument by cases, and *reductio ad absurdum* will be excluded under supervaluationist logic. He is right in this, but I wonder whether this is such a problem, just as I disputed whether the degree theory's rejection of certain aspects of classical logic can be bracketed off. The aspects of classical logic that fail to hold will occur only in a restricted range of cases. As such, it will not lead to general abandonment of those principles, and it is unclear just what the costs are. If other considerations show that supervaluation is the right approach to vagueness, then perhaps those considerations will also show that those aspects of classical logic ought to be abandoned for arguments involving vague propositions.¹⁴

In Chapter One, I argue that supervaluationist accounts fail to resolve problems of higher-order vagueness. I argue that either there is a sharp boundary somewhere, or the entire range is not fully determined by the facts of the matter. There will be a sharp

¹⁴ A similar point is made by Keefe (2000), 181.

boundary if, at some point, the metalanguages become precise, and exclude a certain portion of the series (say, from 1 to 50) from being admissiblyⁿ (for any n) lavender. The problem comes out in a statement by Fine: He writes (1975, p. 297): “Anything that smacks of being a borderline case is treated as a clear borderline case. The meta-languages become precise at some, but no pre-assigned, ordinal level.” As Williamson demonstrates, resorting to this argument does not solve the problem. Every patch is indistinguishable from the patches adjacent to it; thus there is no way that patch p could smack of being a borderline case and p+1 not smack of being borderline: “The process has no stable limit short of including all cases.” (Williamson 1994, p. 161) The admission that at some point, we can get a perfectly precise metalanguage is an admission that there really is a last candy which is ultimately, definitively lavender. If this is so, it is unclear why we need the heavy formal machinery of supervaluation—why not just admit that there is a last lavender candy at the first level? Williamson urges that it is better simply to have an epistemic theory without the complex supervaluational structure.

One way to phrase this problem is to claim that supervaluation must face a higher-order Sorites paradox that shows that either:

- (1) for some m, and all $n > m$, $(\exists c) (\text{def}^n(L_c) \ \& \ \sim \text{def}^n(L_{c+1}))$, or
- (2) for some m, and all $n > m$, $\sim(\exists c) \text{def}^n(L_c)$

In English, either the final leftmost endpoint, at some order of vagueness, will be (1) some candy, say #200, which is defⁿ lavender all the way up, or (2) there are no candies, even candy 1, which are definitelyⁿ red for any n greater than some m. I believe, following Williamson, that (1), which Fine accepts, makes supervaluationist no better than the epistemic view. Instead, I believe that something like (2) might be feasible;

though as far as I know, no one in the literature has considered it seriously – Williamson does not consider it, even for the purpose of rejecting it.

Keefe, however, holds that the metalanguages never do become precise; if this can be the case, she could avoid being committed to either (1) or (2). On her view, “the vague is not reducible to the non-vague” (208). The meta-languages in which we state the precisifications will themselves be vague, so when faced with a paradox at one order of vagueness, one can always appeal to a meta-language of that language. She writes: “If there is no general objection to the claim that the sequence of metalanguages for metalanguages is potentially infinite, then what is the difficulty with adding ‘and each of those languages is vague’?” (208)

Keefe is correct that the sequences of metalanguages for metalanguages is potentially infinite; but unfortunately for the supervaluationist account, the Sorites series are *in practice* discrete and finite, and this is enough to cause serious trouble. When one adds an ‘and each of those languages are vague’, it has the effect of claiming that, within a discrete Sorites series, that there is a range of valuations (at the level below) that are admissible. So, if the first order borderline is such that the border-area is the range [401, 600], then adding ‘and that range is vague’ entails that there will be a second order border-area surrounding 401, considering just the lower-bound (as I will do in all that follows – parallel arguments can be made for the upper bound). Let’s say that range is [376, 425]. Then, if we add another vague metalanguage, then that entails that there is some range surrounding 376, say, [361,390]. For each additional metalanguage, either the metalanguage extends the lowest bound further towards 1 or it remains the same. (Fine’s stability condition requires that each higher-order valuation admit the lower-order

valuation; Fine 1975, p. 126.) Either this will result in an infinitely iterated end point (as in my (1) above), or it will lead eventually to the first member of the Sorites series itself (as in my (2) above). Appealing to the potential infinitude of metalanguages does not help when Sorites series are finite. Hence there is no alternative to (1) or (2).¹⁵

Keefe seems to sense this concern in her discussion of Horgan's Forced March argument. She defends the view that metalanguages never become precise by claiming "the nature and depth of vagueness ensure that not all indeterminacy can be removed" (211). As a consequence of this, there will be "questions without answers". She proposes this as a means to resolve Horgan's forced-march paradox. But this "what, me answer?" response does not succeed. When will the first unanswerable question be asked? Is B(0) bald? That seems answerable in that it is true (Keefe seems to hold this). What about B(1)? Etc. It seems that at some point, the Keefian supervaluationist being marched through will refuse to answer; but when? Why begin refusing to answer at that point?

Michael Tye, who defends a three-valued solution, offers a response to the forced march that is extremely similar to Keefe's: (1994, p. 205 fn 27) "Horgan assumes that it is legitimate for him to force the defender of the [3-valued] logic to face a question of the form 'Are B(n) and B(n+1) alike in truth value?' for *every* pair of adjacent statements in the Sorites sequence, one after the other. But this can be legitimate only if *each* such

¹⁵ Sanford argues (personal correspondence) that the example I have chosen here is unrealistic. He writes that higher-order borderline cases cover a smaller range than the prior lower-order ones, and that there are likely very few cases of second order borderline case (much fewer than 50 cases), although we are unsure where the second-order borderline cases are. I think Sanford is right that for each individual who is brought through a forced-march Sorites, the number of borderline cases should decrease. But for us, *sub specie philosophiae*, the higher-order ranges represent the region in which users may admissibly place the relevant lower-order boundary. If it is true that we do not know where the second order borderline region is for any given individual, that means that there is a somewhat wide range which may be considered third-order borderline cases. This is why I have chosen a region as wide as I have, though it should be noted that no individual user should choose a wide third-order borderline region.

question has an answer. And, in my view, it is simply not true that each question has an answer.” What does Tye mean by it being ‘illegitimate’ for Horgan to force a 3-value logician to answer a question? Horgan could just sit the person down and start asking questions; if the logician leaves the table, that is her own prerogative, but that leaves her account with a important unexplained boundary. As with Keefe, we must ask: at what point will the person leave the table? Why?

Another consideration concerning supervaluation is that on Fine’s account, which Keefe follows to some extent, all valuations are classical, but they could instead be many-valued. Keefe mentions (2000, p. 200) Sanford’s (1993) suggestion along these lines. This kind of view may have the benefit of not forcing the degree-theorist to provide one specific truth value for each claim. However, the objection, using the forced-march paradox, seems to hold against this version of supervaluation as well: what point is the first case at which it becomes admissible to some degree to not call the candy lavender to a 100% degree? What facts about our usage could possibly indicate that point?

Though Sanford’s account is still awaiting formalization, it seems likely that this kind of supervaluationist view can be placed on a par with one that employs classical logic. The precisifications, on a supervaluationist view, are not intended to, in themselves, represent the meaning of natural language sentences. If that’s the case, then why have any preference for any kind of valuation over another? As with the different kinds of degree theory, I shall not attempt to answer that, but my suspicion is that equal cases can be made for different kinds of supervaluational models. Furthermore, the argument I gave in the first chapter, concerning how to define the adm and def operators, creates more uncertainty about which supervaluationist account (if any) is correct— what

really should count as an admissible valuation? Even if Keefe is correct that it is all right if more than one supervaluationist formalization may successfully describe vague predicates, at least something should be said about why there is such a vast plurality. Although, as above, Keefe is no advocate of formalism, she, like her opponents, still begins with the workings of a formal structure, and then argues that it will fit the nitty-gritty details concerning natural language. My approach, in the coming chapters, will be to begin with the examination of language and semantics (and metaphysics and epistemology and theories about rules and normativity), and then, only at the very end, to add on to it whatever formal structure seems best.

“I love a broad margin to my life.”

-Thoreau, *Walden*

CHAPTER 3. THE EPISTEMIC VIEW

The epistemic view, brought into contemporary discussion by Roy Sorensen and Timothy Williamson, has emerged as one of the most-discussed proposals for a resolution to Sorites paradoxes. On it, vague predicates have sharp boundaries, but we do not know where the sharp boundaries are located. It is an elegant and highly counterintuitive way to resolve Sorites – the major premise, and, with it, the OD→PD and boundarylessness intuitions, are denied without apology.

There are three main motivations for holding the view that its advocates cite: first, the difficulty in establishing any *other* solution to Sorites paradoxes that maintains those intuitions speaks in its favor. Second, the principle of bivalence should not be abandoned, and if it is maintained, it entails that there are sharp, unknown boundaries. Third, an independently motivated margin for error principle (from reliabilist epistemology) may give intuitive grounding to the view. I believe that none of these three actually provides sufficient motivation for the view. Concerning the first intuitive motivation, I am somewhat sympathetic to the epistemicist's view: I agree that there is, at present, no other account that works. However, I believe this is a reason to continue to explore new options, and not settle with a highly counterintuitive view just because it is the only one under which paradoxes are actually resolved. I shall not discuss this point further, except insofar as this entire work is an effort to provide an alternative to all of the accounts on offer. In 3.1, I discuss the kind of commitment that Sorensen and Williamson have to

bivalence, and I argue that the grounds they cite for maintaining the principle of bivalence are not enough to motivate an epistemic view. First, the pragmatic grounds they give in favor of bivalence are not strong enough, and second, maintaining bivalence does not necessarily favor an epistemic view over some other options. In 3.2 I argue that Williamson's margin for error principle is false, but I shall then reconstruct a principle similar to it that is true and does the same work as Williamson's original principle. But I shall then argue that Williamson's claim that the margin for error principle provides independent motivation for the epistemic view obscures the fact that it is the argument that Williamson gives in showing how the principle applies to vague predicates, and not the principle itself, that motivates the epistemic view.

What is left as the basis for the epistemic view is this very argument, which concerns the mysterious relationship between meaning and use. I discuss this at length in 3.3 through 3.6. I shall argue that eliminativism, by which our ordinary language terms do not refer to anything, is preferable to the epistemic view. (This is just a stage in a dialectic: in Chapter Five, I reject eliminativism.) The epistemic view's claim that there is a precise answer to which candies are lavender and which are not is contrary to a common sensical assumption that for observational terms like 'lavender', there are no truths of which we are in principle unaware. It is hard to conceive how it is possible either for there to be a matter of fact that a borderline candy is lavender or for there to be a matter of fact that a borderline candy is not lavender. Of course, saying that it is hard to conceive of such a way is no *argument*, although some authors seem to use this consideration as an argument (Sainsbury, 1990; Horgan, 1994, 1998). I build this

intuitive consideration into an argument by discussing the broader connections the epistemic view has to the realism/anti-realism debate.

3.1 The Epistemic View and Bivalence

A good way to open the discussion is to cite Williamson and Sorensen:

Is epistemicism about vagueness therefore like David Lewis's modal realism and Graham Priest's dialetheism – hard to refute, hard to believe, the victim of the incredulous stare? One crucial difference is that modal realism and dialetheism, unlike epistemicism, are revisionary in logic... Epistemicism employs a different methodology: one holds one's logic fixed, to discipline one's philosophical thinking. It is its opponents who reject the discipline. The epistemicist's hunch is that in the long run the results of the discipline will be more satisfying from a philosophical as well as from a logical point of view. (Williamson 1997, pp. 217f)

Classical logic and semantics are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains. In these circumstances it would be sensible to adopt the epistemic view in order to retain classical logic and semantics even if it were subject to philosophical criticisms in which we could locate no fallacy. (Williamson 1994, p. 162; also cited by Wright, 1995)

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. (Williamson 1997, p. 215)

We should not retreat from standard logic to rescue speculative hypotheses about how language operates. Change in the web of belief should be made at the most peripheral portion available. Beliefs about how language works are far more peripheral than beliefs about logic. After all, anti-boundary beliefs emanate from philosophy of language, not linguistics or some other scientific discipline. Instead of changing logic, we should change our opinions about how language works. (Sorensen 2001, p. 8)

As one can see, a unifying theme for both Sorensen and Williamson is to keep bivalence at all costs, and then to build a theory of Sorites around it. I believe that this kind of *logical snobitivist* argument in favor of bivalence is spurious. Are humans really better at logic than at philosophy? In introductory logic classes everywhere, students are told that

one must take sentences in natural language, translate them into a symbolic language, and then, using classical logic, construct proofs, etc. Carl Posy, in instructing his teaching assistant (me) about how to teach logic, would often remark about how doing translation always involves some amount of *lying*. Classical logic works only once we have translated sentences from natural language into a formal language, but this is no easy task. I shall continue this theme shortly.

I will not respond to Sorensen's potshot¹ placing philosophy of language below linguistics and other sciences here. I would rather ask that my arguments in the philosophy of language in what follows simply be given a fair hearing. But I will remark that it does not speak well of logic, linguistics, and other sciences if a defender insists upon excluding considerations from other fields out of hand in order to maintain their own balance. And as we shall see, the epistemic view's commitment to bivalence itself relies upon controversial assumptions within philosophy of language and metaphysics.

As an adherent to bivalence myself, their remarks seem off-target for a very specific reason: There is not one relevant principle of bivalence, but *at least two*:

(BP) Every proposition is either true or false.

(BU) Every assertoric utterance is either true or false.

(Assertoric utterance can here mean: a speech act in which the utterer says that the world is a certain way.) Those who maintain that we should keep bivalence must distinguish

¹ In Sorensen's defense, his whole book has quite a humorous tone to it, so as a philosopher of language, I do not take as much offense from his remark as perhaps I would if it were uttered by another individual in a different context.

between these two principles, and Sorensen tends to obscure the difference.² But what is unclear is whether, in our commitment to classical logic and bivalence, and our placing it at the center of our web of belief, (BP) is central, or (BU) is central. If (BP) is what we must unconditionally accept because it too deeply embedded in our belief structure, it does not entail that (BU) must also be accepted unconditionally. But it seems that at issue is not whether to accept (BP), but whether to accept (BU).

What is not often discussed by the epistemicists is what exactly the costs are of abandoning one of the principles of bivalence. Logic may provide the foundation for mathematics, but it seems that only (BP) is relevant there. (Of course, considerations from intuitionist mathematics call (BP) into question.) In Chapter Two, I speculated that the costs of using supervaluationist logic and degree-theory logic might be minimal. For the “snobitivist” argument to work, further examination of the costs is required. Neither Williamson nor Sorensen has given a close analysis of what exactly the costs will be.

² Sorensen discusses (2001, p. 8) Derek Parfit’s view by restating an example from Parfit’s:

Suppose a club is ‘revived’ after years of inactivity. Have the people reconvened the same club or have they merely started another club? Unless the original club had rules specifying how it can only be reconvened, there is no determinate answer to this question. ‘Though there is no answer to our question, there may be nothing that we do not know’ (citing Parfit 1984: 213).

This Rationale conflicts with the law of bivalence: every proposition is either true or false.

This is a non sequitur. To say that there is no answer to the question ‘Have the people reconvened the same club or have they merely started another club?’ is not to say that there is a proposition that is neither true nor false. Rather, it is merely to say that there is a sentence that is neither true nor false. An account of its lack of truth may appeal to the thought that ‘club’ has thereby been shown to be semantically defective, and thus sentences using the term do not express propositions. The preceding line in Parfit states: “The claim ‘This is the same club’ would be *neither true nor false*.” But Parfit does not argue that that there is a proposition that is neither true nor false.

The distinction between (BP) and (BU) is important because it seems that any of the non-classical solutions to Sorites are compatible with accepting (BP). For example, supervaluationist views may reject (BU), but accept (BP): we can construe sentences as not expressing single propositions, but rather that they partially refer to each member of a set of propositions, which themselves are classically bivalent (see McGee and McLaughlin, 1995, and McGee, 1997). Propositions can then be held to be bivalent but sentences are not. It should be noted that Fine himself calls the principle of bivalence “the principle that every *sentence* be either true or false” (1975, p. 121, italics added), but we do not need to agree with Fine’s assessment in order to accept supervaluationism. Seen this way, supervaluationist logic supplements classical logic, rather than replaces it. (Sorensen (2001, p. 8) commends modal logic because it supplements classical logic, but scorns how alternative logics for vagueness only subtract inference rules.) Degree theory, as well, could be seen as an account of a continuous, non-bivalent level of a certain kind of propositional attitude individuals take towards propositions which are themselves bivalent. Sainsbury (1986) explores this approach, using a notion of degree of belief (unsuccessfully, I believe). There may be conclusive reasons to maintain (BP) at all costs, but this does not mean that we must accept (BU) regardless of the costs. But whether to accept (BU) is precisely what is at issue.

Williamson for his part does a very good job distinguishing (BP) from (BU) (1994, §§ 7.1-7.2), and his acceptance of both does come from arguments, and not just from an unconditional adherence to them. I mean to suggest merely that the argument from the indispensability of bivalence is perhaps an argument in favor of (BP), but it is

not an argument in favor of (BU), and it is the analysis of (BU) that really matters in questions of vagueness.

Terry Horgan rejects classical logic without remorse. Horgan's own view admits that there is contradiction in our practices. Though I disagree with that view (I shall discuss it briefly in Chapter Five), he gives an amusing and provocative expression of a sentiment I for the most part share:

It is time to get over our Victorian hangups about the allegedly inevitable evils of logical incoherence, and to acknowledge the real attractions of sado-semantic dominance relations among mutually unsatisfiable semantic standards. Arguably, this sort of kinkiness is what's been going on all along in the case of vagueness, right under our noses; we should not be shocked. (1998, p. 317)

Eventually, I will agree with Williamson on the question of bivalence, though not without putting it into serious doubt and seriously exploring the kinds deviance that Horgan recommends. In Chapter One, I do make a 'quick argument' in favor of bivalence. On this argument, propositions are defined to be bivalent; hence (BP) is an analytic truth. But the conditions that I set for a proposition are quite high: a proposition says that the world is a certain way. On the assumption that sentences express propositions, (BU) is true, but it may be difficult to support the claim that human sentences express propositions. I return to this in Chapter Five.

What are the consequences if we do unconditionally adhere to (BU)? What then should be said about Sorites? Both Williamson and Sorensen use (BU) to argue that there thus must be *unknown* boundaries. The argument, not ever stated explicitly, is something like:

- (1) Statements such as 'that is lavender' are always either true or false.

- (2) In certain cases, we do not know whether the statement ‘that is lavender’ is true or false. We do not know which candy is such that it is lavender and the one next to it is not, but, given bivalence, there must be one such candy.
- (3) Therefore, there are unknown boundaries.

But here is another way to go, after accepting (BU). I shall call this the Q argument:

- (1) Statements such as ‘that is lavender’, if they are assertoric sentences, are always either true or false.
- (2) In certain cases, we do not know whether the statement ‘that is lavender’ is true or false.
- (3) Hence their status as assertoric sentence is threatened – perhaps they do not say that the world is a *certain* way, even though we do indeed intend such statements to say *something* about the world.
- (4) Thus, on pain of inconsistency, we must *stipulate* that they say that the world is a certain way if our intention to say something about the world is to be realized. (BU) is only applicable if we stipulate a precise boundary for vague terms.
- (5) But these boundaries, given that they are stipulated, are not unknown.

Epistemicists use (BU) *descriptively*, not *normatively*. They claim that our vague sentences are, and have always been, precise. But on my normative construal of vague sentences, all (BU) indicates is that sentences *must be* either true or false if they are to be assertoric, and vagueness demonstrates that our sentences are (probably) not assertoric unless we take action and make them so. This is a result far different from what Williamson and Sorensen use of (BU).

It might be useful to compare Williamson and Sorensen to another philosopher who accepts bivalence in spite of vagueness – Quine, in “What Price Bivalence?” (1981). It is Quine from whom I derive the Q argument. In that short paper, Quine lays out what I consider to be the most common-sensical, philosophically grounded, and internally coherent approach to the Sorites paradox. It is unfortunate that neither Quine nor anyone

else, to my knowledge, has developed his approach into a full theory of vagueness. Quine adheres to bivalence because of “the simplicity of theory it affords” (1981, p. 91), but he also acknowledges the costs of accepting it – awkwardness in Sorites cases. Following Wright (1975), Quine notes that observational terms will be imprecise, because of the limits of our perceptual acuity. This imprecision leads to Sorites paradoxes. What we must do, according to Quine, is

arbitrarily stipulate, perhaps, how few grains a heap can contain and how compactly they must be placed. What had been observation terms are arbitrarily reconstrued, on pain of paradox, as theoretical terms whose application may depend in marginal cases on protracted tests and indirect inferences. The Sorites paradox is one imperative reason for precision in science, along with more familiar reasons. (1981, p. 92)

This is indeed very close to the view I myself advocate. One point I shall dispute in Chapter Five is the claim that the reconstruction of the imprecise term into a precise one is ‘arbitrary’.

Quine discusses other questions that we do not know the answer to: whether the number of blades of grass on Harvard Yard at dawn of commencement day, 1903, is odd, or whether there is a hydrogen atom located within a certain distance of a certain point.

Quine writes:

one feels differently about the question of the heap, or of baldness: that it is a mere question of words, to be settled by a stipulation... In what way, then, are the questions of heaps and baldness matters of convention, and the other matters of fact? One way to bring out the contrast is in terms of our physical theory itself, in full acceptance of bivalence. Namely, the number of the blades of grass and the presence of a hydrogen atom are physically determined by the spatio-temporal distribution of micro-physical states, unknown though it be. Where to draw the line between heaps and non-heaps, on the other hand, or between the bald and the

thatched, is not determined by the distribution of microphysical states, known or unknown; it remains an open option. (1981, pp. 93-94).³

Quine's approach highlights a certain methodology: accept bivalence, analyze our linguistic practices as best we can, and then fit them within the constraints of bivalence. Williamson and Sorensen accept bivalence, and then [re-]describe linguistic practice such that it fits the constraints of bivalence, even though it does seem to. Quine's view accepts that linguistic practice does not fit bivalence, and so we must posit⁴ that there are precise demarcations.

Another philosopher writing on vagueness who maintains bivalence is Unger (1979). Unger does not argue in favor of bivalence inasmuch as he doesn't consider other options. But he uses Sorites paradoxes to draw a much different conclusion from Williamson, Sorensen, and Quine: to show that there are no macro-level physical objects. If there are tables, reasons Unger, then they must be composed of atoms. Removing one atom from a table will always leave a table. But iteration of this Sorites step will lead to a paradoxical result that a table may be composed of one atom. Hence there are no tables

³ Sanford, in comments to me on this passage, makes a very helpful distinction between *relative* precision and *absolute* precision. While the expression 'blade of grass' may be more precise than 'heap', Sorites paradoxes can still be made on it. This consideration might convince Quine to place 'blade of grass' in the same category as 'heap', rather than in the same category as 'hydrogen atom' – it goes to show that 'blade of grass' won't have a place in our final physical theory. However, to me, such 'relatively precise' terms might be important in that while they may or may not have a place in a final physical theory, they may still have a place in our *best current* account of how things are. This will relate to issues I discuss in Chapter Five.

⁴ Quine uses 'posit' here rather than 'stipulate' because it will be impossible for humans to stipulate a precise reference for 'table'. That's also why I use "in principle" in the preceding sentence – we cannot in practice stipulate a definition for table. Quine might be wrong about this point: 'table' is defined functionally, and so we can stipulate a table as something that can hold up the weight of a certain arrangement of items on top of it for a certain amount of time; we need not phrase our stipulation in micro-physical terms, though of course whether an item satisfies our functional definition of a table will be dependent upon its micro-physical structure. This is an important issue, and I shall return to it in Chapter Five.

(see 1979, pp. 236-238). Williamson, in Chapter Six, argues against this view (with only partial success, I believe), but its existence shows that there are indeed alternatives to the epistemic view even if we accept bivalence. In this chapter I shall argue in favor of this kind of eliminativist view.

It strikes me as if epistemicists accept their view because it seems that other solutions are unfeasible. At a crucial point, Williamson defends the epistemic view's account of the connection of meaning and use, as I shall soon discuss, on the grounds that no one else has given a successful account of it either. I find this reason to continue to search for an even better alternative, and not for accepting any current view. My goal here has just been to show that one who accepts bivalence need not accept the epistemic view. This is the converse of what Wright demonstrates – that one who accepts an epistemic view need not accept bivalence.

Williamson does make one argument against alternative logics that seems to me to be misguided. He writes (1994, p. 191):

Formal semantic treatments of vague languages – many-valued logics, supervaluations and the like – are characteristically framed in a meta-language that is conceived as precise. Thus one cannot say in the precise meta-language what utterances in the vague object-language say, for to do so one must speak vaguely; one can only make precise remarks about those vague utterances. Since the expressive limitations of such a meta-language render it incapable of giving the meanings of object-language utterances, it can hardly be regarded as adequate for a genuine semantic treatment of the object-language.

