The Dissertation Committee for Ivan J Hu
certifies that this is the approved version of the following dissertation:

Epistemicism

Committee:

____
Richard M. Sainsbury, Supervisor

____
Johan A. Kamp, Co-Supervisor

____
Daniel A. Bonevac

____
Robert C. Koons

____
Joshua Dever

____
Diana Raffman

Epistemicism

by

Ivan J Hu, B.A.

DISSERTATION

Presented to the Faculty of the Graduate School
of the University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

The University of Texas at Austin
May 2015

For my parents

Acknowledgements

The greater part of this work grew out of many memorable hours of rich and lively conversation with Daniel Bonevac and Hans Kamp. I am indebted to them both. My best times of productivity regularly came from talking with Dan. Over these years, the extent to which I have benefited from those moments—learning from his intellectual creativity, sheer breadth of knowledge, and contagious enthusiasm when discussing all things philosophical and more—is beyond repayment. My debt to Hans is equally deep. His uncanny ability to quickly get to the heart of the issue when discussing any philosophical matter, with all seriousness and rigor yet always with a constructive attitude (a rare quality), set for me an example of true profundity tempered with intellectual humility.

I am privileged to have a dissertation committee consisting of so many teachers whose own contributions to the philosophy of vagueness are each noteworthy in their own right. Much of my thinking on vagueness is owed to reading Mark Sainsbury's writings on the subject. His insightful feedback has always challenged me to fine tune my own work while achieving more in less through leaner economy of expression. I owe much to Josh Dever, whose support and availability proved invaluable on many occasions and whose brilliance in pedagogy never ceases to amaze. Diana Raffman was kind enough to join my committee at a late stage, though I had already learned much from reading her work on vagueness in previous years. Rob Koons was the first professor I met as a graduate student. I am grateful for all the advice and encouragement he has given me over the years, especially regarding the importance of knowing one's own broader philosophical outlook.

I have also enjoyed many discussions about vagueness with other faculty members not on my committee, including David Sosa, Cory Juhl, Ray Buchanan, Galen Strawson, and Jon Litland. Thanks also to my fellow graduate student colleagues who have offered comments and feedback related to this work: Fatema Amijee, Kate Ritchie, Richard Davis, and Brian Cutter.

I should also thank John Campbell, whose supervision during my undergraduate years did so much to sharpen my philosophical writing and instill within me the discipline of being thorough. I probably would not have pursued scholarship in philosophy, or even the greater humanities, were it not for Gary Hendrickson, whose encouragement in high school was pivotal when I first encountered philosophy through literature.

For my dearest family and friends: I thank my parents for being among the few Chinese parents on this earth who would ever let their son major, never mind pursue a doctorate, in the humanities. Their love and support over the years in allowing me to find my own path has been invaluable. Thanks to my Berkeley CKG1 peers for believing in me, especially Brian Chang. Thanks to Greg and Mary Jane Grooms, whose home has been a haven of constancy and rest during this otherwise often disjointed journey of graduate school. Finally, thanks to Janie for supporting me while I got this done, through both sacrifice and celebration.

Ivan J Hu
The University of Texas at Austin
May 2015

Preface

Back in 2011, while reading some of the latest on vagueness—by David Barnett (2009), Brian Weatherson (2006), Patrick Greenough (2003), Cian Dorr (2003), Crispin Wright (2001)—I was shocked to discover that people were denying what seemed the most plausible part of epistemicism: the claim that vagueness entails ignorance, or what I call the UNKNOWN principle—e.g. if it's vague whether somebody is bald, you don't know whether he's bald. Such objections went either undiscussed or unnoticed. Evidently no one deemed it worthwhile to defend the epistemicist theory, even against criticisms that (I felt) were surely wrong. So I took the task upon myself. What emerged was my prospectus thesis: a longer, unwieldier version of Chapter 1 ("Epistemic Gaps") written in response to Barnett's (2010) claim that UNKNOWN fails to capture anything *essential* about vagueness: in a hypothetical scenario where our cognitive faculties were improved so as to enable knowledge, vagueness might still be possible, since people's intuitions could still be vague over vague matters; it just wouldn't entail ignorance.