This seems much too quick. If we replace 'vague' with ambiguous and 'precise' with 'unambiguous', in the second sentence in the above paragraph, it reads: "One cannot say in the unambiguous meta-language what utterances in the ambiguous object-language say, for to do so one must speak ambiguously." This is obviously false: ambiguous

utterances are easily characterizable in unambiguous language. But for supervenionists such as Fine, Lewis, and McGee and McLaughlin, vagueness just *is* a kind of ambiguity. One might dispute that vagueness is ambiguity, but if it is, then one *could* say in a precise metalanguage what utterances in the vague object-language say. Does Keefe's version of supervenion, in which the metalanguages never become precise, make it any better in Williamson's eyes than versions of supervenion in which the metalanguages do become precise? It is hard to see how a form of supervenion gains expressive power by not giving precise specifications.

One other final note here is that what will become relevant is whether *minimalism* about truth and meaning is correct, and what theory of vagueness (if any) is entailed by it. Minimalism has a long history and is still a prevalent view, and it seems to underlie much of what is said by Sorensen, Williamson, and Wright. I mention this view briefly below, and shall return to it again in Chapter Five. Williamson takes something of a minimalist stance – he seems to take our sentences at face value. But I shall argue, against minimalism, that because of Sorites, we must venture to do the hard work in the philosophy of language to explore the meaning of our terms. As I discuss below, *minimalism about the meanings of terms in a discourse is appropriate only if there are no problems arising in the discourse*. When there are problems, it demands that we examine the meaning of the terms in the discourse, and be prepared to abandon them.

3.2 Margins for Error and Epistemicism

3.2.1 Williamson's Margin for Error Principle

Williamson states in his preface (1994, p. xi) that a margin for error could provide an independent motivation for the view. He writes in this brief autobiographical section:

For years I took this epistemic view of vagueness to be obviously false, as most philosophers do... However, I continued to think about the epistemic view, for the standard objections to it did not seem quite decisive. It was not clear that they did not assume a suspect connection between what is true and what we can verify. It then struck me that the notion of a margin for error could be used to give a specific explanation of ignorance of the sharp boundaries of our concepts, and the epistemic view began to look more plausible.

Margin for error principle are based upon the idea that in order for one to know something, one must be a reliable indicator of its truth. For example (Williamson, 8.2), if I attempt to estimate the attendance in a large stadium, and I claim that there are at least 20,000 people in the stadium, and it turns out that there are exactly 20,000 people there, though what I uttered was true, I did not know that it was true, because I was not a reliable indicator of it – had there been one fewer person in attendance, I would not have been able to tell the difference, and still would have reported that there were at least 20,000 people in attendance. Williamson formalizes the principle as follows:

A margin for error principle is a principle of the form: 'A' is true in all cases similar to cases in which 'It is known that A' is true. (1994, p. 227)

Williamson gives a nice metaphor for this thought:

Believing is often compared to shooting at a target, the truth. The comparison is not quite apt, for the truth is a single point (the actual case), like a bullet, while the proposition believed covers an area (a set of possible cases), like a target. Instead, the believer's task may be conceived as drawing a boundary on a wall at which a machine is to fire a bullet. The belief is true if the bullet hits the bounded area. If truth is a hit, knowledge is a safe hit. That is, the point of impact is within the bounded area and not

so near its boundary that the bullet could very easily have landed outside (had a light breeze blown). For example, a hit might be safe just in case every point on the wall less than an inch from the point of impact is within the bounded area. The one-inch margin inside the boundary corresponds to the cases in which 'B' is true but unknown; when this margin is removed from the bounded area, the remaining area corresponds to the cases in which 'B' is known. (1994, p. 228)

In the context of Sorites paradoxes, Williamson gives an instance of the principle:

(!) If we know that n grains make a heap, then $n-1$ grains make a heap. (1994, p. 232)

If this margin for error principle is true, then a consequence is that there will be true statements that we do not know to be true, which is what the epistemic view claims.

Williamson then argues that the fact that we do not know the locations of sharp cut-off points make us "fall under the illusion that such points do not exist." (1994, p. 247). So, Williamson uses the margin for error principle to explain how there can be unknown truths and to give an error theory of why his opponents might (erroneously) believe that there cannot be such unknown truths.

Though much of Williamson's writing on vagueness has received a great deal of response in the literature, there has been very little written on his use of the margin for error principle. Williamson does not use it as a deductive part of his argument, but as an independent motivation for the epistemic view. It doesn't persuade me of much. For claims about the number of people in a crowd, or height of a tree, or point of impact of a bullet, some kind of margin for error principle seems quite reasonable; these are matters of prediction and estimation of facts that, in Quine's terminology, are uncontroversially physically determined by the spatio-temporal distribution of micro-physical states. Given that these subvenient facts are beyond our ken at the appropriate moments, there should

be a margin for error when we report our beliefs. But for vague terms such as ‘red’ and ‘heap’, we do not seem to be estimating what we believe to be a precise fact of the matter. The intuitive thought is that they are conventional in some respect. The question is: can Williamson’s invocation of the margin for error principle in itself convince one who agrees with Quine’s assessment that questions of whether or not something is a heap is a convention, and not determined by micro-physical states? I cannot see how it could (and I do not know of anyone, aside from perhaps Williamson himself, who has been persuaded by it).

Another way to express the difference between terms like ‘red’ and terms like ‘the number of blades of grass on Harvard Yard at dawn on commencement day, 1903’ is using Wright’s principle of evidential constraint, (EC), which posits that there is a necessary connection between the feasibility of knowing the truth of a statement and its truth, for observational terms like red. Consideration of Williamson’s margin for error principle does little to rebut (EC), if one maintains a commitment to (EC) more strongly than to the margin for error principle applied to vagueness. I’m surprised no one has brought this point up against Williamson – perhaps the reason is that it is difficult to respond to an argument that says that such and such a claim should persuade you of this other claim – except by saying ‘it doesn’t persuade me’. Below, I shall argue that the argument that leads up to Williamson’s application of the margin for error principle to vague predicates is what might be persuasive, and not the margin for error principle itself.

3.2.2 The Error in Williamson's Margin for Error Principle

Putting aside these concerns for the moment, it can be demonstrated that the margin for error principle is false, and it indeed fails as a general principle in epistemology, and not just when applied to Sorites paradoxes. Peter Mott (1998) has criticized Williamson using the same sort of objection that I will be discussing, but Mott handles the objection in such a way that it leads Williamson, in a response (2000), to interpret Mott in such a way that Mott's argument is easily refutable. I believe that this is partially Mott's fault and partially Williamson's. To avoid getting caught up in that confusion, I shall develop my argument against Williamson independently of Mott's, and I shall note connections to Mott's argument.⁵

One inadequacy in Mott's article is that it fails to provide an initial intuitive motivation for why we should suspect that the margin for error principle must be incorrect – and Williamson at a crucial point (2000, p. 80) argues that Mott has provided no reason to believe that the margin for error principle fails, apart from Mott's formalism. But there is an easy way to capture what is wrong with it. Williamson's point about the target is that in order for one to *know* that the bullet will land inside the bounded area, it must be the case that (to repeat)

the point of impact is within the bounded area and not so near its boundary that the bullet could very easily have landed outside (had a light breeze blown). For example, a hit might be safe just in case every point on the

⁵ I should note that I read Mott's article and Williamson's response to it after formulating my own views on the margin for error principle. In the end, I do believe that the second of the arguments that Mott gives, when interpreted correctly, is a proper objection to the principle, but Mott phrases it in such a way that it invites Williamson's misinterpretation. I should also note that Mott heralds his defeat of Williamson's margin for error principle as a blow against the foundations of the epistemic theory (1998, p. 503) and Williamson agrees (2000, p. 77) that if Mott's argument were successful, it would be such.

wall less than an inch from the point of impact is within the bounded area.
(1994, p. 228)

The motivation for the margin for error principle is that one cannot know something that could very easily have been false. The fallacy in the above argument is that Williamson has implicitly assumed that no breeze has *in fact* blown. From the premise that the point of impact is just inside the boundary, Williamson concludes that it very easily could have landed outside. But this inference is invalid.

To show this, let me introduce the notion of a *conspiring variable*. A conspiring variable is an attribute of an epistemic situation that jeopardizes the correctness of a judgment. Conspiring variables are those things external to the judge that, depending on the values that they take, might give the judge bad epistemic luck. In every epistemic situation, there will always be some conspiring variable, because we are not omniscient, though this is not to say that the conspiring variables will always have values that are significant in any way. In order for us to know something, the conspiring variables could not easily have values that would make the judgment false. We should not have a criterion of knowledge that one knows only if it is not *possible* that the conspiring variables will have values that will make the judgment turn out false, because that would set the bar for knowledge too high, and would too easily lead to skepticism. Rather, the (reliabilist) criterion for knowledge should be that one knows only if it is *unlikely* that the conspiring variables will have values that will make the judgment turn out to be false.

In the example with the target, which I shall hereby turn from a metaphor into a real example, let's assume that there is one conspiring variable: the strength of the breeze. Let's say that the circular target is three inches in diameter, and the judge places it (or

draws it) on a wall such that the judge believes that a bullet fired from the machine will hit the target. Let's assume that it is not rare for there to be a breeze that will blow the bullet up to an inch in one direction, but only very rarely does a strong breeze blow the bullet more than an inch. I shall say that the *normal range* of the conspiring variable is between 0 and 1 inches (denoting the value of the breeze by the amount of inches it blows the bullet). Such an example is intended to be in line with Williamson's supposition that there will be a one-inch margin for error – given a one-inch margin within the target, it is unlikely, though not impossible, for the conspiring variable to have a value that will make the judgment come out false.

In a case where there was no breeze, and the bullet lands less than an inch from where the judge has placed the boundary of the target, Williamson is correct that the person's judgment could quite easily have been wrong – a breeze could easily have blown it an inch off its course so that it could have landed outside the boundary. Hence the person is not a reliable judge and does not know that it will land inside the boundary. But let's consider another case where there *is* a moderate breeze that blows – say a breeze of .8 inches. Let's assume for this example that the person is a reasonably reliable judge, such that had there been no breeze the bullet would have landed near the center of the target, 1.4 inches from any boundary. But because of the breeze, the bullet lands less than an inch from the boundary – .6 from the edge. According to Williamson's claim, the person does not know that it will land inside the boundary, given that it lands within the one-inch margin. But the fact that it landed within the one-inch margin does not show that the judgment was unreliable – the judgment is unreliable only if the bullet lands

inside the one-inch margin *and the conspiring variable actually had a minimal value*⁶. If the reason that it landed within the one-inch margin is that the conspiring variable had a value in its normal range, which caused it to land in the margin rather than in the safe area near the center, then the fact that it landed within the one-inch margin does not show that the person is likely to have been wrong – it would have been very unlikely that there would have been a breeze that would have blown the bullet more than 1.4 inches. Despite the bullet's landing .6 inches from the boundary, the person still knew that it would hit the target. Williamson's margin for error principle is false.

What matters for knowledge is whether the person's belief is *in fact* reliable. What the one-inch margin really can be used for is to test potential judges in, say, indoor situations where there is no breeze. If you need a good target placer, and you know that breezes usually blow up to an inch, then you can test a person's reliability by having a controlled situation such that you only accept target placers who, when the conspiring variable is set to zero, have a margin for error that is within the normal range of the conspiring variable that you are controlling for, i.e., one inch. If there is no controlled testing environment, it is difficult to determine whether a person is a reliable judge if one does not have a wind-gauge. If one sees that in an outdoor placement, the bullet comes within an inch of the boundary, one might not know whether or not the person is reliable.⁷ To know it, one must also know whether the reason it came within the one-inch

⁶ Specifically, a value of less than .4, which is calculated from the equation $1.0 - (\text{distance from boundary})$.

⁷ Of course, if one gives the placer multiple tests, then one can get a good sense of whether the placer is reliable. My argument has assumed that reliability can be revealed in one trial; I do not see this as

margin is because there was no breeze and the person got lucky that there wasn't, or because the person predicted quite well, and it only landed within the one-inch zone because the conspiring variable acted up.

There are two possible objections that I would like to discuss here. First, an objection to my argument would be to claim that, given that *was* a light breeze that blew, there could easily have been a slightly stronger breeze that would have blown it further off course. But this objection relies on analyzing the counterfactual in a metaphysical, rather than epistemic manner. Given the starting assumption that it was/is unlikely for a breeze to carry it more than one inch, it was/is unlikely for a breeze to carry it more than one inch even though the actual breeze was, say .8 inches. (The preceding sentence is analytically true.) Sanford comments (personal correspondence): “Actually, the sentence might be false. Consider the following possibility: most hours the wind speed never exceeds .5; most hours when the wind speed exceeds .75 it exceeds 3.0.” Sanford’s scenario is one in which the nearest possible worlds to one in which the breeze was .8 are ones in which the breeze is more than 3.0. And in the example, the actual world’s breeze is .8, so it seems as if it is likely for the breeze to exceed 1.0. So let me clarify: even if the wind speed was 3.5, in Sanford’s scenario, it still was unlikely that the wind speed exceeded 1.0 – it’s just a case where an unlikely thing happened. If we analyze the counterfactual as Sanford seems to be suggesting, we lose the ability to say, in any situation, that something unlikely has occurred – of course the nearest possible worlds to the actual world will be similar to the actual world! That is not the result we want. The

a flaw of the account, especially if the target analogy is read somewhat metaphorically, so as to cover one’s general tendencies to be “on target”.

actual world is one in which it is unlikely for there to be a breeze of more than 1.0, from the perspective of a target-placer prior to placing the target. (Which is the relevant perspective to explore for this example.) So even if, metaphysically, the nearest possible worlds to the actual world are ones in which the breeze is greater than 3.0, epistemically, it is unlikely for that to happen, and it is this epistemic possibility that is at issue here. So I stand by the claim that “Given the starting assumption that it was/is unlikely for a breeze to carry it more than one inch, it was/is unlikely for a breeze to carry it more than one inch even though the actual breeze was, say .8 inches” is analytically true.

The second point is that the judge herself need not know exactly how reliable she is. *All that matters for knowledge is whether a person’s true belief is in fact reliable.* This is where the confusion in the Williamson/Mott dialectic occurs. Williamson (2000, p. 79) describes Mott as claiming that my knowledge of the following claim:

(3) The tree is between 50 and 60 feet high.

comes from a deduction from two other claims, *both of which I know to be true*:

(4) I estimated that the tree is 55 feet high

(5) If I estimated that the tree is 55 feet high then it is between 50 and 60 feet high.

Williamson rightly criticizes the use of (5) within the larger scope of Mott’s argument (not restated here) on the grounds that one’s knowledge of (5) must have a margin for error – one can’t *know* that one’s margin for error in estimates is exactly 5 ft.⁸

⁸ Mott says (1998, p. 499): “I *estimate* that the tree is 25 feet high and then *judge* that it is between 20 and 30 feet high. We do not claim knowledge for the estimate (even if it is correct), but only for the interval judgement.” The statement here that we claim knowledge of our accuracy interval may have encouraged Williamson’s interpretation. But elsewhere (*ibid.*, p. 502), Mott does state explicitly that one does not need to know the value of one’s own margin for error. As such, I endorse his argument. Williamson’s interpretation of Mott may also be skewed by Mott’s flawed first argument (*ibid.*, pp. 498-

There is no reason to suppose that I myself must know (5) in order for me to be a reliable indicator. It simply must be true of me. Of course, it would be odd for me to have become a reliable estimator without believing that I am somewhat reliable, or without having some understanding of the limits of my reliability. But even if I do not know that I am reliable, I still may be reliable. (I shall argue for this point more strongly momentarily.) The more common scenario is where one becomes a reliable judge about targets because one has *some* beliefs about how the wind might carry a bullet. (One who is careful to never place a target on a wall when she believes that the bullet is unlikely to hit the target might refuse to place the 3-inch target on the wall if she believes that the wind, which she cannot detect in advance, normally blows bullets between 0 inches and a foot.) But there is no reason to suppose that one must need to know the *precise* (and, by presumption, *correct*) reliability principle that (in the case being discussed) there must be a one-inch boundary in order for one to know that the bullet will land in the target. One must simply *be* reliable.

It is worth noting that reliability principles seem to be vague: what facts about the matter make it such that the determination for what is to be considered *highly unusual* is set at 1.0 (as I have stipulated) rather than 1.0001? The vagueness of reliability principles is not my concern here, though I do in fact believe it makes for a serious problem for the concept of knowledge. The above example uses the notion of ‘estimating’ rather than ‘believing’, and I should say something about this. If one estimates that it will land in the exact center of the target, and one is reliable, then it will land within an inch of the center.

499) against Williamson, which Williamson (2000, pp. 77-79) does interpret correctly and rebut successfully.

But this is not to say that one believes in this situation that it will land on the target. As above, one might withhold belief because one does not know how reliable one is.

Presuming, as I have been, that what is required for knowledge is that one have a reliable *belief*, and not just an *estimation* that is likely to be correct, in order for one to have knowledge, one must come to believe that the bullet will land on the target. Hence, in order for one to know that the bullet will land on the target, one must estimate that the bullet will land in the center, and believe that if one estimates that the bullet will land in the center of the target, it will land within 1.5 inches (the radius of the target). Hence, one must *believe* something like (5) in order to know that the bullet will hit the target, it is wrong to assume that the person must *know* (5), or even know a principle with a value slightly different from (5).

As I have emphasized, what really matters is whether the person is reliable. In fact, the constraint that for an individual to know something, the person must know that she is reliable will place an impossible constraint on knowledge. On this constraint, to know anything – such as (5) – I myself would have to know that I am reliable in having that belief – in this instance, a belief in (5). To know that I am reliable in believing (5) would then require me to know that I am reliable in believing that I am reliable in believing that (5). But this will lead to a regress, whereby I can never have enough knowledge to begin to know anything. Provided that we reject skepticism, it must then not be necessary for one to know one's own reliability principle. Rather, one must simply be reliable.

3.2.3 A Reconstructed Margin for Error Principle Applied to Vagueness

Although Williamson's margin for error principle is false, it seems like something like it must be true. After all, it is only in unusual circumstances in which I have knowledge without leaving any room for error. From a perspective of someone who is auditing my reliability as a placer, if all the auditor knows about a situation are that

- (1) The normal range of the conspiring variables is from 0 to 1 inches
- (2) I believe that the bullet will hit the 3-inch target
- (3) The bullet lands 0.1 inch from the edge

then the auditor will suppose that it is unlikely that I knew that the bullet would hit the target. In order for an auditor to determine whether I knew that it would hit the target (if the auditor has no way of knowing which way the wind did blow) she would have to figure out whether it is more likely that I am an unreliable judge or that the wind blew at a level of .9 or greater. The conclusion from this is that, given (1), (2), and (3), it is quite possible, even likely perhaps, that I do not know that the bullet will hit the target.

To review, Williamson's margin for error principle is

- (ME) 'A' is true in all cases similar to cases in which 'It is known that A' is true. (1994, p. 227)

But if my objection is correct, (ME) should be abandoned, but it can be replaced by a principle that states:

- (ME*) 'A' is highly likely to be true in all cases similar to cases in which 'It is known that A' is true.

More specifically:

- (ME**) 'A' is true in all cases similar to cases in which 'It is known that A' is true, *except* for those cases where the knower has a reliable belief that 'A' is true (i.e., a belief that will be true as long as the conspiring variables take a value within their normal range), and the conspiring

variables in fact take a value that is very close to, or greater than, the maximum of their normal range.⁹

Let's examine how Williamson supposes that margin for error principles relate to Sorites cases. I am unsure why Williamson defends the stronger, and false, (ME) rather than a weaker, and true, (ME**), which he may have retreated to in response to Mott's argument. I shall show that (ME**) has the same persuasive power, for his purposes, as (ME). To connect the margin for error principle to vagueness, Williamson argues that meaning of vague terms must supervene on use. By this:

A slight shift along one axis of measurement in all our dispositions to use 'thin' would slightly shift the meaning and extension of 'thin'. On the epistemic view, the boundary of thin is sharp but unstable.¹⁰

Suppose I am on the 'thin' side of the boundary, but only just. If our use of 'thin' had been very slightly different, as it easily could have been, then I should have been on the 'not thin' side. The sentence 'TW is thin' is true, but could very easily have been false without any change in my physical measurements or those of the relevant comparison class. Moreover, someone who utters the sentence assertively could very easily have done so falsely, for the decision to utter it was not sensitive to all the slight shifts in the use of 'thin' that would make the utterance false. (1994, p. 231)

(For reasons of clarity, in what follows I shall replace Williamson's 'thin' example with my 'lavender' example.) Williamson concludes that because the meaning of 'lavender' could vary without one knowing it, if candy 400 is just on the 'lavender' side of the

⁹ There are rare cases where a knower is extremely reliable – more reliable than one needs to be to have knowledge (her estimates are accurate well beyond the normal range of the conspiring variables) – but the conspiring variables take a value outside their normal range, but the knower still gets it right. These cases should also be considered knowledge, even though they do not fit the letter of this definition. To refine the definition, we need an accuracy value for each knower (AV_k), which is the knower's accuracy (measured in distance from the edge) when the conspiring variables are set to 0. K knows that 'A' is true if AV_k is above the maximum of the normal range of conspiring variables (hence K is reliable), and AV_k is greater than the value of the conspiring variable in the actual case, regardless of whether the conspiring variable is in the normal range or not (hence K's judgment is true).

¹⁰ This seems highly reminiscent of Graff's view, though Graff does not make the connection between her view and Williamson's.

boundary then one *cannot* know that candy 400 is lavender, because if one believes that candy 400 is lavender, one could very easily have been mistaken about the truth conditions for 'lavender'. Hence there is a range of possible cases where 'Candy n is lavender' is true, but we do not know that it is true. This discussion has assumed that there is a sharp boundary, which Williamson's opponents will reject. However, it seems that Williamson's point here is that if meaning supervenes on use, then if candy 400 is the last candy for which 'Candy n is lavender' is fully true, then we can never know that it is true. This helps him establish, first, a break between what we know to be true, and what is true, regardless of one's theory of truth. Second, Williamson will try to appease his opponents' intuitions that there can't be a sharp boundary (i.e., the boundarylessness intuition) by claiming that what we really intuit is that there can't be a *known* sharp boundary, and so this argument, if it is successful, establishes that conclusion.

Under my modified (ME**), the main portion of my summary of Williamson's argument should be re-written:

Because the meaning of 'lavender' could vary without one knowing it, if candy 400 is just on the 'lavender' side of the boundary, then one *probably does not know* that candy 400 is lavender, because even if one believes that candy 400 is lavender, one could very easily be mistaken about the truth conditions for 'lavender' *on the assumption that it is not the case that both (a) one is in fact reliable, and (b) the conspiring variables took a value at the maximum of their normal range.*

Let's consider adding the following controversial lines:

Given what we know, it is unlikely that (a) and (b) are met. Hence there is most likely a range of candies where 'Candy n is lavender' is true, but we do not know that it is true.

If one believes that meaning supervenes on use, and truth conditions are variable in the way that Williamson describes, then these two lines are true. Williamson's own epistemic view claims that there *must be* a range of candies in which statements like 'Candy n is lavender' are true but we do not know it, and then a precise boundary that we are unaware of, and then for some range on the other side of the boundary, 'Candy n is lavender' is false but we do not know it. If that is the epistemic view, then the rewritten argument above does not prove it.

However, the rewritten argument does demonstrate a slightly different epistemic view: that there is a probably a range of candies such that both 'candy n is lavender' is true and that we don't know that 'candy n is lavender' is true. In the cases where it is such, the epistemic view as Williamson has it applies. The one scenario in which there are no candies n in which both candy n is lavender and we do not know that candy n is lavender is if we are reliable in saying that 'candy n is lavender' and the conspiring variables are at maximal values. Importantly, in this scenario, there *must be* candies where 'Candy n is not-lavender' is true and we do not know it. In order for one to know that 'Candy n is not-lavender' is true, one cannot easily be wrong about it. The only way for there to be no cases in which 'Candy n is not-lavender' is true and we do not know it is if we are reliable and the conspiring variables for that claim are at maximal values. But this cannot happen in the same scenario as the one in which there are no candies n such that both candy n is lavender and we don't know that candy n is lavender: the conspiring variables for 'candy n is not-lavender' are of course the same types of variables for 'candy n is lavender', except for one crucial feature: the value of the variable for one must be the negative value for the other. Translated to the target example: in those cases

in which you are reliable, and know that the bullet will hit the target and it barely does (because of a breeze), you could not have known, had the gun been aimed very slightly further in the same direction as the breeze, that the bullet would have *missed* the target, because your judgment that it would have missed came true only because the breeze was favorable to you.

The conclusion of this is that there must be some range of candies in the border area in which we do not know the truth value of the statement: ‘Candy *n* is lavender’. We can never know where the boundary is between lavender and not-lavender, even if we might, in rare cases, know, for the last lavender candy, that it is lavender, and in rare cases know that, for the first non-lavender candy, that it is not lavender. There are no cases in which both of these scenarios come true. This is all that Williamson needs out of a margin for error principle, for it gives him the result he wants.

3.2.4 Margin for Error Principles Marginalized

The next question is: how persuasive is all this against one who does not believe in Williamson’s epistemic view in the first place? The margin for error principle was supposed to provide motivation for the epistemic view. I will take it for the moment that Williamson’s opponent is one who rejects that there can be unknown, and unknowable, truths about terms like ‘lavender’. Simplified, Williamson’s opponent is one believes that Wright’s (EC) is a conceptual truth – for observational statements like ‘this is lavender’, if the statement is true, then it must be feasible to know that it is true. Williamson’s argument has a conclusion that denies (EC), so if the opponent comes to believe that Williamson’s argument is sound, then the opponent should give up (EC).

The point that I'd like to bring out is that the (ME) principle (or whatever version of it works, such as the (ME**) principle) is *irrelevant* to Williamson's argumentative structure. The claim that the margin for error principle can be applied to questions of whether we can know the truth of vague sentences is a *consequence* of Williamson's argument that meaning supervenes on use in a highly complex manner.¹¹ That very argument, which concludes that if a term had a slightly different use (or if we had slightly different dispositions to use it) then it would have a slightly different referential structure, is enough to show that we can never know the precise reference of our terms – we can never discover all the facts about our use and dispositions to use terms that are relevant to meaning. But when one advocates (EC), one does not advocate it on the grounds that a statement can be true only if it is feasible to know that it is true by a complete examination of our practices. Rather, the feasibility of knowing that a statement is true, as implied by (EC), comes from the fact that we, as observers, have a special connection to the truth conditions for statements like “that's red” in a way that we do not have a special connection to the truth of judgments about where a bullet will land. The analogy with reliabilism in epistemology is broken unless Williamson can show that we do not have such a special connection. (Frankly, I am unsure who bears the burden of proof. It seems both Williamson and the proponent of (EC) need to prove their claims about how statements like ‘this is lavender’ can be known to be true.)

Though he does not state them explicitly, there are three relevant parts to Williamson's argument using the margin for error principle:

¹¹ This argument happens quickly in his 1994, and perhaps realizing this, he expands on it in his 1996.

- (1) Meaning (i.e., truth conditions) supervenes on use.
- (2) Meaning changes slightly if use changes slightly, in some complex, corresponding, manner.
- (3) Hence meaning could be slightly different, and if so, we could easily be mistaken in our judgments.
- (4) Hence we cannot know the precise meaning of our terms

Williamson gives the supervenience thesis (1994, p. 203):

- (*) If x has exactly the same physical measurements in a possible situation s as y has in a possible situation t, then x is thin in s if and only if y is thin in t.¹²

(1) seems right, and I shall not question it here, though I will ultimately reject it below. If (2) is correct, then (EC) must be false. (2) seems *plausible*, though notice that (1) can be true even if (2) is false. Williamson does not seem to consider the possibility that meaning supervenes on use, but the same meaning can be *multiply realizable* given different uses. One possible such account is a direct-realist-plus-externalist view which may deny (2), such that meaning is determined by external objects, so slight changes in use don't change reference. Meaning is fixed more by the external world than by use of either a community or a person. This account may be true for natural kind terms, for which there is a division of linguistic labor (as in Putnam 1975, Kripke 1980), but this view seems implausible for vague terms. (Williamson agrees that it is implausible – 1994, p. 269; 1997, p. 332.) I shall discuss this issue in greater detail in 3.5 and in Chapter Five.

¹² As Sanford notes (personal correspondence), what is thin for an ocean-liner mooring cable is not thin for a surgeon's sutures. So we should interpret Williamson's supervenience thesis to be relativized to contrast classes. This relativizing might create further problems for Williamson, but I shall not dispute it.

Another way for meaning to be multiply realizable upon different uses is if (EC) is true: a proponent of (EC) would claim that truth conditions cannot vary in undetectable ways, though perhaps use does vary in undetectable ways. Given this, there may be different possible worlds in which use is slightly different but meaning is not. Williamson needs to give an *argument* for (2), for it is not entailed by (1). But whatever argument he gives in favor of (2) will itself undermine (EC), regardless of the margin for error principle. Hence what is important is whether Williamson or his opponents have a better account of the connection between meaning and use.

3.3 Interlude: The Normative Semantic Problem

Another realist position has a problem similar to Williamson's problem of connecting meaning to use. I have in mind a kind of moral realism, as discussed by Michael Smith in his *The Moral Problem* (1994). The moral problem is a problem that arises for a moral realist when the realist's claim of the objectivity of moral judgments is combined with the practicality of them. I shall argue in a moment that what I call the *semantic problem* for Williamson is a problem that semantic realists must face when the (supposed) objectivity of factual judgments (such as 'this is red') is combined with the practicality of them. For a factual judgment, if it's objectively true, then it's precise. But if it's practical, then it's imprecise. For a moral judgment, if it's objectively true, it's true independent of anyone's desires. But if it's practical, then it must connect to individuals' desires.

Here is Smith's statement of the moral problem:

- M1. Moral judgements of the form ‘It is right that I Φ ’ express a subject’s beliefs about an objective matter of fact, a fact about what it is right for her to do.
- M2. If someone judges that it is right that she Φ s then, *ceteris paribus*, she is motivated to Φ .
- M3. An agent is motivated to act in a certain way just in case she has an appropriate desire and a means-end belief, where belief and desire are, in Hume’s terms, distinct existences.

The apparent inconsistency can be brought out as follows: from (1), the state expressed by a moral judgement is a belief, which, from (2), is necessarily connected in some way with motivation; that is, from (3), with having a desire. So (1), (2), and (3) together entail that there is some sort of necessary connection between distinct existences: moral belief and desire. But (3) tells us that there is no such connection. Believing some state of the world obtains is one thing, what I desire to do given that belief is quite another. (1994, p. 12; see also McNaughton, 1988.)

Here is a statement of the semantic problem, created by adding a few notes within a quote by Williamson. (The numbering and the brackets are mine; the rest is a direct quote.)

- S1. The meaning of a declarative sentence may provisionally be identified with its truth-conditions [which are precise, i.e., have a sharp boundary],
- S2. and its use with our dispositions to assent to and dissent from it in varying circumstances,
- S3. [and, for there to be a precise boundary reflected in use, we would need to have a disposition to assent to a sentence when applied to one member of a Sorites series, and dissent from it for the next, but we don’t have such a disposition, and so] the complaint is that the epistemic view of vagueness sets truth-conditions floating unacceptably free of our dispositions to assent and dissent. (From 1994, p. 205)

Williamson, of course, accepts my bracketed addition to (S1). The part of (S3) that I add is precisely what Williamson will reject. I add it to make it more explicit what the objection to the epistemic view is, so that below I can more easily show what it is that Williamson rejects. The moral problem arises from the entrenched Humean assumption

(stated in M3) that there is no necessary connection between belief and desire, but for moral realism to be true, there must be such a connection. The semantic problem is slightly different: it starts with the entrenched assumption that there must be a necessary connection between meaning and use (noted in S3, where Williamson says that it is unacceptable for truth conditions to float free from dispositions), and then shows that if there is a sharp objective boundary in the truth conditions for vague sentences (as indicated by the realist epistemic theory, as in my addition to S1), there cannot be a necessary connection between meaning and use.

I don't wish to dwell too long on this comparison, because there are specifics in the two cases that differ, though I do believe that further examination in another place may prove rewarding. I will note that while Smith's answer to the moral problem might possibly provide an answer to the semantic problem, it is not a promising avenue to explore. Smith solves the moral problem by appealing to a kind of ideal moral observer. This kind of answer has its analogies in semantics with Peirce's view of an ideal scientific community, which was revived in the early 1980's by Putnam. I do not have an argument against such a theory (not a good argument, at least), but because the view has not reached common acceptance, I shall not discuss it further.

There is another way to tell the normative semantic problem: Rather than focus on use, focus on perception. Concepts are precise, and we understand concepts. Let's assume that perceptions all have conceptual content. Hence, if there is a sharp boundary reflected in our understanding of a concept (an observational one, such as 'red' or 'heap'), it would have to be perceptible. But it is not perceptible – our perceptual mechanisms are not acute enough. This shows that concepts and perceptions are separate existences, which may be

an interesting anti-McDowellian argument. I think it works, but I cannot discuss it further here.

3.4 The Connection Between Meaning and Use

What can be said in favor of (2) in Williamson's supervenience argument?

'Lavender' is a public term, and so its meaning does not vary solely with one person.

There must be a set of truth conditions based upon the use of a term in a community; we latch on to them, but we don't know the specifics of use. If truth conditions did depend on one individual at a given time, then it would be difficult to account for how the person could ever err. Considered on its own, (2) does seem plausible (on the assumption that vague terms do have meaning). Williamson (1994, p. 216): "In some sense we create our language, but it does not follow that it is in every respect open to our gaze. Why should the boundaries of our terms not be invisible to us?" Even one who rejects bivalence or Williamson's claim that vague predicates have sharp boundaries can still accept that if the semantic model for the reference of 'lavender' must account for the use of the term by all individuals in the linguistic community, then (2) is quite plausible. If so, there is, contra (EC), no necessary connection between any individual's beliefs about 'lavender' and truth conditions. The reason one might accept (EC) in the first place is that 'lavender' is an observational term, applicable in casual settings. But if the truth conditions are a function of the use by the general public, then there is no reason to suppose that anyone could know the truth conditions (unless the term represented a natural kind). But we still use terms like 'lavender' successfully even though we do not know the precise reference because we still have some imprecise knowledge of its truth conditions. (Williamson

discusses how we understand vague terms without knowing their precise reference in his 7.6, and I accept his account there.)

Williamson himself phrases the most common objection to the epistemic view:

A common complaint against the epistemic view of vagueness is that it severs a necessary connection between meaning and use. Words mean what they do because we use them as we do; to postulate a fact of the matter in borderline cases is (it is charged) to suppose, incoherently, that the meanings of our words draw lines where our use of them does not. (1994, p. 205)

Williamson's response to this kind of objection is the following:

Every known recipe for extracting meaning from use breaks down even in cases to which vagueness is irrelevant. The inability of the epistemic view of vagueness to provide a successful recipe is an inability it shares with all its rivals. Nor is there any reason to suppose that such a recipe must exist. (1994, p. 207; also see 1996, p. 233)

In agreeing with Williamson about (2), the most I say is that (2) is *plausible* – if truth conditions are mysterious, then we can't know for sure that (2) is true. I would like to see more of an argument for (2), but I believe that the considerations given cast serious doubt on (EC): the motivation for accepting (EC) is to account for how truth conditions of observational must be dependent upon us (who are imprecise perceivers). But (2) is consistent with truth conditions for terms like 'lavender' depending upon speakers. I am satisfied with Williamson's argument, *taken conditionally*: IF terms like 'lavender' have truth conditions, THEN the truth conditions are unknowable. But why should we accept the epistemic view rather than an eliminativist view if truth conditions are so inscrutable?

In fact, in one place, Williamson's argument is almost explicitly in the conditional form. It rests on a rejection of other accounts of meaning:

The epistemic theory of vagueness makes the connection between meaning and use no harder to understand than it already is. At worst, there

may be no account to be had, beyond a few vague salutary remarks. Meaning may supervene on use in an unsurveyably chaotic way. (1994, p. 209)

If it is indeed true that meaning supervenes on use, then it must occur in some way. But no one else has given anything close to a successful account. On a supervenience account, our use picks out a range of truth conditions, and not a single precise truth condition. I believe that this helps a bit, because it is slightly more palatable to claim that use picks out a range of possible truth conditions than to claim that use picks out one particular truth condition. Still, supervenience is problematic, for reasons of higher-order vagueness: what exactly is the range of admissible uses? Or, if there is higher-and-higher-order vagueness, what exactly is the range of ranges? The supervenience theorist would need to explain how her account pertains to use, in such a way that the account does not fall prey to higher-order vagueness. This will prove to be as difficult, or even more difficult, than the epistemic theorist's task.

Though I believe that Williamson's argument is successful in demonstrating that his view is better than the others being offered, it does not persuade me to accept his epistemic view. For if acceptance that meaning supervenes on use leads us to believe that there are unknown boundaries for observational terms like 'heap', then we should then explore the question of whether meaning really does supervene on use, and whether there are any truth conditions at all.

My argument will be that if there are supervenience principles, we must know them. But we do not know any. Hence, there are no supervenience principles. On the assumption that for there to be heaps, there must be supervenience principles between heaps and grains of sand in order for there to be heaps, there are no heaps. It seems to me

to be much more plausible to deny the existence of heaps than to accept that there are unknowable bridge laws. Eventually, I shall argue that we want statements like ‘there are heaps’ to be true. And because there is no determinate fact of the matter connecting use to truth conditions, we can say that though we do not know any supervenience principles, we can stipulate some ourselves. Or, at least, posit that there are some, without going through the trouble of stipulating. If we just posit, and not stipulate, it might be that the actual instances of the supervenience principles are beyond our ken (we haven’t gone through the trouble).

3.5 Interlude: Sorensen’s Epistemicism

In two works (1988, 2001), Roy Sorensen argues in favor of the epistemic view. These books are filled with many ingenious arguments and humorous examples, which I shall for the most part not discuss. I believe that Williamson’s position is more clear and better argued. But Sorensen’s epistemic view differs from Williamson’s in a relevant way: Sorensen denies (2001, p. 178) that meaning supervenes on use. On Sorensen’s view, it is implausible to suppose that there are unknown supervenience principles. Instead, Sorensen holds a form of externalism by which meaning for all terms is not determined by use but by the objects and properties in the external world that the terms refer to. Williamson’s realism is a community-based realism, and on it, there are what Sorensen calls ‘relative borderline cases’ – meaning is relative to use. Sorensen’s own view is more of a direct realism – vagueness doesn’t depend upon cognizers, and there are what Sorensen calls ‘absolute borderline cases’. Sorensen argues that this move escapes objections against Williamson. Sorensen builds his case more out of rebuttals of

opposing views than positive arguments in favor of his own view, except for the logical snobitivist argument that we must keep bivalence. But as I have argued, maintaining bivalence does not commit one to the view that there are unknown sharp boundaries. I shall put aside discussion of this kind of externalist realism, because in Chapter Four I shall discuss a view of reference that seems consistent with Sorensen's account – Ruth Millikan's biosemantic theory (though Sorensen doesn't discuss biosemantics and Millikan doesn't discuss vagueness).

3.6 Vagueness and Realism

Let me recapitulate the status of the argument in this chapter. Arguments for bivalence, even if true, do not entail the epistemic theory. The revised margin for error principle, in itself, fails to motivate the theory, but Williamson's argument concerning the connection between meaning and use might. There are good reasons, stemming from the public nature of language, to believe that (EC) is false. If that argument is correct, then it may even sidestep the issue of whether (2) is correct – whether slight differences in use entail slight differences in meaning is unclear, but unimportant. What matters is that there is no known connection between meaning and use. What is left open at this stage is whether it is more plausible to claim, given all this, that the epistemic view is correct, or that there are no such things as heaps. I believe that the answer is intricately related to subtleties in the realism/anti-realism debate.

As I discussed in Chapter One, in 1983, Putnam claimed that vagueness shows that realism is false, but he underestimated the task of providing an anti-realist account. Most writing on the topic refers to Williamson's view as being realist, though he himself

doesn't use that term, as far as I know. Famously, the realist/anti-realist debate is difficult to characterize: there are perhaps more authors who claim that the realism/anti-realism debate is confused¹³ than those who take a position on one side or the other. Though I agree that there is no single position that all those who call themselves realists hold (likewise for anti-realists), there are some tendencies in the debate that I'd like to highlight. I hope to show that some insights arising from the debate may be used to reject the epistemic view, and I also hope that my discussion may shed more light on the debate itself.

Let's take it that Williamson has shown that if there are things such as heaps, and if meaning supervenes on use, then the epistemic view, or something like it, must be the best option. (Alternatively, if Sorensen is correct, then the epistemic view is correct even if meaning does not supervene on use.) I shall hereby put aside supervenience questions, and focus on the question of whether there are heaps and lavender candies.

I'd like to begin the process of answering this by asking a broader question: What is the best explanation of our use of terms like 'heap' and 'lavender'? The objection I'd like to consider is not whether the epistemic view has a poor account of the connection between meaning and use – that much, Williamson admits. Rather, I'd like to consider the possibility that the epistemic view has a poor account of what exists. The problem for the epistemic view may be not primarily in philosophy of language and semantics, but in metaphysics, and physics, too. Horgan says it well (1998, p. 178):

Given the broad outlines of contemporary physics, it would appear that even a complete physics-level characterization of the world would fail to

¹³ Rorty, Putnam, and Davidson, to name three.

provide any non-arbitrary way to precisely delimit the extension of terms like ‘tall’; and nothing one can point to in the special sciences suggests that facts specifiable in special-science vocabulary could play this role either. As far as one can tell, our best empirical theories of the world just do not posit hidden facts that determine precise boundaries where there is apparent vagueness.

If this is true, what is the best explanation for our beliefs about ‘tall’? In what follows, I do not state explicitly criteria for what it is for an explanation to be the ‘best’ explanation; such a project would take me well beyond the scope of the dissertation. I do, however, argue that some explanations are better than others, and the strength of my arguments below will have to rest on what I hope are common assumptions about what better and worse explanations are.

An analogy for the kind of argument I’d like to make is from a realism/anti-realism debate in meta-ethics. Gilbert Harman argues (1977) that the best explanation of our beliefs about moral claims can be given without positing the existence of any moral reality. Harman (1977, p. 120):

You can observe someone do something, but can you ever perceive the rightness or wrongness of what he does? If you round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it, you do not need to *conclude* that what they are doing is wrong; you do not need to figure anything out; you can *see* that it is wrong. But is your reaction due to the actual wrongness of what you see or is it simply a reflection of your moral “sense”, a “sense” that you have acquired perhaps as a result of your moral upbringing?

Harman answers that the best explanation of your belief in the wrongness of the act is not that there exists an objective property of wrongness that has caused your belief. Though Harman ultimately argues for a form of moral relativism, the foregoing argument is a strong argument in favor of moral skepticism or anti-realism. The approach I’d like to take with terms such as ‘heap’ is to argue that if the best explanation of our use of ‘heap’

makes no reference to the existence of heaps, then we have good grounds for saying that there are no heaps.

There are many other instances of this form of argument. If I walk out of my advisor's office believing that there is an elephant in the department lounge, how should I evaluate the belief when I realize that I have it? If I happen to walk back into his office, and notice on the bookshelf a book titled *How to Hypnotize Your Students Into Believing that There Is an Elephant in the Seminar Room or Department Lounge*, I have good grounds at that point for abandoning my belief that there is an elephant. The best explanation of my belief is not that there is an actual elephant in the lounge, but that there is no elephant and I have been hypnotized.

This same form of argument is perhaps the most effective means to argue against the existence of God – if we can explain individuals' beliefs in the existence of God without appealing to the existence of a supernatural being, then we have good reasons for denying that God exists. However, there is a difference in the two cases: *we* can understand the meaning of 'there is an elephant in the department lounge' because elephants exist (skeptical doubts about vagueness put aside for the moment) and the department lounge exists. Hence, the belief attribution should be read as: 'Avram believes that there is an *x* such that *x* is an elephant and *x* is in the department lounge.' For God, it is not so simple: how can one who denies the existence of God explain the belief attribution 'Yolanda believes in God'? If it is I who am stating the belief attribution 'Yolanda believes that there is an *x* such that *x* is God', and I do not believe that God exists, and I am correct in this, then this statement will need further evaluation if it is to have a chance of being true, given that my statement will have a non-referring term.

One potential way to understand this belief attribution is to claim that Yolanda believes that there exists a thing which fits a certain description (of being omnipotent, etc.). But it is not likely that this form of explanation will be satisfactory, either. First of all, it is unclear whether we really understand the terms of the explanandum (omnipotent, etc.) and use them to refer to anything. And it is unclear whether we can give a complete description which is in fact the description of God under which Yolanda believes that such a thing exists.¹⁴

Fortunately, I do not have to solve this thorny issue at present. It seems that some account of the concept of ‘God’ can be given in order for the atheist to properly make her ‘inference to the best explanation argument’ against the existence of a referent of ‘God’. And it is worth noting that whether or not one believes in God might vary with this very point. Perhaps the best argument in favor of the existence of an intelligent creator is that it is a more likely explanation of the existence of the universe than that given by physicists. I shall employ the ‘God’ example below, but at issue for my purposes here is what the best explanations of the use of terms like ‘lavender’, and ‘heap’ are. If those explanations can be given without a claim that there exists a property that lavender refers to, then we can say that statements such as “Candy 1 is lavender” are not true, hence avoiding Sorites paradoxes. What is worth discussing is the status of sentences such as

(C) ‘Lavender’ refers to lavender.

¹⁴ There seems to me to be much philosophical work yet to be done in explaining belief attribution statements. Perhaps the previous paragraph reveals that I reject the notion, stemming from Quine, that co-referring terms *cannot* be substituted *salva veritate* within belief contexts. If I can truthfully say ‘Lois believes that Superman can fly’, then I can also truthfully say ‘Lois believes that Clark Kent can fly’, because *my* use of ‘Superman’ and ‘Clark Kent’ do refer to the same person. This, of course, is an issue best left for another place; though my argument in the text does rest on my own aberrant view of belief attributions, I believe that similar kinds of argument can be made under a more traditional theory of belief attributions.

This statement, taken by itself, is not unambiguously true. For it to be true, we would need assumptions about how to translate from the object language into the meta-language. ‘Lavender’ might be a term of a foreign object language such that it is not translatable as the meta-language’s ‘lavender’, so perhaps the statement is false. Or ‘lavender’ might not refer at all in either object or meta-language, and if it does not, (C) is not true because the last word in it is non-referring. So in giving explanations of terms like ‘lavender’, it will be of little help to use statements such as (C). Furthermore, even if (C) is true, it would not help as an *explanation* of ‘lavender’ – it adds no content to our understanding of ‘lavender’.¹⁵

Sentences such as (C) prove to be important in characterizing the realism/anti-realism debate. I do not claim that the view of the debate that I shall ultimately arrive at will accurately characterize everyone, or even most people, who are involved on either side of it. My underlying purpose is to get the tools needed to resolve Sorites paradoxes, and not to have a specific analysis of the realism/anti-realism debate. I will not attempt to produce an analysis that maps on well to all the nuances of the debate. If you’d like, you can understand the distinction I make between realists and anti-realists in what follows as my stipulative definition (and thus not subject to dispute), and not a description of a real debate. I’d like to think that these comments cast light upon an actual debate, but what

¹⁵ This isn’t quite true. (C) can be a statement that ‘lavender’ is (potentially) a referring term, and not a connective. One wouldn’t say:

(O) ‘or’ refers to or.

So, in declaring (C), one is giving a piece of an explanation of the use of ‘lavender’ – it is not a sentential connective. In fact, it may even say something philosophically controversial about ‘lavender’: that it refers to a property. Nominalists would reject this. Throughout this section, I say that terms like ‘lavender’ refer to properties. If one is a nominalist, one can replace ‘lavender’ with the proper name of an object, as in ‘Mt. Baldy’ refers to Mt. Baldy.

matters most is whether the view I argue for, which I declare to be a certain kind of anti-realism, is correct.

Here is a bold (but vague) first attempt at a thesis about the difference between realists and anti-realists:

- (R1) Realists believe that we cannot explain our use of terms like ‘lavender’ without some kind of statement like (C) – ‘lavender’ really does refer to a real, objective, mind-independent property of lavender. Anti-realists, on the other hand, argue that we cannot take the reference of terms like ‘lavender’ at face value – the metaphysics of the notion of ‘real property’ is fraught with problems. (C) may still be true on some anti-realist accounts, but the explanation given for its truth will not involve correspondence to a real external property.

I have in mind here Devitt (1984) as the realist. His definition of realism is:

“Realism: Tokens of most current common-sense, and scientific, physical types objectively exist independently of the mental.” (1984, p. 22)

I am unsure who fits the characterization of anti-realism, though I shall discuss anti-realism below in 3.6. On this view, the difference between realists and anti-realists is in how they go about answering questions such as “What is the best explanation of the use of the term ‘lavender’?” Hence the realism/anti-realism debate is not about metaphysics (contra Devitt), and is not about truth (contra Dummett), but in *methodology*. The terms of Devitt’s characterization can be held by philosophers with much different approaches to metaphysics (in particular, Wittgensteinian minimalists). Contra Dummett, bivalence can be rejected by those whose methodologies are similar in spirit to Devitt’s. Realists assume at the outset that for at least most of our terms, we refer to a real mind-independent property or object. Anti-realists do not make such an assumption at the outset, though ultimately they may agree in a roundabout way that there is no sense in saying that objects are mind-independent.