It took 70 pages (including eight whole pages of endnotes) to respond to Barnett's one counterexample. The reason being: it took quite some work to figure out the logical structure underlying Barnett's arguments. It wasn't trivial, for instance, spelling out how the determinacy (clarity) operator used to express claims of vagueness was supposed to interact with a knowledge operator and with conditionals. But that's what's exactly at issue because UNKNOWN involves all three. What I did was take some of the methods of reasoning about vague claims, which I had picked up from Greenough and Barnett, and apply these exhaustively to a whole range of claims—response-dependence conditionals (if it seems that *p*, one can know that *p*), ambivalence (if it both vaguely seems that *p* and vaguely seems that $\neg p$), intuition-completeness (if *p*, it will seem that *p*), connections between clarity and knowledge (if *p* is knowable, *p* is clear, or vice versa)—tracing out all their possible interrelations, not just when things were vague, but (more crucially) when things were higher-order vague. What I discovered was that you can derive quite a lot just by using the minimal resources of a KT modal logic for determinacy (clarity) and knowledge. For instance, the type of response-dependence Barnett wanted, where one is never quite determinately ignorant about vague matters, can be shown to be compatible with UNKNOWN and any broader encompassing epistemicist theory. It merely imposes an extra constraint: you never have definite borderline cases. That's enough to get you the same results, without meddling with the intuitive principles already in place linking vagueness and ignorance.

I took this as evidence for Greenough's (2003) fundamental insight that some principle connecting vagueness to unknowability—contrary to the skeptical attitude of recent work—really *could* serve as a minimal theory of vagueness, or something like it. This was the seemingly undeniable part of epistemicism. To that extent, challenging epistemicism on *epistemic* grounds struck me as unthinkable. I found no reason to believe that any of the objections to UNKNOWN posed by Wright (2001), Dorr (2003), or Barnett (2010) were insurmountable. For their alternative proposals for what the relation between vagueness and unknowability should look like instead all represented variants of the same general type of view concerning the epistemology of vagueness, the overall structure of which was something that seemed easily recoverable on epistemicism—one merely had to deny the existence of knowable or definite borderline cases.

I later came to think there were independent reasons to question, on epistemic grounds, this entire family of views advanced by Wright, Dorr, Barnett, and others. In Chapter 2 ("Vague Unknowns"), I go through each of their alternative proposals one by one and show why there's independent reason to find them implausible. What's more, I came to find what Williamson himself had to say about the relations between vagueness and ignorance deeply unsatisfactory, or downright puzzling at best. Since then I've come to doubt whether Williamson's account should even be taken to be representative of the epistemicist perspective on vagueness at all (see Chapters 2 and 4; see also discussion in Outlook).

That said, I should pause here to pay my intellectual respects. It was while reading Timothy Williamson's *Vagueness* (1994) in a seminar on vagueness co-taught by Hans Kamp and Mark Sainsbury in the spring of 2010 that I realized I *had* to work on vagueness. No other philosophical topic held so great an aura of enigma, for so ostensibly simple a problem as the Sorites paradox. Williamson's book made this salient. The unparalleled rigor and clarity found in his discussion of the subject was infectious.

The first paper I ever wrote on vagueness was for that seminar: an embryonic version of Chapter 4 ("Margins for Error in Meaning"). I wrote it five years ago, while still very much within a Williamsonian frame of mind and earnestly trying to press out the wrinkles in his theory. Although I've since then come to think there are some fundamental, irremediable differences in my approach in addressing the epistemology of vagueness or even what it means to be an epistemicist theory, I still find it valuable to confront Williamson's theory on its own terms, if not because it raises a host of fundamental questions about the epistemology of meaning.

My general line of argument there is this. There's a real worry about whether the margin for error principles, so crucial to Williamson's explanation of vagueness-related ignorance, won't also rule out semantic knowledge, if the meanings of our vague expressions are themselves in constant flux. Even Sainsbury in his illuminating (1997) comment on *Vagueness* tries to come to the rescue by developing Williamson's externalist idea of semantic knowledge as linguistic induction. But I think that Williamson must ultimately give up his claim of *exact* knowledge of meaning. By the lights of his own view, inexact knowledge is knowledge of something falling within a range, where for each thing within that range, you don't know that it's not *that*. But this is exactly the sort of structure you get on Williamson's picture of meaning: a range of candidate meanings for 'bald' arranged by their different cutoffs, where each remains an epistemically possible meaning for 'bald', since you don't know that 'bald' doesn't mean *it*.