The first objection to realism is that it requires an ineliminable *faith* in our ability as knowers. Why assume at the outset that humans refer to anything with our terms? To respond, realists may appeal to the success that humans have in navigating around the world, using language as we do. This might be bolstered by an appeal to evolutionary considerations that show that it is no accident that our terms refer – if they did not, we would not have survived as a species. I’m not persuaded by this response, unless more specific considerations about how exactly an explanation of our evolutionary success requires that language is used to refer to objects and properties. What I would require for such an explanation is a well spelled out scientific account, for particular terms, of how the term refers to a real object in the environment. I would want the realist to engage in discussion about the best explanation for the use of the term. If, at the end of the explanation, the existence of a property such as lavender is required for the explanation of our use of ‘lavender’, then it might be acceptable. (I discuss Ruth Millikan’s view in some detail in Chapter Four.)

Hence there is room for a distinction between two kinds of realists: ‘end-of-explanation’ realists and ‘beginning-of-explanation’ realists.¹⁶ I’d like to dismiss, for the moment, realists in the latter category, for saying “‘lavender’ refers to lavender” at the beginning of an explanation of lavender will be no explanation at all. (Though I shall ultimately argue that this kind of attitude may be plausible, but only if it arrived at through anti-realist means.)

¹⁶ In the former camp I place Millikan (1984, 1993) and, for other reasons, Richard Boyd (1980). In the latter I place Arthur Fine (1986).

An analogy might be helpful here. Nicholas Sturgeon argues (1988), contra Harman, that the existence of moral properties *is required* for the proper explanation of our moral beliefs and of actual human behavior. Sturgeon gives several examples, one of which is (1988, p. 245):

An interesting historical question is why vigorous and reasonably widespread moral opposition to slavery arose for the first time in the eighteenth and nineteenth centuries, even though slavery was a very old institution; and why this opposition arose primarily in Britain, France, and in French- and English-speaking North America, even though slavery existed throughout the New World. There is a standard answer to this question. It is that chattel slavery in British and French America, and then in the United States, was much *worse* than previous forms of slavery...

Harman says that assumptions about moral facts seem “completely irrelevant” in explaining moral observations and moral beliefs, but on its more natural reading that claim seems pretty obviously mistaken about these examples.

Sturgeon then goes on to argue that these moral facts (such as about slavery being worse) are not reducible to other facts, but their existence does supervene on natural facts. Hence Sturgeon advocates a non-reductive naturalism.

I do not wish to get into all the details of Sturgeon’s response to Harman, but it is worth pointing out that Sturgeon’s argumentative structure in (1988) is extremely similar to Williamson’s (in 1994, Chapter Seven, and 1997), as we shall see in a moment. There seems to me to be a response to Sturgeon that he does not give enough consideration to: what if we can give an explanation of what it is for slavery to be *worse* in one situation than in another without adhering to any *normative, moral* facts? It is quite plausible to suppose that one can explain it in terms of the kinds of treatment of slaves, and our ‘moral sense’ about this treatment. Sturgeon admits that his argument doesn’t refute moral skepticism, but that he does not have to (1988, p. 237 and p. 253). I believe that

Sturgeon is wrong about this burden-of-proof claim – if the best explanation of any moral statement can be given without reference to any moral properties, then moral skepticism is quite plausible.

The lesson I'd like to draw is the following. Assuming that what he says is historically accurate, I would admit that Sturgeon is correct that the explanation 'a reason why there was opposition to slavery in North America is that it was worse than elsewhere' is a good explanation, given what one is expecting out of the explanation when the original question about objections to slavery arises in a natural context. We can grant Sturgeon that, but then it seems plausible to ask whether there is an explanation of this explanation that does not include an irreducibly normative component. If so, that would give us a better explanation of our beliefs, and one that does not employ irreducible moral properties. This is the kind of explanation that the philosopher is looking for when asking about whether moral properties exist. *Sub specie philosophiae*, we demand a more complete explanation of the belief. And it seems plausible to suppose that there can be such an explanation that is consistent with the skeptical hypothesis that is in fact a better explanation than the one that leaves moral facts unreduced and unanalyzed.¹⁷ However, the questions that arise next are: why demand such an

¹⁷ This is a rejection of one of Sturgeon's main theses, and I do not attempt to give a full response here. A full response would require a detailed analysis of the fact/value distinction. However, I should note that one of Sturgeon's main arguments seems to fail. Sturgeon analogizes moral theory with microphysical theory (1988, pp. 251-252), and argues that if this form of argument leads us to be moral skeptics, then the argument should also apply to microphysical theory – but of course we do not want to be skeptics about microphysical theory. I believe that this analogy is misleading: in the case of beliefs about moral facts, we have a very plausible alternative skeptical explanation that can account for why we have beliefs about moral facts without positing the existence of moral facts. Given the problems of moral theory, at the very least, this skeptical explanation competes quite well against the moral realist explanation. With microphysical theory, on the other hand, we have no plausible alternative explanation of why we have

explanation? Why would such an explanation be *better*, even if it could be given? I do not wish support *moral* skepticism by answering this question here, though I believe the moral skeptic has good resource for responding.

For the case of vagueness, the reason why there is a demand for a reductive explanation should become apparent. We start with our use of terms like ‘heap’. We realize that there are Sorites paradoxes, and there is a seeming incoherency if we generalize principles of our use of the term. At this point, we should demand an explanation of our use of the terms. Williamson’s response is that of the beginning-of-explanation realist. After arguing that vagueness is ignorance, he writes (1994, p. 202):

The epistemic theorist is sometimes asked: of what fact could we be ignorant? The answer is obvious. In the present case, we are ignorant either of the fact that TW is thin or of the fact that TW is not thin... there is no general requirement that vague words be definable in other terms.

This kind of argument strikes many as implausible. I hope that the following considerations demonstrate why. We must ask whether the metaphysical picture that the epistemic view paints is mysterious. Williamson considers this objection:

A further question is sometimes asked: what kind of fact is the fact that TW is thin? The question is a bad one, for a reason unconnected with vagueness. We do not have a proper taxonomy of facts, not even of precise ones. When a taxonomy is provided, it will be time to say what kind of fact a vague fact is. (1994, p. 202)

Though I believe that Williamson is correct that we do not have a complete taxonomy of facts, it does seem that certain claims of fact have more legitimacy than others. We have

beliefs about microphysical entities, and so we should not fall so quickly into skepticism in that realm even if the microphysical explanations that we give are problematic.

Williamson’s argument is similar to Sturgeon’s. He accepts the claims at face value, and says that if the conditions were different such that it is not a heap, then we wouldn’t call it one. He argues that disputing borderline cases will get us into skepticism about the clear ones, but we do not want to be skeptics about clear ones.

something of a taxonomy – facts about the color of gold are facts about a certain natural kind based in its subatomic structure; facts about musical scales are facts about soundwaves; facts about balls and strikes are facts about baseball conventions. The first two of these examples are scientific discoveries, and the third is a stipulation. What kind of fact is a fact about whether TW is thin? Of course, we know that facts about thinness are facts about one's girth, but we don't have a precise way of specifying the higher-level facts based upon the lower-level facts. We haven't had any scientifically legitimate discovery about the precise extension of thinness. In the three examples I give, they are precise in the relevant ways.

When something is claimed to be a fact, and we do not have a semblance of what kind of fact it is, what we should do is *investigate*. Hence, I think something based in Wright's approach is exactly right here: when we are faced with a situation in which we have no idea about the factual status of a claim, we should not make any positive assertions about it. If Wright's logical maneuvers are legitimate (and if he can provide an account of higher-order vagueness), then I believe that there is no reason to accept Williamson's epistemic view over Wright's: if something is a mystery, we should not make positive assertions about it. Extending Wright's account beyond where he himself takes it, I claim that in the face of mystery, we must investigate the mysterious fact further. If we can reduce it to some known kind of fact, then we can grant it legitimate status. If we cannot, then what matters most is whether our *best* metaphysical explanation indicates that it is a fact or not, and if there is a competing hypothesis of why we might believe that it is a fact when it is not. Perhaps the investigation can be empirical (the possibility that evolutionary semantics might make a contribution to semantics is

evidence that empirical considerations may be relevant), and likely at least some of the investigation will be *a priori*. But investigate we should.

So, as part of an investigation, we should ask, once more: what is our best current explanation of terms such as ‘heap’, ‘lavender’, and ‘thin’? Once we begin to answer this, it seems that the best explanation requires no objects/properties to which they refer. The kind of explanation I have in mind is that once upon a time, someone saw something with his or her limited perceptual mechanisms, and later on had a similar perception, and instigated a practice of using a term to make an official connection between the two perceptions/items. So a word has come into practice, but without there being a precise object/property to which the term refers. Perhaps an example with ‘Mt. Baldy’ would help – proper names might be an easier avenue to demonstrate the point. What best explains my use of ‘Mt. Baldy’? Is it that a precise portion of the landmass of the San Gabriels leads me to use the term? No; we see a summit (i.e., the highest point in the general area), and have some reasons for speaking of some vague area around the summit (this can be analogized with paradigm cases for properties), so even though we know that we can’t define ‘Mt. Baldy’ precisely, we use it anyway. The land, of course, exists, but there are no natural demarcations, and no demarcations set up by our use of the term. In fact, the very reason why we have terms like ‘Mt. Baldy’ is that we cannot easily specify a precisely demarcated area including and surrounding the summit. If this is correct, our use of ‘Mt. Baldy’ is *essentially* imprecise. The best explanation of our use of ‘Mt. Baldy’ is not that there is a precise object ‘Mt. Baldy’ that plays some causal role in our use of the term.

Note that my argument isn't really about natural kinds, as is Wheeler's (1975). A convincing appeal to something's being a natural kind is a good way to show that it's a real property, but it's not the only way. I do not wish to delve into debates about natural kinds, but it seems as if terms such as 'ball' and 'strike' can be given a precise reference even if they do not represent natural kinds. We make, and use, conventions all the time. Though perhaps conventional words cannot be associated with a human-independent natural kind, they can be associated with a relational property (relating some aspect of the world to something about us) that is naturalistically respectable.

My argument here is incomplete without a more positive picture of the best explanation for our use of terms. I shall save the positive picture for Chapter Five. But in the context of the epistemic view, I believe no further argument is needed to make the point that there are no precise boundaries in nature that correspond to our use of the term 'Mt. Baldy'. For Williamson himself agrees with much, though not all, of what I write above. At the end of his book (1994, p. 269), he writes:

The nominalist suspects that properties, relations and states of affairs are mere projections onto the world of our forms of speech. One source of the suspicion is a sense that we could just as well have classified things differently. Vagueness is indeed one manifestation of the fact that our classifications are not fixed by natural boundaries. The vagueness of singular terms suggests that if the nominalist conclusion did follow, it could not exempt the category of objects. The boundaries of a particular mountain reflect our language no less than do those of the property of mountainhood.

(Williamson ends by noting that although the boundaries of a mountain are determined by our word for it, this is not to say that language created the mountain.) Given that "our classifications are not fixed by natural boundaries" – the sense of 'natural' that he intends is not spelled out – they must be fixed by our use of terms. But if this is where the

dialectic has gone, it seems as if the eliminativist wins, because there are no precise boundaries fixed by our use of terms (and there are no imprecise boundaries fixed by our use of terms, on the presumption that supervalueation and other moderate indeterminist views fail, as discussed in Chapter Two). On the contrary, it is because we, as language users, do not fix precise reference with observational terms that we brought the terms into language in the first place. I cannot see how one could motivate this view of reference except if one has argued that all alternative views, including eliminativism, have failed. In fact, I too believe that eliminativism fails, but that does not cause me to recoil into being a ‘beginning-of-explanation realist’; instead, I shall take this as a dialectic which will lead to what I hope is a nuanced and correct anti-realism.

There are two problems with eliminativism. The first concerns its consequences, some of which Williamson discusses (all too cursorily) in his Chapter Six, that I shall discuss in my Chapter Five. There are other various reasons for why there must be a Mt. Baldy. Perhaps a principle of charity must be true, such that it must be the case that most of our utterances to be true. If the facts of the matter do not determine a precise referent for ‘Mt. Baldy’, why cannot *we* determine a precise referent? I take this kind of consideration seriously, will eventually agree with the conclusion (though I do not base it upon the principle of charity).

The second objection, which I have already touched upon, is that the terms of my explanation of ‘Mt. Baldy’ themselves are problematic. In my explanation, I use terms such as ‘landmass’, ‘peak’, ‘visual image’, ‘we’. Surely these are fraught with problems (due to vagueness or otherwise). How good is this explanation I have given? Is it any

better to give such an explanation than to leave ‘Mt. Baldy’ unanalyzed, and thus maintain some truth conditions for its application that often do hold?

Importantly, the answer is *no*. Though the terms of the explanation are problematic, we should not recoil to an acceptance of first-order terms such as ‘Mt. Baldy’ which have been shown to be Sorites-susceptible. To see why, let’s return to the God example. Probably, the best explanation the atheist can give to undermine Yolanda’s belief in the existence of God appeals to mental states, persons, causes, ideas, etc. For reasons not connected to vagueness (and some connected to vagueness), these terms are problematic – the metaphysics of beliefs and ideas and causes are all problematic in many ways. We need some means of rejecting the reference of terms in a discourse despite having no full scientific theory of everything. The reason why we abandon belief in the veridicality of some discourses seems to be the very reason I give for rejecting reference for terms like ‘Mt. Baldy’: the *best*, albeit incomplete, explanation of the terms of the discourse in question make no reference to objects to which the terms refer. Why do most philosophers and scientists not believe that the important terms used in astrology refer? Because the astrology runs into trouble when combined with other empirical considerations, and the best explanation of the discourse of astrology does not grant that its terms refer to real phenomena. Though astronomy has not been a completely reduced, final science, there is nothing about the deficiencies in scientific accounts in astronomy that would make us return to astrological explanations.

Williamson responds (1997, pp. 330-33) to an objection similar in spirit to mine that claims that in order for there to be precise reference for an observational term, its precision must be ‘cognitively accessible’. Williamson rejects this view, rightly, on the

grounds that the demand that all truths must be explicable in a precise language depends upon an overly scientific attitude. But the mere request for an explanation of terms when the term is Sorites-susceptible in itself seems not to be overly scientific, and the claim that we can say that the term does not refer when our *best* explanation of the term does not indicate that the term refers is not in itself overly scientific either. The hope is that the explanations of our terms will build *towards* a scientifically or empirically viable explanation of everything. But there is no requirement that we be at a final stage of inquiry in order to explain away terms in a discourse.

Much of the motivation for the eliminativist view I am arguing in favor of here is based upon the analogy with the best explanation of 'God'? But there may be a disanalogy: the terms used in my best explanation for 'Mt. Baldy' are themselves observational terms. The same *kind* of term is used in the explanation as in what is being rejected. In the explanation concerning someone's belief that 'God' refers, I use terms such as 'belief', 'we', etc., but I do not use terms with the same supernatural bent as the one whose reference is being denied. There may be a concern that unless I can give a scientifically legitimate story (as Williamson might demand for a rejection of the macro-level terms), I will face a problem. Though I do not have a complete response to this point, the general form of my response is that for terms in the explanations themselves, Sorites paradoxes are not relevant to questions at hand. When, and if, they become problematic due to Sorites, we can deal with them. But we need some place to stand upon. All terms come with promissory notes that their reference can potentially be explained, and the terms in my explanation of 'Mt. Baldy' are less Sorites-susceptible than 'Mt. Baldy'. For instance, 'landmass', being a mass term, does not depend upon a

precise spatial boundary. There will be an issue with how many atoms together are enough to be considered a landmass. But there are two important things to note about this: first, it seems that whatever the outcome of this question is, it will not turn out in such a way that shows that ‘Mt. Baldy’ does refer. Second, though this may lead towards an eliminativist rejection of ‘landmass’, what we are doing at this point is *science*, as the explanation of one’s use of ‘Mt. Baldy’ is an empirical investigation. And science ought to proceed by taking a hypothesis, confirming it, and then questioning the terms of the hypothesis, and thus, gradually, better and better theories will develop. Sorites paradoxes, as Quine understood, set constraints on proper scientific investigations.

I have hardly mentioned anti-realism in the preceding remarks, but I am now in a position to say something about it. Some anti-realists despair of our ability to give satisfactory lower-level explanations of our terms that will be satisfactory. I have in mind here the first three chapters of Putnam, 1981¹⁸ — though notably, Putnam calls himself there an *internal realist*, and not an anti-realist. I consider Putnam an anti-realist because he does not assume the realist premise of the existence of mind-independent objects at the outset of his investigations. Anti-realists, in a Wittgensteinian spirit, may then choose to become minimalists about reference, and this may lead to an acceptance of observational terms as referring to an objective world (see especially Wright, 1992b). This kind of anti-realism has complicated the realism/anti-realist debate, because anti-realists may end up accepting the letter of Devitt’s paradigmatic realist claim, though in a much different

¹⁸ See my “Brains in a Vat in a World” (unpublished) for further elaboration of this theme.

spirit. I believe that a focus on methodology is the right way to distinguish between these two classes of philosopher.

These anti-realists may resemble beginning-of-explanation realists, because they have rejected the notion that there can be a full explanation of our terms, and so they accept the legitimacy of ordinary vocabulary. This may be a reason why Williamson himself does not proclaim himself as a realist: it is unclear whether he is a beginning-of-explanation realist or a minimalist anti-realist. The only difference between the two seems to be a difference in the amount of metaphysics the philosopher explicitly *rejects* en route to accepting ordinary vocabulary. Though these two views only differ from each other by matter of degree, they are sharply different from end-of-explanation realism, which itself is very similar to reductionism: if our best empirical explanation of our use of a term indicates that it refers, then it refers.

My view is different from all these views. Against the beginning-of-explanation realism and anti-realism, we should not assume at the outset, or recoil after an attempted investigation into believing (respectively) that terms like ‘heap’ refer to heaps. Against end-of-explanation realism (not that anyone in the literature holds this reductionist view), I deny that our best explanation of such terms succeed in providing a reason to believe that terms such as ‘heap’ refer. I have not taken a position about whether we can get to a final explanation of our terms. Instead, I have argued, like the anti-realists, at the outset of our investigations both empirical and *a priori*, we should not assume that our vague terms refer and we also should not assume that our best explanations, though not final, indicate that they do not refer.

The arguments I have given in this chapter that favor eliminativism over realism are ones that naturalists should accept, and my project might appear to be a strongly naturalistic one. However, in the next chapter, I shall argue that naturalized semantics fails. I claim that it should be no surprise that the facts of the matter (concerning language use and the structure of the external world) do not determine a precise reference for our vague terms, because the former are non-normative facts and the latter are normative ones. (This may even call into question scientific language.) Vagueness, as I have argued, is not a special case of the indeterminacy of reference. It is simply *noticed indeterminacy*. Sorites paradoxes simply reveal potential problems that lurk underneath the surface of all language use. It is more than just, in Horgan's terms, that there is kinkiness in vagueness; there is kinkiness pervading all language.

Oth. "Swear thou art honest."
Des. "Heaven doth truly know it."
Oth. "Heaven truly knows that thou
art false as hell."
Des. "To whom, my lord? With
whom? How am I false?"
Oth. "Ah, Desdemon. Away, away,
away!"
- *Othello*, Act IV Scene ii

CHAPTER 4. VAGUENESS AND NORMATIVITY

In the previous chapter, I argue that there are no natural boundaries that determine a precise reference for terms like 'red', 'heap', and 'Mt. Baldy'. I also argue that our language does not determine precise boundaries where external nature does not. In the first two chapters, I argue that for terms to have any reference at all, they must have a precise reference – there must be some precise fact that determines the reference. Hence my conclusion, to this point, is that terms such as these do not refer. This argument was based upon a plausible assumption (which I shall ultimately reject) that meaning supervenes on use, which should be accepted by all naturalists, and some non-naturalists. Vagueness makes it *apparent* what is the case even for non-vague claims – that reference is not determinate.

Though this view of vagueness is not common, I'm certainly not the first to argue for the indeterminacy of meaning. Quine (1960) is perhaps the best known advocate of this thesis. His famous 'gavagai' example purports to show that there is more than one possible meaning for a term, and the facts of the matter not determining which of the possible meanings is the correct one. For reasons I shall not delve into in detail, I believe that his 'gavagai' example (and other examples similar to it) are unsuccessful. In short, it fails because it relies on a controversial assumption that a linguist cannot discriminate

between two possible meanings for ‘gavagai’. Quine only allows the linguist to observe the behavior of the language-users. Of course, that is all a linguist can do without becoming immersed in the language. But Quine does not permit introspective evidence to play any role in the determination of meaning, and this strikes many who are less behavioristically inclined than Quine as overly strict. Vagueness makes the indeterminacy thesis much stronger, because we assume from the start that there is no way, even for a sophisticated linguist viewing linguistic behavior *sub specie philosophiae*, to determine which of two (or more) precise propositions (if any) a speaker is expressing, even given all the possible evidence of both behavior and internal mental states.¹

The purpose of this chapter is to cast my skeptical conclusions about vagueness within broader discussions about the nature of meaning and normativity. The problems that we face because of vagueness should not be surprising, given the other problems, and so my conclusion should not seem as radical. Importantly, casting the problems of indeterminacy in the context of these other issues will help guide us towards a resolution to problems of vagueness, and ways to avoid the skepticism/eliminativism argued for in the preceding chapters.

In 4.1, I discuss both Kripke’s Wittgenstein (following Blackburn 1993, I shall call this character KW). Though I am somewhat concerned about Ludwig Wittgenstein himself (LW), and I shall make some remarks about him, I am not ultimately interested in which interpretation of LW is most true to LW’s own intentions. I shall use the rule-

¹ A similar point is made by McGee, 1998.

following considerations, and extend them, to further elaborate on the argument given in the previous chapter. I do not intend to say very much original about LW's view, except in the extent to which I apply the rule-following considerations to vagueness. I shall then argue in 4.2 that there is an is/ought gap in semantics that renders naturalized semantics untenable. In 4.3 I discuss Ruth Millikan's view as an example of naturalized semantics, and though I criticize some specifics of her account, I argue that some similar considerations can be given against all forms of naturalized semantics. In 4.4 I discuss a plausible way of dealing with the is/ought gap in ethics, and suggest that a similar kind of account can be given for semantics. This forms the first part of an argument that a non-naturalized semantics may be tenable, and that we need not be semantic eliminativists in the face of problems of vagueness and indeterminacy.

4.1 Kripke's Wittgenstein, Vagueness, and the 'New Riddle of Rule-Following'

On the surface, there are several similarities between problems of vagueness and problems of rule-following that KW elaborates upon. The challenge posed by Sorites paradoxes is that even if we were to know all the facts about a language user's linguistic behavior and dispositions to use terms like 'lavender' and all the facts about the external world that lead to the user's use of 'lavender', we will be unable to tell whether certain candies ought to be called 'lavender' or not. Armed with the problem of higher-order vagueness that (I argue) refutes moderate indeterminism, I extend this argument to include all candies, and hence advocate a radical form of semantic indeterminism. On KW's rule-following considerations, which Kripke refers to as 'Wittgenstein's Paradox', "The important problem for Wittgenstein is that my present mental state does not appear

to determine what I *ought* to do in the future.” (1982, p. 56. All Kripke citations in this section are from his 1982.)

I will not rehearse all of KW’s arguments, but it will be helpful to summarize the basic problem.² Suppose that I have never calculated $68+57$ before (see Kripke, pp. 8-10). When I do so, I get the answer ‘125’. I encounter a skeptic who suggests that the answer I should have given is ‘5’. The skeptic claims that in the past, I have used ‘+’ to denote the following *quaddition* \oplus function:

$$\begin{aligned}x \oplus y &= x+y, \text{ if } x,y < 57 \\ &= 5 \text{ otherwise}^3\end{aligned}$$

What fact about my past usage can be cited to demonstrate that in the past, I had always meant addition rather than quaddition with ‘+’? If we try to provide an interpretation of the rule for ‘+’ that we have been using, we may be faced with another skeptical inquiry about that interpretation. Writes LW (*Philosophical Investigations* §198):

Any interpretation still hangs in the air along with what it interprets, and cannot give it any support. Interpretations by themselves do not determine meaning.

The point is that any time we try to state what a rule consists in, that statement will itself be in need of interpretation, and so we will not have proven that we mean *plus* and not *quus*.

² My characterization of Kripke’s argument borrows from Hale, “Rule-Following, Objectivity, and Meaning” (1997).

³ I believe there is a problem in stating the problem this way, given that it uses the ‘+’ symbol to express what I may have used to mean by the ‘+’ symbol. This may be part of the problem in interpreting Wittgenstein; perhaps there is no way to phrase the supposed paradox. I will assume in what follows that Kripke’s phrasing is acceptable, though I have some qualms about doing so.

Other critics have analyzed this portion of LW's argument differently from Kripke. Although I am most concerned here with Kripke's analysis, it is worth mentioning the primary alternative, which I shall call the *Oxgenstein* view (see Baker and Hacker 1984, McDowell 1984). On this view, Wittgensteinian is not stating a skeptical problem. Instead, he is using rule-following considerations as a *reductio* of the notion that rules of language are in need of an interpretation. Rather, we have a practice to use language and there is nothing wrong with the practice. On this view, Wittgenstein is arguing against philosophical attempts to explicate the meaning of our terms. There is much to commend about this view, and about it as an interpretation of Wittgenstein – my own interpretation of Wittgenstein is mostly in accord with it. But as I have argued in the previous chapter, we should not uncritically accept our first-order language when problems with it arise.

A possible Wittgensteinian solution to Sorites paradoxes would be to claim that Sorites paradoxes only occur in contexts within philosophical discussions. In such contexts, language goes on holiday, and we need not be concerned about problems that occur within that abnormal language game. One reason I believe this kind of anti-philosophical solution fails is that it is unclear when, exactly, language goes on holiday. There may be a Sorites paradox on 'Language has gone on holiday at time x'. Importantly, ordinary language is not entirely first-order. There are theoretical intuitions about language that ordinary speakers possess that come to the fore in Sorites paradoxes, and that influence normal usage. We should not simply exclude Sorites paradoxes and other problems with semantic indeterminacy as mere philosophers' fancy. Furthermore,

we should hold out hope that semantic analysis does lead to further understanding of our language.

On KW's view, the rule-following considerations show that there is no fact of the matter about whether I am following a rule (cf. 68), and so there is a strong threat of scepticism. According to Kripke, the rule-following considerations connect with Wittgenstein's private language argument. If there is only one speaker of a language, the fact that the speaker's internal states do not determine the speaker's meaning indicates that nothing else indicates this. But according to KW,

All that is needed to legitimize assertions that someone means something is that there be roughly specifiable circumstances under which they are legitimately assertable, and that the game of asserting them under such conditions has a role in our lives. No supposition that 'facts correspond' to those assertions is needed. (78)

Statements are legitimately assertible, as KW proceeds to argue, when there is a community:

There is no objective fact – that we all mean addition by '+', or even that a given individual does – that explains our agreement in particular cases. Rather our license to say of each other that we mean addition by '+' is part of a 'language game' that sustains itself only because of the brute fact that we generally agree (97).

Blackburn, in "The Individual Strikes Back" (1993), argues against KW that including features of community agreement does not help at all. He argues that the same sceptical arguments that the solitary language user's terms have no determinate meaning also may be used to show that even communities of speakers who take themselves to be agreeing about their practices in the language game also do not speak language with a determinate meaning. For what can show that the members of the community are in agreement? Any effort to show that they are in agreement will require that they have

some independent characterization of what it is they agree upon, and this is precisely what is put into question by KW's original sceptical arguments. Blackburn writes (1993, p. 223): "My community may all suddenly start saying that $57+68=5$, but this does not make me wrong when I continue to assert that it is 125." If this is correct, and I believe it is, then a solution other than KW's is needed to reject KW's sceptic.

A possible straightforward solution to the rule-following paradox is to claim that although we ourselves cannot state the rule, we still have dispositions that allow us to follow rules. Hence there are facts about us, albeit unknown to us, that indicate what we mean by a term. Kripke makes two arguments (22-37) against dispositional accounts: First, for addition, we cannot have dispositions to add numbers so large that it would take more than a lifetime to do the calculation. Second, sometimes we are disposed to make mistakes. As Hale (1997, p. 372) summarizes, "The dispositional proposal fails to capture the essentially *normative* aspect of meaning... there is no room for a needed contrast between the answers I *would* have given and the answers I *should* have given."

Blackburn himself (1993, pp. 219-221) argues that a dispositional account of one's continued language use may be promising – Kripke's arguments against dispositional accounts are too brief and have been criticized elsewhere in the literature (See Papineau, 1987). Against the first objection, we are disposed to give answers to mathematical problems that do not require years to solve, and for problems that would take too long to solve, it still is the case that we have a disposition to reiterate procedures that we are disposed to use, even if we cannot in practice realize this iterated disposition. Blackburn gives a successful analogy: let's assume that a certain brittle piece of glass is disposed to break when struck against hard surfaces. Let's assume that it would

decompose before reaching Alpha Centauri, and hence it is not physically possible for it to break when struck against a rock on Alpha Centauri⁴. But this should not disqualify it from having the dispositional property of brittleness. Just because a certain calculation or event is physically impossible does not entail that no dispositional account can be given to cover these cases. Against the second objection, Blackburn agrees with Kripke that we may have dispositions to be mistaken in certain cases. But Blackburn notes that we also have dispositions to withdraw answers and substitute other answers upon further investigation, and this may provide an error-proof dispositional theory.

I am unsure whether these considerations rescue the dispositional account. But it appears that vagueness provides a further problem for it. We are disposed to call candy #1 ‘lavender’. But this does not suffice to provide precise truth conditions for ‘lavender’. What exactly are our dispositions to use that term? To claim that we have dispositions to use a term to describe certain items requires that one specify what those items are; but because of higher-order vagueness, that is just what we cannot do. Which is the last item in a Sorites series for which we are disposed to call the object “lavender”? There seems to be no good answer to this. Furthermore, if we attempt to get around this by positing a degree-theoretic account of dispositions – that we are slightly less disposed to use the term as we move gradually down the Sorites series – we still will run into the same problems facing degree-theoretic accounts. Namely, on what grounds can one give to claim that we are 71% disposed, rather than 72% disposed to use a term to apply to certain object? Though the notion of a disposition does seem to play an important role in

⁴ Are there rocks on Alpha Centauri? I don’t think so.

explaining our use of terms, proponents of dispositional accounts must still find a way to resolve Sorites paradoxes.

My own interpretation of LW agrees with Kripke's that LW relies upon notions of custom and practice. LW does not explain exactly what he means by these notions. Contrary to Kripke's interpretation, and in accord with an important part of the Oxgenstein view, I believe that the rule-following considerations were not intended by Wittgenstein to make a sceptical point about the possibility of semantic normativity. Rather, Wittgenstein was making an argument about the necessity of having a outward practice in grasping a rule; merely thinking that one is following a rule is not sufficient for following a rule. The most important aspect of grasping a rule for a term is in one's ability to use the term. I believe that Wittgenstein's main point in saying that "it is not possible to obey a rule 'privately'" (§202), is to contrast private thought against repeated action, which happens to be publicly observable. Grasping a rule of language requires know-how. This is in direct contrast to Peacocke's interpretation of Wittgenstein. Peacocke writes (1981b, p. 72): "The conclusion of Wittgenstein's arguments about following a rule is, in his own words, that "'obeying a rule" is a practice' (§202). By 'practice' here he means the practice of a community." Peacocke cites a line from Wittgenstein's *Remarks on the Foundations of Mathematics* to support his interpretation. Rather, I believe that by 'practice', Wittgenstein does not have in mind, necessarily, the practice of a community. For my understanding of Wittgenstein, the most important line in the text is in §201: "there is a way of grasping a rule which is *not an interpretation*, but which is exhibited in what we call 'obeying the rule' and 'going against it' in actual cases." This is not specifically about public agreement, or even correctness or

incorrectness, but rather, about what it is to grasp a rule – to take an action in a real situation. Questions of normativity, on my reading, become secondary (if they matter at all to Wittgenstein). Wittgenstein seems to take for granted that there must be rules of language.

Still, I believe that the challenge posed by KW is a formidable one. It is difficult to see how there can be normativity in language if the skeptical part of KW's program is successful. If Blackburn is correct that agreement within a community is not sufficient to solve the problem faced by the individual – then we will be in need of another solution to the rule-following paradox. How should one interpret Wittgenstein's notion of *custom*, which he repeats in several places, most importantly §198: "a person goes by a sign-post only in so far as there exists a regular use of sign-posts, a custom"? Are *custom* and *regular use* normative notions? If so, they will likely be in as much need of explanation as *meaning*. If not, how could they provide a solution to the paradox? If a dispositional account succeeds, then that may be able to provide backing for what it is for a use to be regular. Failing that, it is difficult to see how there can be normativity in language.

Given that my argument from the previous chapter concluded that ordinary language predicates fail to apply, it is worth examining how the strategy I have employed using Sorites relates to KW's use of the rule-following paradox. The primary difference between my arguments and KW's is that I have been focusing on best explanations that one, *sub specie philosophiae*, could give for linguistic behavior, and KW focuses on an individual's own assessment of the rules of a language. There are two aspects of the difference. The less important one, if Blackburn is correct about community language being in as much trouble as private language, is the step from one user's examination of

her own terms to the examination of terms done *sub specie philosophiae*, which attempts to examine the use of multiple users. The more important step is in going from an analysis of the rules of language to a broader causal analysis of why we use terms.

I attempt this broader approach because I do not take for granted that there are rules of language, and in this I agree with KW. Because of vagueness, it is more fruitful to analyze the possibility of rules of language not by focusing on one's introspective judgments about what it is to follow a rule, but by focusing on the actual origins of linguistic practices, and determining whether the best explanation of our practices indicates that there are rules that we follow. Because of vagueness, I answer in the negative. One way to phrase my argument is as a more general problem; I consider it to be a "new riddle of rule following". I am motivated by the consideration, stemming from problems of vagueness, that even given a full understanding of all the relevant facts of our language use, we still do not know what the meaning of the term 'lavender' is. The old rule-following considerations show that given introspective judgments, we cannot say what the rule is. The new riddle is that even if we add knowledge of facts about practices, we still cannot say what the rule is.

It might be helpful to discuss a fanciful example. Let's imagine a community of individuals⁵ who are *vague-sighters*. (I intend them to be philosophical relatives of blind-sighters – those who claim to have no conscious visual experiences but are still able to give accurate answers to questions requiring vision.) These individuals believe that they cannot say where the precise boundary is between lavender and not-lavender when asked.

⁵ Given that I accept Blackburn's individualism, the story can be told equally well with one vague-sighter who repeats a test.

They would rather not answer the question. But if we force them to give an answer, it turns out that they have a remarkable consistency in their responses. In fact, unbeknownst to them, all their responses when forced always draw a precise boundary located between, say, candies 421 and 422.⁶ This seems to be a case where a rule is being followed, but the followers do not grasp the rule. Let's also assume that one can give a naturalistic explanation of this behavior – perhaps, that being able to make such discrimination was an important part of the evolutionary history of the individuals in the community.

According to KW, these people are correct because their use agrees with each other. According to a dispositional account, the meaning of 'lavender' is determined by our dispositions to use terms, and in this case, there are precise dispositions. This kind of example seems consistent with the paradox of rule-following as KW phrases it, but is such that the notion of a rule is not threatened. There must be a rule of behavior that the language users themselves use. LW's point can be made quite well with the example of the vague-sighters: we cannot state the meaning of our terms, but our terms do have meaning. The so-called paradox, given the details as I have given them, seems resolvable either by KW's methods or by a dispositional account. On the former, our words have meaning because there is actual agreement. On the latter, our words have meaning because we have dispositions to use them. The rule-following considerations do not seem to threaten normativity in this logically possible example.

⁶ The fact that usage makes a division between 421 and 422 does not prove that there is a precise boundary given by use; perhaps there is disagreement between individuals where exactly, in the spectral range between 421 and 422, the boundary is. I shall assume that, as far as is possible to test, vague-sighters do indeed demarcate a precise boundary.

There are two relevant questions now: is there any reason to suppose that our actual practices are consistent in some way that we cannot, as users, state, and second, if our practices are consistent in this way, whether they are normative, such that we *ought* to continue the practice. The first question is an empirical one that is likely false: ordinary humans are not vague-sighters. Whether ordinary humans have a more complex suite of dispositions that determine some precise (though not bivalent) model that we cannot ourselves understand by introspection is a question that I have not given a definite answer to, though problems of vagueness make this doubtful.

Let's assume, though, that some dispositional account can be given. For sake of simplicity, let's assume that we are indeed vague-sighters. The question remains as to whether it is a normative fact that, given precise usage, the truth conditions for 'lavender' are precise in just the way that vague-sighters draw a boundary. What is at issue is whether we can take facts about use and draw a normative conclusion from the facts. Surely, it would be absurd to say otherwise, though it strikes me as no more absurd than claiming that one might use 'plus' to denote quaddition rather than addition. Of course, that possibility is worth taking seriously. This is the new riddle of norm following: Let's grant that there are consistent practices and customs, common dispositions, and agreement. Can we then conclude that there is now a normative fact of the matter? Is there a rule that dictates that we must follow our rules of linguistic behavior? Consider a rule of law that claims that citizens must obey the rules of law. That rule would be useless: either citizens will have a practice of obeying the laws, or not. Either way, the new rule will not add anything.

One possible response would be to say that meaning *just is* use. The meaning a word gets is determined by its use, and its use just is this way. Hence it is true that the precise boundary between lavender and non-lavender is at 421. This response relies on an anti-realism that seems to be mistaken. What if, for some reason, everyone has a consistent disposition to err on a particular math problem, and are not able to revise it. Would that make math wrong? No; it would mean that there is a conflict between our dispositions and our acceptance of Peano's axioms. There will be an inconsistency, a real fact of the matter than our dispositions go against application of addition under Peano arithmetic. This is more difficult than I can do justice to, and does relate to problems in philosophy of math, but I hope that it is enough to cast doubt upon the claim that even if there is universal agreement, there might still be universal error. This is a realist assumption, and I am unsure whether I want to rest upon it.

What is at issue is whether we can take our best explanations of linguistic behavior, and extend these explanations into prescriptive claims. What could ground the claim that we ought to continue using language as we have, or that those in the past whose usage differed from most others ought to have conformed? There does seem to be a systematic reason why there cannot be a full justification for that: ought statements cannot be deduced from is statements. What can one say to another who, knowing that everyone has called 421 the precise boundary of lavender, calls it at 423? On the one hand, because language is used to communicate, there is some compulsion to use language as others do. But this does not cover every aspect of our use of a term. It doesn't seem absurd to reject the way everyone else does it. There seems to be something about our beliefs about the linguistic *ought* that we do not just accept 'because everyone else

does it' as enough grounds for claiming that I ought to do it, too. The difference between Quine's inscrutability arguments and KW's is that KW adds facts about our internal states, and there still is no easy way to see how normativity arises. The difference between KW's arguments and mine is that I add considerations about actual use and dispositions, and there still is no easy way to see how normativity arises. That is the consequence of what I call the 'new riddle of norm following'.

4.2 The Is/Ought Gap in Semantics

Though deflationists about reference and truth may claim otherwise, an implicit presupposition of the Sorites paradox is that semantics is normative. For a deflationist, Sorites paradoxes must not be primarily about language, or about the referential structure of the word 'heap' or 'lavender'; they concern heaps and the color lavender. Leaving no space between one's semantic analysis of the term 'heap' and one's analysis of heaps will lead very quickly into an epistemic view. I agree with Sorensen, who writes: "If deflationism is correct, then epistemicism is hard to avoid" (2001, p. 165). The arguments in the previous chapter against the epistemic view may be taken to cast doubt upon deflationism: the best way to resolve Sorites is to analyze the semantics underlying terms such as 'heap', and not take it for granted that 'heap' refers to heaps. Deflationism will always have trouble when there are problems in the first-order discourse, and given that most ordinary language terms are vague, deflationism about reference will need some way to account for it. I am unsure if deflationism is salvageable, but the epistemic theory is not.

My own approach in the previous chapters does set one criterion for what must be the case in order for a term to refer: the term, if it is a general term like ‘heap’, must refer to a precise set of objects; if it is a proper name like ‘Mt. Baldy’, it must refer to one single precise object.⁷ This requirement is extremely difficult to meet, and it seems that few or none of our ordinary language terms meet it. Hence my argument has been for a kind of eliminativism. My argument has presupposed that there is no way to go from facts such as about wavelength and about actual use of terms (and dispositions to use terms), to provide a single model for the reference of terms like ‘lavender’. It also presumes that there is no precise set of admissible models, and so a supervaluationist approach will not work either. It appears that there is no way to go from these non-normative facts about the world and about behavior to normative principles about how terms ought to be used.

Of course, explanations using the notion of *wavelength* might be problematic, because likely, a Sorites can be given using ‘wavelength’ (though I am unsure exactly how). In the previous chapter, with an analogy to arguments against the existence of God, I argue that this is no reason to go back and accept that the use of the terms being explained, such as ‘lavender’. This is the precise point at which I reject Williamson’s account: he despairs of the eliminativist move, claiming that it leads to a scientific reductionism. But I claim that just as one can reject ordinary beliefs about God and not be

⁷ To review: The moderate indeterminist wants both boundarylessness and classification. In order for there to be classification, there must be some kind of barrier between the positive and negative cases. The moderate indeterminist explains the barrier by appeal to some kind of imprecise boundary. But because of higher-order vagueness, there is no way to make the notion of an imprecise boundary coherent – it implies that there is a range of cases which are not determinate. To claim that there is a range implies that there are boundaries to that range. If one claims that there are imprecise boundaries to that range, then there must be some other range of where that boundary may be located. But this argument is not infinitely iterable, given that Sorites series are finite and discrete.

a reductionist, one can also reject ordinary beliefs about heaps. I have left it as an open question what to do about the terms in the lower-level explanations, and I will not fully answer that until Chapter Five. In order to answer it, we must look at structural features of the project of semantics.

If semantics were not normative, then, given the failure of deflationism, we would be able to resolve Sorites very easily with a semantic approach: Explain away the minor premise, that Candy #1 is lavender, by claiming that though we normally (in a frequency or dispositional sense, not a normative one) use 'lavender' to refer to the color of Candy #1, it is not the case that Candy #1 *is* lavender.

The problem for semantics comes from the combination of the assumption that semantics must be naturalizable with the assumption that there are normative semantic facts. But on this assumption, we are led into the epistemic view, which fails for reasons that naturalists should be comfortable with (that according to the best explanation of our linguistic practices, we do not refer with our terms). Perhaps we can still be naturalists, but we cannot both be naturalists and semanticists (of ordinary language), given that the best explanation for our terms is that they do not refer. The further problem for naturalized semantics is that a supervenience approach seems to fail: to say that Mt. Baldy supervenes upon some portion of the landmass in the San Gabriel Mountains will lead also into an epistemic view. Which precise portion is it? There is nothing about the totality of beliefs and practices of the individual and facts about the land itself that determines the precise reference of 'Mt. Baldy'.

The general problem for normative semantics is that it is an attempt to get an ought from an is. This should seem highly familiar to those dealing with normative ethics, for essentially the same problem occurs there. Here's what Hume has to say:

I cannot forbear adding to these reasonings an observation, which may, perhaps, be found of some importance. In every system of morality, which I have hitherto met with, I have always remark'd that the author proceeds for some time in the ordinary way of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when of a sudden I am surpriz'd to find, that instead of the usual copulations of propositions, *is*, and *is not*, I meet with no proposition that is not connected with an *ought*, or an *ought not*. This change is imperceptible; but is however, of the last consequence. For as this *ought* or *ought not* expresses some new relation or affirmation, 'tis necessary that it shou'd be observ'd and explain'd; and at the same time that a reason should be given, for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it. But as authors do not commonly use this precaution, I shall presume to recommend it to the readers; and am persuaded that this small attention would subvert all the vulgar systems of morality, and let us see that the distinction of vice and virtue is not founded merely on the relations of objects, nor is perceiv'd by reason. (*Treatise*, Book III, Part I, Section I)

Here's an attempt to state the problem for semantics:

Every naturalized system of semantics relies on a presumption that we begin with facts about the external world, and facts about our the history of our language use, our actual language use, and our dispositions to use language. And then, they go from these kinds of claims about regular patterns to claims about words referring to items in the world, where this claim of reference is taken to be normative: one *ought* use the term in such and such a way. But this *ought* expresses some new relation, and some explanation is needed for how it can be a deduction from others. I am persuaded that this small attention would subvert all the philosophical systems of semantics, and that the distinction between whether a term refers to an object or not is not founded merely on the relation of objects, nor is perceived by reason.

Famously, Hume proceeds to engage in normative ethics after his warning about the vulgar systems of morality. I think it is right not to take the is/ought gap seriously unless there is a first-order reason why the relevant normative practice has problems. For

example, concerns about widespread moral differences, which Hume himself does not sufficiently countenance, have led to debates in contemporary meta-ethics about the is/ought gap and the foundations of normativity. I believe that the Sorites paradox is just such a problem for semantics. Sorites paradoxes show that our ‘semantic sense’ is insufficient in giving us an adequate picture of normative semantics in the way that Hume believed (naïvely, perhaps) that appealing to our own moral sense could give an adequate picture of ethics.

I’m unsure how deeply the parallels go, and if the is/ought gap in semantics is simply the same problem as the is/ought gap in ethics, or if it is merely an analogy. There do seem to be two different entrées into the normative – one for how one should act, and a separate one for what one should believe. Pragmatists like Rorty and recent Putnam might claim that it is one problem. At the very least, what naturalized normative semantics and naturalized normative ethics have in common is that they seem to commit a logical flaw: deriving an ought from an is.

There are many alternative meta-ethical accounts that are designed to avoid this problem. Perhaps ethics is not an effort to derive an ought from an is; we simply start with acceptable ‘ought’ claims and then draw further conclusions. Is the same kind of view available to the naturalized semanticist? The problem is that we may still be inclined, in the face arguments given by the moral skeptic, to continue our practice of making moral claims. The moral skeptic might not undermine our confidence in our moral judgments. But for semantic oughts, when faced with Sorites paradoxes and the proposed eliminativist resolution of them, though we might still continue our practices of

using language, we may do so believing that we aren't referring with our terms. It strikes me as being easier to give up on semantics than to give up on ethics.

4.3 Millikan's Teleo-Semantics

Ruth Millikan incorporates evolutionary considerations in giving a realist account of meaning. She sets up her account as a solution to the *normativity problem*:

Each of these [alternative] theories of representation, if taken bare, runs into exactly the same problem, namely the problem of accounting for cognitive errors: misperceptions, false beliefs, confused concepts, bad inferences, unrealized intentions, and so forth. Each founders over the distinction between the facts of cognition and the norms of cognition. Call this the normativity problem. (1993, p. 3. All Millikan citations in this section are from her 1993.)

There seem to be two separate problems that Millikan points to in the above passage: the first is the problem accounting for errors in cognition; the second is accounting for the norms of cognition. As we shall see, Millikan believes that by giving an evolutionary account of errors in cognition she has thereby solved the problem of accounting for normativity. Though I believe that she is reasonably successful in the former project, she does not do enough to show how there might be normativity. One can give an account of errors without supplying an account of truths.

Millikan claims that evolution designs us with certain features, for example, a heart, as having functions. Likewise, our brains have functions:

Might not evolution have designed our brains to produce representations, in accordance with certain highly abstract rules of projection, by certain methods of physical interaction with what is represented? Could it be that evolutionary design is what hooks a percept or a thought onto the state of affairs it represents? This is the point I will argue. (6)

Millikan seems to view 'evolutionary design' as bestowing normativity.

One reason that I choose to discuss Millikan is that she has an explicit analysis of KW's rule-following considerations, which she takes to be a challenge to semantic naturalism, as I do. Her analysis of it is also a helpful way to bring her full view into light. She summarizes what she takes to be the problem with giving an account of rule-following: "The problem is to account for the *normative* element that is involved when one means to follow a rule, to account for there being a *standard* from which the facts, or one's dispositions, can diverge." (p. 216)

Millikan's example in the analysis is of the male hoverfly, which hovers in place awaiting a female hoverfly to pass by. The male hoverfly darts after just about any object that is approximately the same size and speed of a female hoverfly – such as projectile dried peas, male hoverflies, blocks of wood. The male hoverfly is better off (reproductively) if, in the interest of a hasty chase, it often pursues objects that are not female hoverflies, than if it were to take more time to evaluate whether the object in its sights is a female. Millikan distinguishes between the *proximal hoverfly rule*, which is of the form: "If there is an object of size X flying across your visual field at velocity Y, then go get it!", and the *distal hoverfly rule*, which is of the form: "If you see a female hoverfly, go get it!". The male hoverfly follows the proximal rule in order to follow the distal rule.