Let me offer an analogy. I used to travel frequently between Mainland China and Hong Kong. In the mid-2000s, the currency exchange rates for the Chinese *renminbi* and the HK dollar relative to the US dollar were comparable: roughly 1 USD = 7.75 HKD = 8 CNY. At times I would get confused, and use the wrong rate when converting prices back to USD. No doubt this led to instances of unwarranted consumer overconfidence. On Williamson's picture, meanings for vague terms are like currency rates: they are shared, collectively determined semantic features of our vague expressions that, although constantly fluctuating, nonetheless enable communication, much like how (unpegged) currencies, despite having precisely determined yet freely fluctuating exchange rates, nonetheless continue to facilitate consumption, trade, and exchange. On such a picture, inexact semantic knowledge should be the norm: surely, we can only know the meanings of our vague terms *inexactly*, just as any currency exchange rate, albeit perfectly determinate (down to the fraction of a cent) at any given point in time, can only be approximated. Having a mistaken belief about the meaning of a vague term does not prevent communication of an unjustified idea, any more than a mistaken belief about the real rate for a currency prevents making an unjustified purchase (if only it did!). Any transaction, whether monetary or verbal, is blind to the attitudes of the involved parties.

I must add that, whatever my disagreements with Williamson's epistemicist theory, among the criticisms facing his account I find most pressing, the endorsement of sharp cutoffs for vague predicates is *not* one of them. I have never quite fully understood what is so appalling about the mere suggestion that things stop being heaps past a certain point, especially if this is a reasonable conclusion—indeed, the *only* conclusion, given classical reasoning—to draw from the Sorites paradox. To this day, I must confess, I still fail to appreciate exactly why the view attracts such hostility from philosophers. Testing the same intuitions on untutored "folk" ears has so far failed to reproduce the same levels of prejudice.

Contextualist theories comprise my last main target. Through discussion with Josh Dever I came to see that contextualists too could deny epistemicist principles like UNKNOWN. Although none of the contextualist literature discusses the epistemic issues surrounding vagueness at any great length, the idea would be: you can settle the status of borderline cases differently in different contexts (calling a borderline F 'F' in one place, 'not-F' in another), and to do so is to issue a knowledgeable verdict about the borderline case in question; so vagueness doesn't preclude knowledge after all. Any such permissibilist conception of vagueness threatened to be at odds with the UNKNOWN principle, and therefore with epistemicism at large.

Chapter 5 (formerly "Epistemicism, Paradox, and Conditional Obligation", now "Ignorance and Open Texture") was a partial answer to this challenge. There I reply to Shapiro (2006), who accuses epistemicism of giving the wrong predictions about the normative consequences of vagueness. The details revolve around showing why Shapiro's counterexample doesn't work, since the line of reasoning he uses against epistemicism rests on the same sort of faulty reasoning you use to derive paradoxical results in the miners paradox. The consensus from recent work on the miners paradox (Kolodny and MacFarlane 2010, Willer 2012) is that conditional obligations carry a sort of information sensitivity: what you ought to do may depend on what you know or what you don't know. What I do is apply this lesson to Shapiro's case: since epistemicism tells you that vague matters are unknowable, any information deficits resulting from vagueness will affect your normative situation, and possibly remove obligations you would otherwise have if vagueness wasn't there to keep you from knowing the facts at hand. That's enough to get you the sort of permissibility phenomena Shapiro thinks is closely tied to vagueness, except now the explanation for such vagueness-related permissibility is epistemic rather than semantic. Going this route, however, does require relinquishing some central tenets of Williamson's version of epistemicism—including his commitment to classical logic, at least for reasoning about norms.

In trying to combine epistemicism with open texture, I was staking out a claim on the underexplored topic of the normativity of vagueness. On Dan Bonevac's suggestion, I applied the same logical technique I had developed earlier to this set of issues, to bring out exactly how vagueness interacts with conditional obligations. Shortly afterwards, I started toying with the idea of understanding tolerance in terms of deontic conditionals. The idea was that, assuming we can conceive of the very tolerance principles underlying the Sorites paradox as a series of conditional obligations ("If you call one guy bald, you've got to call the next guy bald"), the ignorance-entailing feature of vagueness should block Sorites reasoning, in roughly the same way vagueness-related ignorance rendered conditional obligations inert in Shapiro's case. What was needed was a suitable logic for deontic conditionals. For that, naturally, I turned again to Dan, whose joint work with Nick Asher in the 90s gave me exactly what I needed. Asher and Morreau's system of commonsense entailment (1990) provided a theory of defeasible reasoning, with a conditional specially designed to handle generics, like "Birds fly". Later work by Bonevac & Asher (1996, 2005) extended this account with an analysis of the deontic conditional, suited for representing prima facie obligations and conditional obligations, in order to solve a number of deontic paradoxes. What I sought to do was borrow the analysis and apply it to problems in vagueness. The result turned out to promise a compelling alternative to contextualist accounts.