Millikan uses this to resolve the rule-following paradox by discussing the following example: Suppose that no male hoverfly had ever been approached by a female hoverfly at a certain range of velocity $[V-\epsilon, V+\epsilon]$, which is a sub-range within the normal range of female hoverfly velocities $[W-n, W+n]$. Could the proximal hoverfly rule be: "if you see something moving between $[W-n, W+n]$, then go get it, except if it's between

[V-ε, V+ε]”? Millikan appeals to claims about simplicity of explanation to argue that the rule really just is “if you see something moving between [W-n, W+n], then go get it”. Millikan quickly extends considerations of hoverflies to considerations of humans. She argues that for cases which we haven’t confronted, like 68+57, simplicity of scientific explanation indicates that we do have a purpose in answering 125, and hence (given Millikan’s identification of biological purposes with language meaning, to be discussed further below), we do mean addition and not quaddition with ‘+’. I find this resolution problematic, because it then brings up Goodman’s famous grue problem. Why is it that scientific explanations that are not disjunctive are better explanations than those that are? Millikan claims (221) that her goal is to resolve the rule-following considerations and not Goodman’s paradox; but both Kripke (1982: 58, 62fn) and Blackburn (1993, passim) claim that the two are intricately connected.

Putting this point aside, Millikan uses the distinction between proximal and distal rules to do important philosophical work. She claims

These truth rules would concern distal conditions under which we should say or think certain things. The truth rules might imply directives with this sort of form: if you have reason to speak (think) about the weather in Atlanta, say (think) “It is snowing in Atlanta” when and only when it is snowing in Atlanta; if you have reason to speak (think) about the color of snow, say (think) “Snow is white” if and only if snow is white. For a simple biological model here, compare worker honeybees. They (biologically) purpose to follow rules of this kind: when dancing, angle the axis of your dance 10 degrees off the vertical if and only if there is a good supply of nectar 10 degrees of a direct line from hive to sun. (233)

She uses this consideration to argue against the ‘verificationist’ anti-realism of Putnam and Dummett. While it is the case that proximal assertibility rules (such as “say ‘P’ iff there is good evidence for P”) concern what is within our own evidential scope, proximal

truth rules are followed only because we have a distal biological purpose, which is to tell the truth.

I think that this is an interesting approach to Dummettian anti-realism – it gives grounds to the claim that truth conditions may in some cases go beyond what can be given in assertibility conditions. I make a very distantly-related argument against ‘verificationist’ anti-realism in Chapter Three (one that makes no reference to biology). Millikan’s view can be used to refute the dispositional theory. Blackburn claims that we may have dispositions to correct errors in our initial dispositions. But the best explanation of our dispositions is that they were given to us by evolution, and we really are disposed to make some errors, just like the hoverfly is. It seems implausible to suppose that we have dispositions to correct *all* our errors, given the very pragmatic nature of our dispositions.

However, I’d like to make several points against Millikan’s view. The first is that Millikan does not note the difference between saying and thinking. I’ll grant here that we have a distal biological purpose to have truthful inner representations of the world, and I’ll even grant that biological purposes determine language meaning. Do we have distal biological purposes in speaking the truth? Certainly, it would not suit Grog well for his evolutionary purposes to lie all the time, because then no one would ever listen to Grog. But does Grog have a purpose in saying “it’s snowing in Atlanta” in all and only those cases in which he has reason to speak of Atlanta’s whether and in which it is snowing there? Certainly not. Humans are more sophisticated than bees, and this is one of the relevant ways in which we are different – it may suit Grog well, for a number of reasons, to be deceptive about Atlanta’s weather. Even if we assume that humans today are acting

upon distal evolutionarily-given purposes, which itself is highly controversial, there is no evidence that our most distal purpose in speaking is telling the truth. We speak to get ourselves ahead, and often this involves us in telling the truth, but does not always require this. If Millikan needs to show that language meaning is given by biological purpose, then to show that language can be characterized by correspondence truth conditions, Millikan will need to show that the ultimate distal goal in speaking is in telling the truth – that the biological purpose of speaking is truth-telling. However, it seems that a proper biological account of human language use will give a very prominent role to the distal goal of proper socialization, and other aspects of language use which are not necessarily connected with one's telling the truth.⁸

For what follows, let's grant that one's single biological purpose in speaking were in telling the truth. The critique of Millikan that I have just developed is something of an internal criticism: it shows that it is not the case that our biological purposes are in speaking the truth. But there is another problem: The world is so complex that it is impossible to regulate precisely how we should speak of it. Here's a speculative hypothesis: The very reason a language capacity arose is that our predecessors were unable, based in casual observation, to grasp the world as it really is. Rules of language are what permit us to speak at all, given that we don't represent the world precisely. The very notion of a true statement comes out of the fact that we're never getting the world precisely correct, and so we need something to condone some incorrectness.

⁸ I consider the foregoing argument good reason to buck the current naturalistic trend in subsuming philosophy of language into philosophy of mind. If our purposes in having a language capacity are separate from our capacity to garner veridical mental states, then there is much less motive to believe that there is a *language* of thought.

More generally, my concern with Millikan's view is that even *stating* the error problem requires an assumption that there is normativity. Other naturalized semantic accounts emphasize *causation*. But the classic problem for these theories is that often a non-cow (e.g., a horse) causes tokenings of 'cow'. Millikan's approach resolves this by making a distinction between proximal and distal biological purposes. As naturalized semantic views go, I believe it is an ingenious way to get around the problem. But there is a more basic normativity problem: how do we know that 'cow' refers to cows, rather than nothing at all? Just because Millikan has provided an elegant theory which, if we identify concept meaning (i.e. truth conditions) with distal biological purpose, shows why we refer just to cows with our term 'cow' rather than cows and an occasional horse. But why must we identify meaning with distal biological purpose?

There are other considerations that relate more directly to issues of vagueness. One problem for Millikan's account is that *perfect* representation is not needed for the survival of a species, and this will make vagueness quite a problem for her account. Millikan writes: "Forming adequate concepts is learning to represent or map in thought what is the same as the same" (11). But vagueness makes it difficult to claim that what we are doing with most of our terms is mapping a real pattern of the world into thought. Rather, in perception, we do not see the world in its actual precise state, and we create vague language so as to do the best we can, given our perceptual limitations. On my alternative account, which I discuss briefly in Chapter One, no external *objects* are needed to account for evolutionary success. An explanation of evolutionary success that makes use of the notion of a continuum of particles seems highly plausible, and seems to me the better analysis. The problem for Millikan is that we can't just assume that the

items that fall under a general term are really the same. Instead, it seems most plausible that what language does is map what is *similar*, though not the same, as the same. But this involves an *error*, given that our sentences are vague and the world is precise. This error is why I argue for eliminativism in the previous chapters.

I have two other criticisms of Millikan's account. Her selection of the male hoverfly as her primary analogy glosses over a potential problem concerning indeterminacy. For the purpose of the male hoverfly's darting mechanism has only one object of its distal goal: hooking up with a female. But there are other more complicated examples: Let's imagine a kind of frog that catches two kinds of flies, which both are about the same shape and size as each other, but are biologically very different from each other. What's the biological story of what the frog represents? There is more than one thing that can satisfy the frog's needs. The world itself doesn't provide a single demarcation of what the frog is built to pick out.

This criticism is in fact a more general form of a common argument against biological accounts of semantics. Loewer makes an argument relating to this point:

Truth-conditional content seems much more determinate and fine-grained than anything that teleology is capable of delivering. This is made obvious by considering that there cannot be any selectional advantage for creatures whose beliefs are about rabbits over those whose beliefs are about undetached rabbit parts; yet our contents are so fine-grained as to distinguish these belief-states. (1997, p. 117)

I don't think there would be problems for naturalized semantics if our best explanations state that our biological purpose in speaking, and what we still do when speaking, is to represent the same things as the same. If it could, perhaps it would give all that is needed for there to be normativity. If it were the case that there are objective similarities in the

world that our perceptual and linguistic mechanism were designed by evolution to represent, then that might provide some grounds for claiming that there is a normative fact of the matter. But, if Quine (1960) and my arguments from vagueness are correct, it does not.

We can have a non-evolutionary interest-relative view, like Graff's. The evolutionary theorist, armed with Millikan's distinction, might try to recast these proximal interests into more distal evolutionary interests. Notoriously, they haven't done a good job to this point. If distal biological rules served as the best explanation of our behavior, that'd be something. That's where most evolutionary accounts have problems – that they haven't, to this point, given great explanations of human behavior. Maybe it's no coincidence that the species they have the hardest time with is also the one with (supposedly) propositional discourse.

But let's assume that the best explanation of our linguistic behavior can be given by showing how we have certain distal biological purposes. Ought we continue to act in accord with these purposes? There does seem to be an is/ought gap here. On evolutionary accounts, certain things are designed by nature to have functions, naturally. The best explanation of the heart is that it has the function of pumping blood. This function is to be construed normatively: hearts *malfunction* when they don't pump blood. In other words, hearts that do not pump blood fail to do what they ought to do. Extending this to our semantic systems, it is said that we malfunction when we use language out of line with both our proximal and distal biological purposes.

I'm not going to give a full discussion of this here. It is true that we *say* those things about hearts when they do not pump blood, and it is true that it helps us to

understand an item when we are told what its purpose is. But at the same time, it does seem as if all biological purposes can be rephrased in terms of hypothetical claims: “if one wants to pass down one’s genes, then one ought to ...” The problem is that there is no necessary reason for accepting the clause in the antecedent. How, then, is biological purpose really normative? On theological theories, God bestows normativity. But can ‘evolutionary design’ really play the role that ‘God’s design’ plays?

There is something like a Euthyphro problem here. Is a behavior *correct* because it was designed by evolution, or was a behavior designed by evolution because it is *correct*? The answer seems to be that those behaviors which have survived to this day have been carefully selected for because they were behaviors that tended to lead to successful reproduction. Hence, the second horn of the dilemma seems more right. But here, *correct* would just mean ‘conducive to reproductive success’. But there are several problems if the evolutionary theorist resolves the dilemma that way: First, there may be other behaviors that are also conducive to reproductive success, and it isn’t at all obvious that our normal use of ‘correct’ seems in line with the evolutionary theorist’s. Millikan herself wants an historical notion of success, so even if a behavior is *now* conducive to reproductive success, it is not our function if it wasn’t a behavior that was designed by evolution. But why have any notion of normative *correctness* if we’ve given a full explanation of our behavior using evolutionary considerations? If Millikan’s evolutionary science does all that she wants it to do, it would have the effect of removing the notion of normativity, rather than explaining it.

If evolutionary theory could prove that there are macro-sized objects, and that we were designed by evolution to represent them, then that might help the evolutionary

semanticist. As I noted above, Millikan's theory just assumes that there are objects. But if could prove that there are objects, then it seems that our biological purposes are irrelevant – we should just use language in such a way as to refer to the objects, regardless of our biological history.

4.4 Normativity and Vagueness

This brings up an important point about the is/ought gap: If our investigations leave us, *sub specie philosophiae*, with 'nothing else to think'⁹ but that we ought to call a certain item 'red', then even though we cannot deduce from premises that it ought to be called 'red', there is no point denying it. In such cases, the is/ought gap becomes *irrelevant*. Unfortunately, there are many reasons why evolutionary semantics does not leave us with 'nothing else to think'. I have not, in the preceding, analyzed other naturalized semantic theories, but (on my understanding of the literature), none seems to have given a complete enough analysis to make the is/ought gap irrelevant. Vagueness is one, though not the only, reason why. The problem of vagueness shows that for a semantic analysis to be correct, it must be precise.¹⁰ But it is implausible to suppose that any naturalized analysis of human behavior will assign the requisite precision to our linguistic behavior.

⁹ I take this expression from Wiggins, 1996.

¹⁰ As a reminder, this is forced by higher-order vagueness. An account which says that there is a range in which humans are unable to make distinctions is then obliged to give precise boundaries for that range, and this seems impossible. Consideration of context may be helpful, but even if one claims that there is no general answer to where precise boundaries are, one is still committed to giving a precise answer

within each context, and this seems impossible as well. On what basis can we create a precise semantic model?

What is the use of appealing to an is/ought gap in semantics when problems for naturalized semantics can be stated without reference to one? If it turned out that we all are vague-sighters, and that there were no other problems for naturalized semantics, would it matter that one cannot deduce semantic oughts from claims about linguistic behavior? My consideration of the is/ought gap was motivated by the thought that even if we knew all the factual information about our linguistic history and our linguistic dispositions, we still could not give precise models for the reference of our terms. As I note above, a similar problem faces ethical theory. What makes the is/ought gap important for ethical theory is that there is in fact moral disagreement. Were there none, then there would be no reason for the is/ought gap to cause any commotion. But the is/ought gap along with the fact of moral disagreement give strong support to the claim that there are no moral truths. Thus Wiggins (1996) (as well as Smith 1994) aims to specify conditions under which there would be no moral disagreements. Wiggins doesn't see this as a refutation of the is/ought gap, but as a means to sidestep it.

Here's what I consider to be a plausible ethical theory, which I argue for in greater detail elsewhere.¹¹ Let's assume what some moral skeptics argue for: there is no fact of the matter about any external ethical truth. We may have 'moral sense' that makes us make assertions concerning like "that's wrong", but the word 'wrong' does not refer to any property. We can give a full explanation of our use of moral terms without needing to assume that they refer to anything real. We can only have primary knowledge of factual claims, and one cannot derive a moral claim from a factual claim, so there are no

¹¹ See my "How to Be an Ethical Quasi-Realist" (unpublished).

knowable moral truths. (And it is hard to understand what an unknowable moral truth could be.) In honor of John Mackie, let me call this the ‘post-Mackian’ position, or the PMP. Let’s assume that I have reached the PMP myself. I then wonder: What should I do next? This question should be taken literally: Now that I am done with the task of meta-ethics, should I move to the shores of a Norwegian fjord? Go shopping for a new umbrella? Stay at home for the remainder of my days? How should I go about answering these practical questions, given that I will be constantly faced with them?

What I hope this shows is that as persons, we are continually faced with decisions about what to do. And even an acceptance of moral skepticism seems to be consistent with facing a decision about what to do next. The “should” in the question “what should I do next?” cannot be anything other than a *normative* notion. Given that there is no ethical truth outside the context of this post-Mackian decision procedure, we may consider ethical truth to denote the best choices to make from this standpoint. Every person who ever asks ‘What should I do?’ (or, perhaps, “How should I live?”), which is everyone, must go through some decision procedure.

How should one in the PMP go about deciding what to do? Presumably, she will survey her desires. She may prefer certain outcomes over others. Her desires will presumably already include other-regarding sentiments, so the result will likely not be ethical egoism. She might also seek advice from others, given that others may have insights that she would appreciate about how she should act. It will also be helpful for her to establish standards for ethical decision making, given that she will repeatedly face the decision of what to do next. Furthermore, she will have to continually make choices with a limited amount of time and information. Though the decision and standards will take

into consideration the views of others, the judgments will ultimately be her own; that is the best one could expect her to do. So the task of ethics, under this new, metaphysically slim foundation, is to solve practical questions—my ethical deliberation simply is deliberation about what I judge to be the best standards are that indicate what the best thing is for me to do, or the best way for me to live.

I hesitate to give a quick schematic because I am concerned that the schematic, without much further hedging, defense, and explanation, will appear both philosophically implausible and practically disastrous. But with that proviso, let me state the beginning of a framework from which I take it to be legitimate to derive an ‘ought’ from an ‘is’. The most basic schematic of my view is that one can derive an ‘ought’ from an ‘is’ in the following way:

- (1) {a, b, ..., n} are the options available to me.
- (2) {a', b', ..., n'} are the expected outcomes of my choosing {a, b, ..., n}, respectively.
- (3) of {a', b', ..., n'}, I judge i' to be the *best* outcome.¹² (Where ‘best’ relates to some facts about my preferences – something akin to ‘my most favored’ – and is not irreducibly normative.)
- (4) Therefore, I ought to do i.

I believe that this kind of inference, properly expanded upon, is valid. The notion of ‘ought’ contains in it a measure of the best option from among the options under consideration. What is important for my purposes is that both governance and guidance will be satisfied. As for guidance, the decisions that we make will indeed be guided by our sentiments and guided by considerations given to us by others. Governance will be

¹² I don’t intend this to assume an overly naïve form of consequentialism, where every action is to be judged on the basis of some *further* consequence. For example, I wish to permit something like the following: the expected (relevant) outcome of my deciding to go dancing is that I will be dancing.

satisfied because ethical decisions can themselves be justified by appealing to the fact that I should do *something*. If one is questioned about why one has done a particular thing rather than something else, then a first-order debate will ensue. Elsewhere, I argue that these kind of first-order debates can be fruitful, even if one believes that one is starting from the PMP. However, I believe that it is not always the case that moral argument will provide a single answer of the right thing to do, and in these cases, one must make a *choice*. But choices are made, most commonly, by having already excluded some potential options. Hence I believe that this account will not lead to a pernicious kind of relativism (by which one cannot condemn Nazism), though it does not claim that there is a single correct way to respond to moral questions.

Though my view seems to be similar to an existentialist account, it is not the same. Though the PMP is one in which there is no single answer that would be agreed upon by every rational individual, and so each individual must make a *choice*, it is not a choice that is completely free (contrary to Sartre). To claim that normativity will ultimately be dependent upon choice does not commit one to the existentialist account of free will. What it comes to is that we are directed, in most cases, in a certain way. The view is similar to Hume's view that all have moral sentiments that favor the public good; however, it need not rely on the degree to which Hume believed there would be moral agreement – it is not quite as optimistic as his account. Though I do not have a more worked-out account of the issues of free-will relating to these points, I do believe that it is a plausible, middle-ground position.

How might these considerations relate to issues of normativity that arise in the context of vagueness? Here's the rough picture: We can choose what the terms refer to.

We have some dispositions, but not dispositions across the full spectrum of cases. Furthermore, even in cases where our dispositions do tend to guide us in one direction, we are always in principle able to override these choices. Alternately, if there *were* complete dispositions across the spectrum, then the is/ought problem wouldn't be a problem because we'd be able give the whole semantic theory in terms of the dispositions, and we would have immediate responses to where boundaries lie in Sorites paradoxes. But we don't. So my middle ground account is to say that while normativity is based in choice, our choices do tend to have a lot of convergence. I will say more about the notion of choice in Chapter Five. On my theory, we are guided to some extent by how we are, and we are governed by the facts about having to making choices. In the next chapter, I will discuss further how this route relates to the eliminativist line, and why we should take this route rather than eliminativism. I will discuss further the notion of choice, and how it relates to the dispositions we already have. Given that choice is implicated at many stages in our semantics, I will claim that normativity is dependent upon choice. But that does not entail that it is totally free.

Man gleaning food between the solemn
presences of land and ocean,

On shores where better men have
shipwrecked, under fog and among
flowers,

Equals the mountains in his past and
future; that glow from the earth was only

A trick of nature's, one must forgive
nature a thousand graceful subtleties.

- Robinson Jeffers, "Point Joe"

CHAPTER 5. THE NORMATIVE CHOICE SOLUTION

In this concluding chapter, I'd like to review what I take to be the most important arguments in the preceding chapters which lead me to what I call the "Normative Choice" solution to Sorites, discuss in outline how I see the solution working, and then tie up some loose ends raised by the theory.

In Chapter One, I discuss some of the background features underlying Sorites paradoxes, and I argue that:

1. Bivalence holds for sentences that express propositions.
2. Vagueness arises from indeterminacy, and is not primarily a matter of slippery slopes.
3. Moderate indeterminism – views that claim that Sorites series are indeterminate for some portion of the series but not for all – fails because of higher-order vagueness.

In Chapter Two, I continue the argument against moderate indeterminacy and discuss other views:

1. Degree theory, though it respects the OD→PD intuition¹, fails.
2. Supervaluation, though it respects some intuitions concerning levels of admissibility, also fails.
3. Graff's interest-relative view is promising but fails because of the vagueness of interests.
4. Wright's newest view, that vagueness is due to *quandaries*, does little more than beckon us to do more work on vagueness.

In Chapter Three, I argue that the epistemic view fails:

1. Neither considerations of bivalence nor of margins for error provide sufficient motivation for it.
2. Though it may not be worse than its current competitors, that is no reason to accept it.
3. Consideration of the best explanation of our linguistic practices gives us strong reason to believe that terms such as 'heap' do not refer to anything – despite the fact that even these best explanations themselves may be problematic because of vagueness or other reasons.

In Chapter Four, I argue that radical indeterminacy should not be surprising:

1. Vagueness provides an even stronger argument for indeterminacy than Quine's 'Gavagai' argument.
2. Kripke's Wittgenstein poses good reason to believe that reference is indeterminate, and an extension of that argument supports my views on vagueness.
3. Millikan's naturalized (teleo-) semantics does not resolve the problems of indeterminacy.
4. There is an is/ought gap in naturalized semantics, and vagueness demonstrates that we should take it seriously.
5. There is a plausible solution to the is/ought problem in ethics, which relies upon the notion of choice. Such a notion can, in principle, be applied to semantic normativity.

In this chapter, I'd like to show how such an account might work:

¹ As a reminder, according to the OD→PD intuition, there cannot be an ontological difference between two items (such as two candies in a Sorites series) without there being a corresponding perceptual difference.

1. The normative choice view can quickly resolve Sorites paradoxes.
2. The notion of semantic choice is not as radically subjective as it might seem – we’re not choosing how the world is.
3. Eliminativism itself is not satisfactory as an explanation of our linguistic practices, because we do assume that we have success using language.
4. Loose ends – concerning determinism; metaphysics, given our choosing truths (and positing them); and the status of our best explanations – are tied.

5.1 The Normative Choice Solution

According to the Normative Choice Solution to the Sorites paradox, we may choose a semantic model for our terms. Because the structure of the external world does not determine a precise semantic model for our vague terms, and because our linguistic history and dispositions do not do so, we ourselves may do so in a conscious way. In this section, I shall show how the normative choice solution can resolve Sorites, and in the next, I will discuss what I take to be some objections to it.

5.1.1 The Basic Bivalent Approach

The simplest way that this can resolve Sorites paradoxes is as follows. Let me repeat the Sorites from Chapter One:

- | | | |
|-------|---|-----------------|
| (1) | L_1 | (Minor premise) |
| (2) | $(\forall p) (L_p \rightarrow L_{p+1})$ | (Major premise) |
| ----- | | |
| (3) | L_{1000} | |

(3) is said to follow from (1) and (2) from 999 instances of universal instantiation and modus ponens. But (3) is false, because

- | | |
|-----|-----------------|
| (4) | $\sim L_{1000}$ |
|-----|-----------------|

On one application of the Normative Choice view, we will accept (1) and (4). At a certain point in the series, say, at $p=411$, we will say that L_{411} and $\sim L_{412}$. Hence, ‘lavender’ will have a sharp boundary, and the major premise is false, as the epistemic theorists claim, but it will be a known sharp boundary – it lies exactly where we stipulate it to be.

The hope is that the traditional worry about there being a sharp boundary will not be applicable against the normative choice view. On the traditional worry, captured in what I have called the *boundarylessness* intuition, the facts about the world and our language use do not determine semantic rules for terms like ‘lavender’ that allow for two things that look similar to have different semantic values. On my view, the semantic rules don’t determine that they do. It is *we* who are claiming that ‘lavender’ applies differently to the two cases, and we are justified in claiming that – it is better to have a stipulated boundary than no boundary at all.

On this account, it is not that we are choosing how the world is. The external world is how it is, independent of these semantic choices. But, if the truth values of our sentences do vary with our choices, and our sentences are supposed to be about the world, then doesn’t my view have the absurd consequence that by making these decisions, the world somehow changes? This potential objection can be averted even without abandoning a correspondence theory of truth. For it can be said that what we are doing in making these semantic choices is stipulating which propositions are expressed by our sentences. On one potential propositional interpretation, ‘ L_{411} ’ is true; on another, it is false. In making the semantic choice, we are simply choosing which proposition we are expressing. As such, it does not threaten ordinary intuitions that our choices cannot affect the physical status of external objects (except insofar as our choices themselves are

physical events with certain uncontroversial causal properties). So, one could potentially hold a correspondence theory of the truth of propositions and still hold the normative choice view. This might happen if one were a kind of reductionist. For example, if one held that our current term ‘lavender’ cannot be reduced to any non-vague microphysical facts, but we can stipulate a set of microphysical states to which it applies, then that view will be consistent with a correspondence theory and a normative choice theory. However, as I discuss below, this kind of reductionism is not the view that I hold.

I claim that on this application of the normative choice view, the OD→PD and boundarylessness intuitions will still be satisfied. We’re not making an ontological distinction without a corresponding perceptual difference – we’re making a semantic distinction. Now, of course, we might, in this spirit, just want to say that it *is* an ontological distinction, using ‘ontological distinction’ in a revised way, in a way that’s informed by the arguments given by the eliminativist. But if we do wish to operate in this spirit, we should at the same time give up on the OD→PD intuition.

5.1.2 Degree Theoretic Normative Choice

On the basic approach given in the preceding sub-section, I kept a bivalent logic and I claimed that we can stipulate which proposition our terms express. But why not claim instead that we can stipulate which *logic* applies to vague terms? I don’t have a conclusive argument against this account, and I don’t want to exclude the possibility of such accounts working, but I regard this approach as unmotivated once the initial eliminativist steps have been taken. Reasons to accept a degree-theoretic logic seem to go away once the consideration that the degree-theoretic semantic model is underdetermined

by the facts of the case is taken seriously. Of course there is a slippery slope from short to tall or from red to orange, and these micro-level facts could conceivably map onto degrees of truth. But this in itself is not reason to adopt a degree-theoretic logic. It simply shows that either we should talk in precise terms at the micro level and give up on the macro-level terms, or we should use a bivalent logic that independently allows for comparatives to be made.

We've seen that vagueness is about indeterminacy, and not slippery slopes. Because of this, nothing seems to be gained by introducing the possibility of a statement being 73% true – the normative choice view can allow for anything that one might intend in calling something 73% true, and it doesn't have the awkwardness associated with these degree-theoretic assumptions and the other difficulties facing a non-bivalent logic. In saying that 'x is red' is 73% true, what might one be saying? That x is redder than most things in a range? That we are not quite prepared to call x 'red', but don't want to say that it's not red? It seems that the normative choice view can account for any of these. In terms of the comparative 'redder than', it can still account, by appealing to micro-level properties, of why x is redder than other things in a range despite not having a choice-independent fact of the matter about which ones are red. We can account for the thought that x is in a border-area by first making a stipulative decision of whether to call x red or not, and then saying afterwards that there really was no intrinsic fact of the matter, and that we weren't naturally inclined to call it one way or the other, but we had to make *some* stipulation.

This isn't a knock-down argument against the use of a degree-theoretic logic, and I won't exclude pragmatic considerations from providing some weight in this matter. But

as far as I can tell, they do not provide enough weight to reject classical logic. (I seem to be persuaded at least slightly by the ‘logical snobitivist’ arguments I attribute to Williamson and Sorenson in Chapter Three.) Quine, too, keeps bivalence on pragmatic grounds, though he recognizes the potential costs. I suppose one way of characterizing my argument here is to show that once we’ve tempered our OD→PD intuition with eliminativist arguments, it is unclear whether there is any cost in maintaining bivalence.

5.1.3 Supervaluationist Normative Choice

There is still something remaining in favor of the intuitions that motivated supervaluation: we do have some semantic intuitions that there may be three (or so) levels of vagueness. At the very least, supervaluation seems correct in that there does seem to be something like a range of admissible valuations, and perhaps a range of possible second-order ranges, and perhaps some higher order ranges. Let’s then consider a supervaluationist view in the spirit of Fine in which we stipulate that there is a range at the beginning that is definitelyⁿ red for all n. There is much to be said for this view. We can stipulate a picture looking something like what follows :




Figure 1. The stipulated supervaluationist spectrum

The good thing about this picture is that it could potentially capture our multi-tiered intuitions about Sorites series. My criticism in Chapter Two against standard supervaluationist accounts is that there is nothing about our practices which indicate where these ranges lie, and that it still must assume that there are precise boundaries

somewhere within the range. But the view that we should stipulate supervenience can accept that there are no precise natural boundaries between the ranges; the boundaries here are just stipulated.

However, I don't believe that this account is better than the basic bivalent account. First of all, it is contrary to the motivation of the original supervenience account, which says that vague predicates are in need of precisification within ranges. This sounds similar to my account – that vague predicates indeed are in need of stipulation – but on my account, nothing even determines the ranges by which acceptable precisifications can be made. Hence, the normative choice view accounts for the intuitions behind supervenience, and rids it of the assumption that makes it incapable of dealing with paradoxes of higher-order vagueness.

Still, there does seem something right about a tiered picture. It does seem like the normative choice view will allow one to select *any* boundary for vague predicates even though it seems like we should, under no circumstances, say that Candy 2 is not lavender. So perhaps we should be stipulating such supervenience boundaries. However, I am uncertain how such an account could work. What would then set the boundaries of the inner ranges? If the leftmost boundary, *a*, denotes the last point at which it is admissible to say that it is admissible to say that it is admissible to say that the Candy is not red, then why bother setting further ranges inside? What is the meaning of these inner ranges? Answering these questions may take us deeper into supervenience logic than can be

undertaken here, but I believe that no proper answer can be given on behalf of the supervenientist.²

I should say a few things about the intuitions behind having some ranges. First, the normative choice view is a *revisionist* view, and so it is not surprising for it to encounter some intuitions which it cannot accept. The general idea is that one should not still believe that there are admissibility ranges if one accepts the rest of the story behind the normative choice view. Second, it might be helpful to introduce here an analogy with ethical theory. On some neo-Humean views, such as that of Blackburn, even though morality is founded in sentiment, and moral claims aren't true in virtue of moral terms' correspondence with a real moral object in the external world, there still is a significant place for moral deliberation. Even though we cannot expect to locate a real moral truth as part of the fabric of the universe, there still is much to be gained from moral deliberation. I share this view.

Though this consideration might seem far afield from the discussion at hand, I believe it is applicable. What should we do, *sub specie philosophiae*, when considering where to stipulate a boundary? Intuitions that there are ranges of admissible values will be part of this deliberation. So although I believe it is not worth accounting for these intuitions as part of a 'normative choice' logic, I do believe that they can play a role in deliberation. I shall have more to say about such deliberation below – it still seems odd to claim that there will be any meaningful deliberation about where to place the boundary if one has already accepted eliminativism.

² Carl Posy urged this point upon me years ago when I was tempted by this kind of supervenientist account.

5.1.4 Three Valued Normative Choice

Though I believe it is not helpful to introduce supervenient or degree-theoretic logics, it might seem reasonable to stipulate a three-valued logic. This would maintain our intuition that there is no sharp boundary between red and not red. I believe that this is a highly plausible alternative to the bivalent view I advocate, and I only have one reason for not adopting it. The reason is just the same as before – once we have taken the eliminativist steps, is there still any motivation to reject bivalence? On a standard three-valued view, the middle range is the range where it is not determinately red and not determinately not red. But once we have concluded that the whole range is indeterminate, then what use is there in having a three-valued logic? One response might be that we will continue to make claims like “I wouldn’t call it ‘red’, but I wouldn’t call it ‘not-red’ either”, and a three-valued logic would account for that kind of intuition quite well. I believe that such claims arise from misunderstandings that the eliminativist move clears up, but I don’t have a knock-down argument that we should reject them. Because of this, I have some sympathies with a three-valued view, though in what follows I shall keep the basic bivalent view.

5.2 Problems for Normative Choice

In this section, I’d like to discuss some potential objections to the normative choice solution.

5.2.1 Aberrant Choices and Semantic Deliberation

In the example I have chosen above, I stipulated that the boundary would be located between 410 and 411. But, given my arguments for eliminativism, on what

grounds could one have for preferring the boundary to be set there rather than, say, between 1 and 2? Given that there are no fully rational/cognitive ways to determine a precise place to stipulate the boundary, can there be any rational means to place it at one point rather than another? If not, isn't that just a *reductio* of my position? What can we say to an individual who insists that the boundary should lie between 1 and 2?

Once again, considerations from ethical theory might help out. On a Humean picture, we cannot arrive at moral truth based upon reason, but there still is a place for moral deliberation. For example, it is important to weed out false beliefs that might cloud one's judgment. If I believe that we should invade another country, and you believe that we shouldn't, and it turns out that my belief is based upon a false judgment about the weapons capabilities of the other country, then we can resolve our disagreement by simply looking at other facts. This is the case even on a Humean picture, in which ultimately judgments about right and wrong action are founded upon the sentiments.

Some Humeans, like Wiggins and Michael Smith³, believe that under certain conditions, we all will converge on the same moral judgments. I take this to be overly optimistic. However, I think there still are good grounds for believing that we can do *better*, morally, if we have more information and make our moral judgments without duress. With that in mind, what I imagine for semantic stipulations is the following. If we, *sub specie philosophiae*, are faced with a Sorites paradox, we can decide where to stipulate a boundary. How might this deliberation go? Though I'm not sure, I imagine it would begin by taking an initial survey of views from normal English speakers about

³ I call Smith a Humean because he too believes that moral judgments are based in our sentiments. However, he also holds the non-Humean view that our sentiments are rationally revisable.

where the boundary should go. Presumably, there would be a range of proposals. Though there will not be universal agreement on a point, there can be some kind of consensus – it is better to stipulate somewhere than nowhere. Against an outlier proposal of say, placing the boundary at 2, one could argue that if one places the boundary there, then the term would be useless because it would have almost no application. In this way, one can in fact use reason to help locate the point. So, if there is a disagreement about where to put the boundary, it would be akin to a first-order debate on an ethical question. Oftentimes, these answers do not hinge upon higher-order questions. But sometimes they do, and so sometimes, we will not be able to reach a single definite answer. But the hope is that this will not often happen, and that most first-order semantic debates about where to place the boundary will not be resolved to the satisfaction of all those engaged in the process *sub specie philosophiae*, even if it would not be satisfactory to every possible rational being.

The view I am advocating is an attempt to straddle the line between a complete indeterminist view and an epistemic view of one sort or another. If one claims that there are facts about the world which do determine that the boundary is located in a particular spot, then that will be a form of the epistemic view. In principle, we could know where the boundary was located if only we knew all the facts about linguistic behavior – so that in principle, rational considerations could give us an answer. But, because the facts don't determine the location of the boundary, and because the coherence of the concept requires some boundary, we must choose a boundary. Can we say that the facts determine a *range*? No, cannot say that, because then we'd be obligated to specify the range. Even though this is counter-intuitive, it is what higher-order vagueness forces us to recognize.

My own proposal is to avoid the discussion of whether facts determine a range. Instead, I see Sorites paradoxes as problems that would require a practical solution involving a stipulation. What I am hoping is that while there are no rules governing the making of such stipulations, there are first-order guidelines that we could follow. We could follow our own semantic intuitions. And I believe the best model for this is similar to a model for ethical deliberation – though there is no external ethical truth, and no set of sentiments that all humans share (even under ideal conditions of inquiry), there still can be ethical deliberation. One might call the process by which we may, *sub specie philosophiae*, stipulate boundaries a process of *thick legislation*.⁴ It is not an *arbitrary* sort of legislation, contrary to how Quine describes it, or as might be attributed to a more existentialist account, but a kind of legislation that has some basis in rational constraints, though it is not fully determined by rational or cognitive constraints. I take this kind of legislation, done *sub specie philosophiae*, to be analogous to the moral deliberation done in the PMP (post-Mackian position) I describe in Chapter Four.

On my account, reason does play a role in our thick legislations, but that reason does not determine a precise boundary. And I believe that our habits play a role, but that habits don't determine a precise boundary. To claim that either reason or habit or a combination thereof does so would be to accept some kind of epistemic view. So, what should we say to someone who chooses some very odd uses for ordinary language terms? The point is that in making our thick legislations, we are in fact setting up a basis for criticizing others, for the simple violation of our thick legislations. So even without

⁴ Thanks to George Graham first, for suggesting that a neologism would be helpful, and second, for suggesting this particular one.

meeting the traditional standards for normativity, we still have the power to discourage others from linguistic behavior that is highly aberrant.

However, I believe that the process might still leave one with some level of *semantic angst*. How does one choose when reason cannot fully tell us what to do? I cannot give an answer to this question in general, but I believe that the *angst* will not be felt on most occasions. When it is felt in particular cases, we can usually overcome the angst and make stipulations. Still, there may be some degree of residual semantic angst, but this does not seem to me to be such a bad thing. It might lead us to strive for a better scientific account of the world. Or it might guide us to a certain kind of humility concerning our understanding of how the world works. As in the quote at the start of this chapter, we should admire nature for its “thousand graceful subtleties”.

5.2.2 Choice and Determinism

In my dissertation proposal, I wrote that the dissertation would include a discussion of the following:

4e. Sorites and determinism

I have used the notion of linguistic choice here without regard to whether choices are themselves merely products of natural aspects of the world. If determinism is true, then choices are merely products of all the naturalistic facts in history, and if that is the case, and choice ultimately underlies normativity, then normativity really can be derived from history. But that seems to contradict my claim that the entirety of historical facts underdetermines normativity. This is a puzzle. In this section, I discuss whether this commits me to a form of indeterminism or to non-naturalistic explanations of choice. I conclude that my account is consistent with the assumption of determinism.

Since I wrote this, I have realized that this is quite a puzzle, indeed. If choice is determined, then it seems that the epistemic view is right in one way: there is a fact of the

matter about all our future choices, even if we don't know in advance what they will be, and so the sharp boundaries, which are determined by our choices, are for the most part, as yet, unknown.

My thinking at that time was that if choice is determined, the story that describes the ways in which it is determined would be so far removed from questions about how we would choose a boundary that the truth of determinism would be irrelevant. The resulting epistemic view would be so far from current naturalized views of semantics that my criticisms of these current naturalized semantics would still hold even if it was not an argument against all possible kinds of naturalism in semantics. I still take this to be true, but it does still seem to point in the direction of the epistemic view. I might be willing to bite that bullet, although I would like to register two problems determinism faces: first, I don't think we can make enough sense of the notion of causation, on which *determinism* depends, and also, if determinism can be made sense of, and is true, then almost everything about our folk understanding of human behavior would be undermined.

Those unargued thoughts aside, it may be helpful to consider the language use of what I will assume are non-free entities. What I'd like to show is that unfree things cannot use a language that is vague like ours. Let me assume for the moment that a hoverfly who snaps instinctively at every passerby of some range of size and velocity is unfree in a relevant way. (If you believe that hoverflies who do so still are free, then imagine a mechanical 'hoverfly-bot'.) Either one of two things is the case: its perceptual and cognitive mechanisms are so good that it draws a sharp line among the ranges, or not. In the former, there is no vagueness in the rule it follows. In the latter case, let's take an example that the smallest object it goes after is 3 mm wide, but because of imperfect

perception, sometimes it fails to go after an object that is 3.1 mm wide. We can assume that evolution made the hoverfly's cognition and perception accurate to a certain extent, but that any more refinement in perception might have been too costly; it does well enough with the cognition and perception it has to be able to reproduce. Should we describe the behavior of the hoverfly as that it is following a vague rule? I believe that that is not the best explanation. The rule might be stated as "go after anything larger than approximately 3mm". However, what we mean by 'approximately' might not be what the fly has in mind with its own innate approximations. Our notion of 'approximately' is informed by our own range of use of the term; one might ask, exactly what 'approximately' means in this context. So the entomologist might give a chart of frequency of pursuit of laboratory hoverflies to repeated stimuli surrounding 3mm. But unless this chart fit some precise function, then would it not be better to say that the hoverfly is not following a rule at all? Again, we might, for ourselves, say, vaguely, that the hoverfly goes after things that are more than about 3mm, but wouldn't the chart itself be a more accurate expression of the hoverfly's behavior?

However, if the scientists could find a point at which the hoverfly begins snapping at things, and then interpolate a probability function that goes from 0 to 1 of whether it snaps at things as they get bigger, then perhaps that could be considered a non-bivalent model for the rule. This would be vague in one sense, insofar as there is no sharp boundary from complete non-application to complete application of a rule, but in another sense, the rule would be very precise: there is still a precise point at which the hoverfly starts snapping.

But can this kind of model be transferred to human language? If our behavior with regard to terms such as 'red' is determined, then it seems that there can be a model for our use of it just like the one I have just assumed there is for the hoverfly. First, there is the issue that not all humans will have the same model. This initial objection can be overcome by using a larger model that includes all humans. But certainly, we will need some constraints: some humans are blind, or do not speak a language at all, or speak a language without an easy translation into the English red. But I will grant that we can find principles means for inclusion of humans as test cases in our model.

Does this show that there is a rule that humans are following with terms like 'red', and that it is a vague, degree-based rule? I might grant that if indeed the frequency by which humans use terms can be neatly reduced to some function, then we can have a model like the one I've assumed for the hoverfly. But an assumption of freedom seems to go against this, in two ways: first, and most importantly, even if the chart is accurate for past use of terms, there will be no guarantee that future use will conform (even if somehow we are able to keep test conditions the same across time). Second, there is nothing that would prevent someone, upon seeing the chart, from intentionally deviating from it. It is always an open question whether we should deviate from the use that is depicted on such a chart. It is not an open question for an individual who is programmed to follow the function the chart depicts. We could never program a computer, even one with unlimited computational and perceptual abilities, to use language exactly the way we do, for there is no exact way in which we use language, and no precisely model-able imprecise way in which we use language.

What I'd like to say is that even if determinism is true, in a broader sense, human behavior still varies in such a way that it can never fit into a function like that I've attributed to the hoverfly, and such a function is the only way we can avoid having to stipulate our use of terms to make them coherent. If it turns out that such stipulations themselves are products of past events outside of our control, then I might be willing to concede that much to the epistemicist.

5.2.3 Stipulations and Posits

Another potential problem for the normative choice view comes from the questions about what goes on in the process of thick legislation. I said earlier that we can choose which precise proposition we are expressing with our sentences. But in fact, we can do no such thing. My motivation for introducing the normative choice view is to be able to resolve Sorites paradoxes – we can choose the members of a discrete Sorites series between which the boundary falls. But this is not the same thing as choosing which precise concept we express with our terms. For we can never specify in complete terms a precise concept – we are simply making a decision in a particular case. As for Mt. Baldy, for example, it is not that we stipulate that a certain set of atom-time-slices composes Mt. Baldy. For we do not have anything near the cognitive capacities to be able to do so. So what are we doing in these stipulations? Sorenson phrases the argument:

Philosophers are strongly attracted to conventionalism. In particular, they think words (with the possible exception of natural kind terms) can have sharp boundaries only if speakers stipulate them. But how would the words used in the speaker's fiat get *their* sharp boundaries? If all of your 'definiens' lacks sharp boundaries, you will not be able to get any sharply bounded words by definition. And if some of your undefined terms have enough sharp boundaries to be used in a definition of a word that has sharp boundaries, then you must have antecedently had all the sharp boundaries

you needed. Stipulation can never increase the number of concepts that have sharp boundaries. (2001, p. 19)

A further problem is that there are reasons apart from vagueness for believing that a macro-level object can never be the same as a collection of atom-time-slices. For example, it can be argued from physics that it is more than merely the atoms that compose Mt. Baldy that make it a mountain – there are fields and other relations between atoms and between molecules that fit in the best scientific explanation of mountains. It might also be argued that for some objects, like a chicken, the causal and evolutionary history determines in part whether an object is a chicken, over and above its material composition. So stipulating a precise boundary or even a precise set of molecules is not the same as stipulating the complete and precise truth conditions for concepts like ‘mountain’ and ‘chicken’.

Perhaps Quine can provide some additional help here – what we are really doing is *positing* a boundary. But what does this mean, in relation to my own view? It seems that what we must do, *sub specie philosophiae*, is not specify precise and complete truth conditions for the concept in question, but merely suggest that in the case at hand, whatever the truth conditions are, they set the boundary at a certain point in a Sorites series. In other words, we posit that there is a mountain to which ‘Mt. Baldy’ refers, and stipulate that the boundary of it lies in a certain place – i.e., somewhere between two members of a Sorites series – if we are pressed by a Sorites. The benefit of this account is that we can solve every Sorites paradox that is presented to us by positing that the object has a boundary between two of the members in a Sorites. But what about those terms that we have not yet been given a Sorites for?

There seem to be two possibilities. The first is to claim, charitably, that all vague terms in ordinary language come with a kind of promissory note, such that their use allows us to use posits in the way I have just described when confronted with a vagueness-related problem. If that is the case, then there is no error in our usage, and the content of our usage is dependent upon our ability to make such posits. Second, our usage of some vague terms might not come with such promissory notes. (It is, perhaps, an empirical question concerning what goes on in the minds of language users when using vague terms.) If that is the case, then we should consider the normative choice theory to be a revisionist project – our language use is defective as is, but we should, from now on, consider our terms to have such a promissory note. Williamson writes (1994, p. 205): “Arguably, if we make stipulations about the future use of a hitherto vague term, we change its meaning, for anyone ignorant of the stipulations fails to understand it in its new sense.” However, if there is a division of linguistic labor, such that meaning depends upon what the linguistic experts say, *sub specie philosophiae*, then even without any further causal connection between individuals *sub specie philosophiae* and regular individuals, the ordinary language terms can acquire such a promissory note immediately.

5.3 Problems for Eliminativism

Even if these clarifications of the normative choice view are successful, the view still is highly dependent upon the arguments in favor of eliminativism. But eliminativism, on its face, seems absurd: how can it be possible to claim that macro-level objects don't exist? The concessive note to common sense that the normative choice view adds – allowing us to consider ordinary utterances to be true despite what seems to be

eliminativist metaphysics – doesn't seem to help the view become less absurd. So in this section, I'd like to discuss some potential objections to eliminativism.

5.3.1 Are There Really Any Eliminativist Explanations?

It might be objected that there cannot be eliminativist explanations of anything. In my attempts in Chapter Three, I used some macro-level terms in my explanations. But if I am really an eliminativist, should I not reject all such language, thus leaving no room for any eliminativist explanations? (On the assumption that we are nowhere near able to give a complete scientific, precise, account of everything.) As I argued in Chapter Three, just because the terms in our explanations are themselves vague does not entail that we should simply accept ordinary first-order language. Even if we cannot have absolute precision in our explanations, we should strive to have greater precision at each level of explanation, and as this occurs, our explanations seem to become better.

Perhaps one way to express this is in terms of promissory notes. The terms in the explanations that are relatively more precise than the terms being explained have slightly different promissory notes: perhaps they have notes that strive towards some kind of reduction, either to something scientifically viable or to a less controversial piece of sensory input. Or the very least, terms in these explanations might wear their promissory notes on their sleeves. For example, in giving an eliminativist explanation of 'landmass', I speculated that 'landmass' might be a scientifically viable term, insofar as we might be able to reduce it into some sort of precise disjunction. Or, for the notion of 'person', which also plays a role in the explanation of 'Mt. Baldy', perhaps we know that we have not given a precise sense to it, and that we may be doubtful that we can ever have a

proper reductive account, but we do have some guesses as to what we might give as stipulations for the notion, and on none of them is it such that we will have to include an object, Mt. Baldy, into our best explanations of 'Mt. Baldy'. Hence I believe that there can be eliminativist explanations, and that they may help guide our scientific enterprises. I believe that the historical association of eliminativist views, especially those in philosophy of mind, with forms of scientism, is misleading, for there can be non-scientistic forms of eliminativism.

There is a further consideration that helps the eliminativist. Importantly, the explanations I have in mind that the eliminativist gives will often make use of the very unit used in the major premise of a Sorites. Consider the following Sorites:

- (1) 100,000 grains of sand, properly arranged, make a heap. (Minor premise)
- (2) If n grains of sand, properly arranged, make a heap, then so do $n-1$. (Major premise)
- (3) Hence, 1 grain of sand, properly arranged, makes a heap.

But this is absurd, because 1 grain of sand cannot make a heap.

I imagine an eliminativist saying that there are no heaps – only collections of grains of sand. If one were to respond to the eliminativist by claiming that the notion of a grain of sand is incoherent, then the eliminativist could respond by noting then that the original Sorites, then, does not make any sense, and so is not something to be concerned with. Does this really work in the eliminativist's favor? I believe it does. Perhaps the objector might not herself be rejecting the notion of a grain of sand, but is suggesting that the eliminativist who rejects 'heaps' but, provisionally, accepts 'grains of sand', is being inconsistent. But the non-scientistic eliminativist I imagine need not have a policy that

only terms used in a final scientific theory can be used in any explanation, and so can escape this objection.

5.3.2 An Argument from Williamson

Suppose I have a twin, TW2, whose dimensions are the same as mine. It seems to follow that an utterance of the material conditional ‘If TW is thin, then TW2 is thin’ is true. Moreover, its truth seems to depend, not on what our shared dimensions are, but just on the fact that we share them. In particular, it seems to be true even if we are both borderline cases of thinness. Since it is true, it must have said something. The conditional says something only because its antecedent and consequent also do... If the antecedent or consequent lacked content, so would the conditional as a whole. Thus ‘TW is thin’ has content; it says that something is the case. (1994, p. 196)

The eliminativist might argue that there is no property of thinness, and hence nothing to give the term ‘thin’ content. This argument purports to show that terms such as ‘thin’ do have content, and a background assumption is that they get their content from the thing that they represent. However, there are problems with this argument.

Williamson’s claim that if a sentence is true, it must say something, is difficult to support. One might argue that tautologies, for example, are true but do not say anything. “If I am a floob, then I am a floob”, on some accounts, may be considered a true sentence, even though ‘floob’ does not have a sense. Second, I think that the sentence ‘If TW is thin, then TW2 is thin’ follows from the generalization that if a thing is thin, then something else with the same dimensions is also thin. This is one condition on the notion of *thinness*. However, one condition is not enough to fix the content of *thinness*. So, the conditional can appear to be true even if the antecedent and consequent do not have determinate senses. I can stipulate that whether something is a floob depends upon its eye color, but this isn’t enough to determine a sense for *floob*. It would seem to allow us to

say that ‘If I am a floob, then my twin (who has the same eye color) is also a floob’ is true, under some natural interpretation. One might object – on the grounds of some theory about conditionals – that if ‘floob’ really doesn’t have a determinate sense, then even though the conditional sentence just mentioned seems to be true, it really can’t be. If this is the case, then we should then say that the sentence ‘If I am a floob, then my twin (who has the same eye color) is also a floob’, insofar as we wish to hold it to be true, really does not have the logical form of a conditional. Rather, though it looks like a conditional, it is merely a way of expressing the condition whether something is a floob is dependent upon its eye color, and this does not require ‘I am a floob’ to have a determinate sense.

5.3.3 Is it Really the Best Explanation?

The argument for eliminativism went like this. First, we must reject moderate indeterminism because of problems of higher-order vagueness. Second, we must consider whether the epistemic theorist’s explanations of linguistic practice, and of the world, are better explanations than the eliminativist’s. (On the assumption that my argument from 5.3a is correct that there is such a thing as an eliminativist explanation.) My claim was that while the eliminativist’s explanations weren’t perfect, but the failures aren’t enough to warrant a simple acceptance of macro-level language. It makes more sense to suppose that macro-level objects do not exist as such than to suppose that either the natural world or our language use sets out a sharp, unknown boundary.

There really is no Mt. Baldy; there just is a big landform. Although using the term ‘Mt. Baldy’ seems to reify a portion of the landform, we actually fail to meet the conditions needed to be able to reify it – Sorites demonstrates that our term ‘Mt. Baldy’

does not have sharp boundaries. Terms must have sharp boundaries in order to have reference because, first, the world itself is not vague in the relevant way such that the lack of boundaries in our concepts simply maps onto the vagueness of the world, and so in order for a term to truly describe the world, it must have a boundary, and second, the argument against moderate indeterminism shows that no sense can be made of the notion of an unsharp boundary. Hence 'Mt. Baldy' does not refer – there is no Mt. Baldy. But then, we can note how we do seem to have some success using the term – it helps us coordinate our activities. So the normative choice theorist says that we might as well keep terms like 'Mt. Baldy', and to do so we can use a notion of normative choice. This is partly a pragmatic argument.

But I can see how this double-cross move might be unsatisfying. It leaves open the question of why there is any pragmatic benefit in using terms at all. In ethical decision making, we are thrust into the future, and have to decide what to do, and this, I have supposed, provides a grounding for ethical deliberation. But for language use, if the eliminativist is right, it is difficult to see how there can be any justification in using language – if there are no real macro-level objects, then how is it that using terms like 'Mt. Baldy' does anything to help us? If using language were intrinsically pleasing, then one might be able to keep a pragmatic account. But it seems like the very reason why language is helpful to us is more than that it gives us pleasure in its own use, but that we use language to communicate with each other truths about the world.

One might argue that semantic normativity just is a form of ethical normativity. We are pragmatically justified in using language because language use is an activity, just like other activities, and it seems to help us, and that is the one pragmatic standard for all

normative justification. In other words, there's only one *entrée* into the normative – linguistic normativity just is ethical normativity. But I think there are two problems with this. First, it doesn't show how it is that language has *any* pragmatic, ethical benefit if terms of ordinary language do not refer. And second, the pragmatic move seems to fail to separate the necessary distinction between helpfulness and truthfulness. This is a classic problem for pragmatism, and I do not believe it has been solved.

I think that the eliminativist may have a response: to question exactly what the conditions for *semantic success* are. If we assume that we are *successful*, then we have introduced normativity into the picture. This is the *entrée* into semantic normativity that we were looking for. If one were inclined towards transcendental arguments, one might even claim that we cannot make sense of our selves or of the world without presupposing such semantic success. The argument might go that one thinks in concepts, and concepts necessarily have a semantic component, and we can only make sense of the semantic component by believing that at least some of the time, we have semantic success. As such, normativity in semantics is not just a choice, but a necessary component of our understanding of the world. I think there is more to be said for this transcendental argument, but I believe it to be unsound: we can make sense of our selves and the world without semantic notions. I will not try to pursue this issue further, but I do believe that it is a possible direction to go in.

Putting the transcendental argument aside, there still seems to be a conceptual and empirical question about what success is and whether there is any success. The eliminativist may respond: what, *exactly*, is success? Unfortunately, it seems that *success* is yet another term that can be Sorites-awayed. Let's take an example: You tell me to

meet you in the Louvre, in front of the Mona Lisa, next Tuesday at noon. Sure enough, we meet there. There seemed to have been semantic success – otherwise, what could explain our both meeting there. But what if I had come one second later? Or was standing an inch further away? That meeting would still, I assume, provide evidence that there was a semantic success. But it seems obvious that a Sorites can be made quite easily out of the notion of ‘was a successful communicative/semantic event’, by adding inches of distance concerning the precise location where we meet. To deem an event as evidence of linguistic success involves a choice, and there are no complete choice-independent criteria for what success amounts to.

The normative choice theorist might agree that if we *assume* that there is success, we can put normativity back in semantics. As such, the only way we can explain success using a term like ‘Mt. Baldy’ is to believe that there really is such a mountain. (And that would entail, if my arguments from previous chapters are correct, that there is a precise mountain.) But the eliminativist may claim that even success is a matter of choice. Still, I think that we *should* choose to say that there is semantic success, and so particular claims are true in case the best explanation of the supposed semantic success involves the existence of some object to which the term refers. That’s what I believe what happens when, *sub specie philosophiae*, we posit objects/properties.

It might be helpful to note that, as I tried to argue in Chapter One, that some notion of *representational* success can (and should) still survive even if one claims that there is no semantic success. (This in itself shouldn’t be too far-out a claim: it seems highly plausible that non-human animals have representational success without conceptual success; I am simply extending this to humans, and taking our use of concepts

as a super-added piece of our evolutionary history.) Rejecting that there is a fact of semantic success need not lead us into Pyrrhonianism. The epistemological picture I advocate is something like an old-fashioned empiricist myth of the given. We do perceive a manifold, which does not come conceptually, and that after that perception, we somehow add concepts. I am aware that this epistemological story is not in vogue today, but the claim that our percepts do come concept-laden will prevent any simple solution to regular Sorites, as well as phenomenological Sorites that I give an example of in Chapter One.

The is/ought gap in normative semantics, for reference, can now be bridged in the following way:

1. We regard our use of a particular linguistic term, T, as giving us pragmatic success.
2. The best way to explain that success is to posit that there is an object/property to which T refers.

Therefore:

3. T refers to that object/property. (Normative semantic conclusion)
- (A similar story can be told for the normativity of truth.)

Given that I choose to accept 1, my account is not simply an eliminativist account.

Choice runs so deep that we can choose to reject eliminativism.

There is one further point concerning positing that is worth mentioning. I mentioned earlier that we cannot simply stipulate a sum of atom-time-slices which we take to be Mt. Baldy, partly because we cannot identify mountains with a sum of atom-time-slices. Part of the problem is what leads to Kant: what, exactly, is the nature of objects in the external world in themselves? It seems that I share Kant's view that

empiricism will leave us without a good account of the real nature of the external world. Is my suggestion that we posit objects simply a move to say that we should believe that there is an world of things in themselves about which we know nothing, but that we should instead consider there to be an ‘empirical’ world that forms the basis of our thinking? Although my view is indeed similar to Kant’s in several ways, it differs from Kant’s in that, first, I reject the transcendental kind of argument I note above; second, I am motivated by best-explanation arguments, which do not in principle rule out our knowledge of at least some parts of the world as it actually is; third, I reject the Kantian/McDowellian epistemology according to which “intuition without concepts are blind”.

5.4 Concluding Thoughts

Let me conclude by briefly reviewing the argument I’ve given in this dissertation. I begin by claiming that views that reject bivalence have one or both of the following serious problems:

1. They still rest on some sharp boundary, even if it's not a sharp boundary between the yes/no cases. This is the problem of higher-order vagueness.
2. They leave it a mystery how the logic actually connects to facts about our natural language use. (e.g., the problem for a standard degree theory that there is nothing in our language use that shows that, say, "That's green" is 72% true rather than 71% true.)

I reject the epistemic theory on the ground that it is an unsuccessful account of what actually is, and that the best explanation of our language use involves no requirement that there be macro-level objects to which our terms refer. At the very least, the eliminativist

explanation is a better one than one according to which there really are objects with precise boundaries, and this seems to be the only alternative to the eliminativist account.

I try to take a methodological middle ground between a naive acceptance of our ordinary first-order language and the unrealistic desire for a full scientific theory of everything. Williamson insinuates that the claim that there are no heaps rests on a belief that the only real things are those that have a place in a reductionist/scientific explanation of everything. I try to show that eliminativism doesn't rest on such a scientific demand – we have plenty of other cases where we've properly gone eliminativist about a discourse even without a full explanation of everything, or even a full explanation of the phenomenon in question.

And I don't have a full argument, independent of Sorites, why eliminativism is the best metaphysical account to take. But I do feel that the Sorites itself should be considered a *reductio* against the existence of most macro-level objects. I take it as a kind of empirical/scientific argument within our analysis of ordinary objects or object terms, and not just an annoying argument that can be applied to any term prior to the knowledge of empirical facts. One needs the notion of a grain, or an atom, or some other small unit of measure, to phrase most Sorites. Prior to Democritus, it would have been impossible to phrase most Sorites paradoxes. (A reason why the original Sorites was about heaps of uncontroversially discrete units.) I take Democritus's atomistic insight as opening the door to a better empirical picture of how the world really is. For the world could have turned out differently – chickens and chairs and mountains might have had complete essences such that one could not remove any part without destroying the whole thing. That something is susceptible to a Sorites is a discovery. It is progress, in a way.

Unfortunately, eliminativism is itself counterintuitive, but the normative choice account is an attempt to get around the problems for eliminativism. What I've tried to sketch in this last section is a theory of truth which is somewhat pragmatic, but that doesn't rest on pragmatism – because, in part, a pragmatist needs some explanation of why it's pragmatically beneficial to use these terms, and it seems that the most likely explanation is that they refer to real objects. So, I've introduced the notion of a thick legislation to explain the normativity of language, and to allow us to keep the truth of ordinary language.

There might be something duplicitous sounding in this account, in a way that a political analyst might denounce. Are there really macro-level objects or not on my account? Some philosophers might deny that we can make sense of general questions like “are there really macro-level objects?” Many have argued against correspondence theories of truth and reference in claiming that we can't even make sense of the conditions of what it really means to refer. These philosophers tend to be more attracted to minimalist accounts of truth and reference. Others, like 1980's Hilary Putnam, claim that the world does not individuate itself into objects, but that humans do that. The Sorites is an argument against that view, as well, unless it is interpreted as the kind of individuation that occurs with thick legislation, rather than with some natural habit of object-making.

My feeling is that Sorites is not an argument against correspondence theory. I think correspondence theory is perfectly intelligible, and probably the theory of truth that best captures most of our intuitions about truth. Sorites is both an argument against a

minimalist theory of reference – we need to appeal to semantic facts to have any hope of resolving it – and of what we take to be truths on a correspondence theory.

Most people who reject correspondence theory take it to be no loss whatsoever. I myself feel the loss – but it is a loss because the world, and ourselves, do not live up to what we had hoped for. We cannot maintain the view that our language reflects the world in a way which corresponds to how it really is. But, for reasons of both charity and explanation, we need to give some positive account of what we are doing in using language, and so I've attempted to give the account of thick legislation to help make sense of how one can say that we do utter truths, and do refer to objects. Ultimately, I hope, that this is not such a bad result – it will help enable us to see the world better – as Quine says, Sorites is a reason for precision in science.

One final note is that I really don't know what to say about *people, life, Avram Hiller, consciousness*. All of these are vague. Should we take an initial eliminativist step about them, too? I really don't know what the best explanations are for these things/terms, and my methodology says that we should take the initial eliminativist step if we have a better explanation of the phenomenon than the one that includes the ordinary ontology (with its requisite unknown/unexplained boundary). But I haven't seen an eliminativist explanation of these things that's better than the ordinary language one. So I'm not arguing for an across the board eliminativism of all macro-level objects. Additionally, it might be the case that at the micro-level, there are objects with sharp boundaries. My argument against the epistemic view was not an a priori argument. What we need to do is determine what our best explanations are of whatever phenomenon is in question. And I hope I have shown that our best explanations of ordinary language usage

do not require that there exist, in advance of our normative choices, things such as heaps and colored candies. But this should not be cause for despair.

Please allow me to posit, now, that this is a complete dissertation. In other words:

The End.

BIBLIOGRAPHY

- Baker, G., and P. Hacker. "Critical Study: On Misunderstanding Wittgenstein: Kripke's Private Language Argument." *Synthese* 58, (1984).
- Blackburn, Simon. "The Individual Strikes Back," in *Essays in Quasi-Realism*. New York: Oxford University Press, 1993.
- _____. "Moral Relativism and Moral Objectivity." *Philosophy and Phenomenological Research* 58 (1998), 195-198.
- Boyd, Richard. "Scientific Realism and Naturalistic Epistemology," *PSA* 80, vol. 2 (1980).
- Burgess, J.A. "The Sorites Paradox and Higher-Order Vagueness." *Synthese* 85 (1990): 417-474.
- Chang, Ruth. *Incommensurability, Incomparability, and Practical Reason*. Cambridge: Harvard, 1997.
- Changizi, M.A. "Vagueness, Rationality and Undecidability: A Theory of Why There is Vagueness." *Synthese* 120 (1999): 345-374.
- Devitt, Michael. *Realism and Truth*. Oxford: Basil Blackwell, 1984
- Dummett, Michael A.E. *Truth and Other Enigmas*. London: Duckworth, 1978.
- Edgington, Dorothy. "Vagueness by Degrees." In *Vagueness: A Reader*. Keefe, R. and P. Smith, eds. Cambridge, MA: MIT Press, 1997.
- Fine, Arthur. "Unnatural Attitudes: Realist and Instrumentalist Attachments to Science." *Mind*, 95 (1986), 149–77.
- Fine, K. "Vagueness, Truth, and Logic." In *Vagueness: A Reader*. Keefe, R. and P. Smith, eds. Cambridge, MA: MIT Press, 1997.
- Frege, G. *Translations from the Philosophical Writings of Gottlob Frege*, Geach, P. and M. Black, eds. Oxford: Blackwell, 1952.
- Graff, Delia. "Shifting Sands: An Interest-Relative Theory of Vagueness." *Philosophical Topics* 28 (2000): 45–81.
- Haack, S. *Philosophy of Logics*. Cambridge: Cambridge University Press, 1978.
- Hale, Bob. 1997. "Rule-Following, Objectivity, and Meaning." In *A Companion to the Philosophy of Language*. Hale, B., and C. Wright, eds. Oxford: Blackwell, 1997. 369-396.
- _____, and C. Wright, eds. *A Companion to the Philosophy of Language*. Oxford: Blackwell, 1997.
- Harman, Gilbert. *The Nature of Morality*. Oxford: Oxford University Press, 1977.

- _____, and Judith Jarvis Thomson. *Moral Relativism and Moral Objectivity*. Cambridge, MA: Blackwell, 1996.
- Heck, R. "A Note on the Logic of (Higher-Order) Vagueness." *Analysis* 53 (1993): 201-208.
- Hiller, Avram. "Are Minds Vague." Presentation at *Towards a Science of Consciousness* conference, Tucson, AZ. April 2000.
- _____. "Brains in a Vat in a World." Unpublished.
- _____. "How to Be an Ethical Quasi-Realist." Unpublished.
- Holtzman, Steven H. and Christopher M. Leich, eds. *Wittgenstein: To Follow a Rule*. London: Duckworth: 1981.
- Horgan, Terry. "Robust Vagueness and the Forced-March Sorites Paradox". In *Philosophical Perspectives, 8: Logic and Language*. Tomberlin, J., ed. Atascadero, CA: Ridgeview, 1994.
- _____, ed. *Vagueness. Southern Journal of Philosophy*, Supplement 35 (1995).
- _____. "The Transvaluationist Conception of Vagueness." *The Monist* 81 (1998): 316-33.
- Hume, David. *A Treatise on Human Nature*. Oxford: Oxford University Press, 2000.
- Keefe, Rosanna. and Peter. Smith, eds. *Vagueness: A Reader*. Cambridge, MA: MIT Press, 1997.
- Keefe, Rosanna. *Theories of Vagueness*. Cambridge: Cambridge University Press, 2000.
- Kripke, S. *Naming and Necessity*. Oxford: Blackwell, 1980.
- _____. *Wittgenstein on Rules and Private Language*. Oxford: Blackwell, 1982.
- Lewis, D.K. "General Semantics." *Synthese* 22 (1970): 18-67.
- Loewer, B. "A Guide to Naturalizing Semantics." In *A Companion to the Philosophy of Language*. Hale, B., and C. Wright, eds. Oxford: Blackwell, 1997.
- Machina, K. "Truth, Belief, and Vagueness." In *Vagueness: A Reader*. Keefe, R. and P. Smith, eds. Cambridge, MA: MIT Press, 1997.
- McDowell, John. "Wittgenstein on Following a Rule." *Synthese* 58 (1984): 325-363.
- _____. *Mind and World*. Cambridge, MA: Harvard University Press, 1994.
- McGee, V. "Kilimanjaro." *Canadian Journal of Philosophy* Supplementary Volume 23 (1997).
- _____. "No Fact of the Matter." Presentation given at the Chapel Hill Colloquium (October 16, 1998), University of North Carolina, Chapel Hill, NC.

- _____, and Brian McLaughlin. "Distinctions Without a Difference." In *Vagueness*. *Southern Journal of Philosophy*, Supplement 35 (1995).
- McNaughton, David. *Moral Vision: An Introduction to Ethics*. Cambridge: Blackwell, 1988.
- Millikan, Ruth Garrett. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press, 1984.
- _____. *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press, 1993.
- Mott, Peter L. "Margins for Error and the Sorites Paradox." *Philosophical Quarterly* 48 (1998): 494-504.
- Papineau, D. *Philosophical Naturalism*. Oxford: Blackwell, 1987.
- Parfit, Derek. *Reasons and Persons*. Oxford: Oxford University Press, 1984.
- Peacocke, Christopher. "Are Vague Predicate Incoherent?" *Synthese* 46 (1981a): 121-141.
- _____. "Reply: Rule-Following: The Nature of Wittgenstein's Arguments." In Holtzman, Steven H. and Christopher M. Leich, eds. *Wittgenstein: To Follow a Rule*. London: Duckworth: 1981b.
- Putnam, Hilary. *Philosophical Papers 2: Mind, Language and Reality*. Cambridge: Cambridge University Press, 1975.
- _____. *Reason, Truth, and History*. Cambridge: Cambridge University Press, 1981.
- _____. "Vagueness and Alternative Logic." *Erkenntnis* 10 (1983): 371-380.
- Quine, W.V. *Word and Object*. Cambridge, MA: MIT Press.
- _____. "What Price Bivalence?" *Journal of Philosophy* 78 (1981): 90-95.
- Raffman, D. "Vagueness Without Paradox." *Philosophical Review* 103 (1994): 41-74.
- _____. "Vagueness and Context-Relativity." *Philosophical Studies* 81 (1995): 175-192.
- Read, Stephen, and Crispin Wright. "Hairier than Putnam Thought." *Analysis* 45 (1985): 56-58.
- Sainsbury, R.M. "Degrees of belief and Degrees of Truth." *Philosophical Papers* 15 (1986): 97-106.
- _____. "Concepts Without Boundaries." Inaugural lecture, Department of Philosophy, King's College, London (1991); Reprinted in *Vagueness: A Reader*. Keefe, R. and P. Smith, eds. Cambridge, MA: MIT Press, 1997.
- Sanford, David H. "Borderline Logic." *American Philosophical Quarterly* 12 (1975): 29-39.

- _____. "Competing Semantics of Vagueness: Many Values Versus Super-Truth." *Synthese* 33 (1976): 195-210.
- _____. "Nostalgia for the Ordinary: Comments on Papers by Unger and Wheeler." *Synthese* 41 (1979): 175-184.
- _____. "The Problem of the Many, Many Composition Questions, and Naive Mereology." *Noûs* 27 (1993): 219-28.
- _____. "Review of Linda Burns: Vagueness: An Investigation Into Natural Languages and the Sorites Paradox." *Mind* 102 (1993): 357-360.
- _____. Unpublished paper on higher-order boundaries.
- Sayre-McCord, Geoffrey, ed. *Essays on Moral Realism*. Ithaca, NY: Cornell University Press, 1988.
- Smith, Michael. *The Moral Problem*. Cambridge: Blackwell, 1994.
- Sorensen, Roy A. *Blindspots*. Oxford: Clarendon Press, 1988.
- _____. *Vagueness and Contradiction*, Oxford: Oxford University Press, 2001.
- Sturgeon, Nicholas L. "Moral Explanations." In *Essays on Moral Realism*. Sayre-McCord, Geoffrey, ed. Ithaca, NY: Cornell University Press, 1988.
- Thompson, Evan. *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*. London: Routledge Press, 1995.
- Tye, Michael. "Sorites Paradoxes and the Semantics of Vagueness." *Philosophical Perspectives* 8 (1994).
- Unger, Peter. "I Do Not Exist." In *Perception and Identity*. MacDonald, G.F., ed. London: Macmillan, 1979. 235-51.
- _____. "There Are No Ordinary Things." *Synthese* 41 (1979): 117-154.
- Wheeler, S.C. "Reference and Vagueness." *Synthese* 30 (1975): 367-369.
- Wiggins, David. "Objective and Subjective in Ethics, with Two Postscripts on Truth." In *Truth in Ethics*. Hooker, Brad, ed. Cambridge: Blackwell, 1996
- Williamson, T. *Vagueness*. London: Routledge, 1994.
- _____. "What Makes it a Heap?" *Erkenntnis* 44, 3 (1996): 327-339.
- _____. "Imagination, Stipulation and Vagueness." In *Philosophical Issues 8: Truth*. Villanueva, E., ed., Atascadero, CA: Ridgeview, 1997. 215-228.
- _____. "Margins for Error: A Reply to Mott." *Philosophical Quarterly* 50 (2000): 76-81.

- Wittgenstein, L. *Philosophical Investigations*. Translated by G.E.M. Anscombe. Oxford: Blackwell, 1953.
- Wright, C. "Language Mastery and the Sorites Paradox." 1976. Reprinted in *Vagueness: A Reader*. Keefe, R. and P. Smith, eds. Cambridge, MA: MIT Press, 1997.
- _____. "Further Reflections on the Sorites Paradox." 1987. Reprinted in *Vagueness: A Reader*. Keefe, R. and P. Smith, eds. Cambridge, MA: MIT Press, 1997.
- _____. "The Epistemic Conception of Vagueness." In *Vagueness. Southern Journal of Philosophy*, Supplement 35 (1995).
- _____. "Is Higher-Order Vagueness Coherent?" *Analysis* 52 (1992): 129-39.
- _____. *Truth and Objectivity*, Cambridge, MA: Harvard University Press, 1992.
- _____. "On Being in a Quandary: Relativism, Vagueness, Logical Revisionism" *Mind* 110 (2001): 45-98.

BIOGRAPHY

Avram Hiller was born on June 11, 1971, in the watershed of the Charles River in the town of Needham, Massachusetts, USA. He received his BA, *magna cum laude*, from the University of Pennsylvania, in May 1993. As a graduate student at Duke University, he received the following awards and fellowships, for which he is grateful: Philosophy Department Graduate Fellowship; James B. Duke Fellowship; Summer Dissertation Research Fellowship; Gerst Instructorship; Kenan Ethics Colloquium Fellowship; and two Graduate School Conference Travel Awards. He is a member of the Society of Duke Fellows. In addition to his graduate work at Duke, he has also been a visitor at Tufts University, the University of Arizona, Harvard University, and the University of North Carolina at Chapel Hill. He has taught philosophy at Duke University, Pitzer College, Claremont Graduate University, and Wake Forest University. In 1989, he won a Silver Medal in Mathematics at the Massachusetts State Academic Decathlon. He received the Camp Kingswood Athletic Director's Award in the summer of 1984. He once made 51 consecutive free throws, though not as a member of the Duke Blue Devils. He sings, plays guitar, takes photographs, and explores wild places.