Critical reaction to contextualist theories of vagueness, not just by epistemicists but generally, is scant, despite their recent growth in popularity among philosophers. Williamson's dismissive and inadequately brief remarks in his (1994:214–5) only target contextualist theories that liken vague terms to indexicals. Contextual variation, he claimed, is no essential feature of vagueness:

fixing the context leaves vagueness intact. Of course, that was before contextualist approaches to vagueness had entered the mainstream philosophical literature. What Williamson didn't seem to anticipate was how few contextualists about vagueness—Soames (2002:445) aside—would now take seriously the idea that vague expressions are to be understood on the model of indexical expressions: it's not the meaning or content of a vague term that changes as the context shifts, only its standards of evaluation that vary with context. Almost a decade later, Stanley (2003) echoes the same misunderstanding. Stanley argues that the standard contextualist strategy for handling the Sorites arguments—predicting that vague terms are, like indexicals, contextually variable in meaning and so express different contents at different stages of the Sorites series—fails to apply to verb phrase ellipsis versions of the Sorites ("If man #1 is bald, so is #2, and so is #3, and…"), which force an invariant reading upon the relevant vague terms, blocking any possible content shifts induced by contextual variation. The idea is the same as Williamson's: vagueness remains when you fix the context—which here just means that the vague predicate must pick out the same property at each step. As Diana Raffman's (2005) and Delia Graff's (2008) replies confirmed, this objection clearly misconstrued contextualism about vagueness, since it mistook the contextualist proposal as advancing some sort of hidden indexical account of vague expressions. Åkerman & Greenough (2009) have argued that Keefe's (2007:286) objections to contextualist views rest on similar unjustified assumptions. Yet besides this small slice of the literature—alongside relatively minor pieces, like Robertson's (2000) critical notice of Soames (1999) and Sorensen's (2008) whimsical review of Shapiro (2006)—there isn't much else by way of criticism for contextualism.

My own discontents with contextualist accounts, mostly borne out of conversation with Mark Sainsbury and reading his papers, was their apparent obsession over the forced march Sorites and the insistence that any reversal of judgment be saved until the very end of the Sorites series. This didn't seem to account for the unease one feels about extending a vague predicate across the borderline range to begin with. A normative interpretation of tolerance, by contrast, did appear able to account for this datum about the phenomenology of Sorites reasoning.

All these pieces came together when writing "Generic Vagueness", now retitled "Defeasible Tolerance and the Sorites" (Chapter 6). It's here where I give the positive proposal. I claim that: (i) tolerance principles are prescriptive norms of judgment, that when properly analyzed take the form of deontic conditionals; (ii) Sorites reasoning is a type of defeasible reasoning; (iii) the defeasibility in question is epistemic in nature: they concern what conclusions may be suspended pending further information or overturned in light of additional knowledge; (iv) there are actually two types of tolerance, obligatory and permissive; (v) the Asher-Bonevac-Morreau framework of nonmonotonic deontic logic suits these purposes: being nonmonotonic, it models the defeasibility of Sorites reasoning, and being deontic, it captures the implicit deontic content of tolerance principles; (vi) the proposal beats out contextualist theories of vagueness, because it better accounts for the phenomenology of vagueness; (vii) it sheds light on two problem cases in the vagueness literature, the forced march Sorites and Mark's paint shop case; (viii) a normative understanding of tolerance is compatible with the existence of cutoffs, where this simply demonstrates the division between the metaphysics vs. the epistemology of vagueness.

Material from Chapters 5 and 6 was presented to audiences at the University of Notre Dame (April 2013), Syracuse University (February 2014), Houston Baptist University (April 2014), University of Toronto (May 2014), and University of Maryland, College Park (Sixth NASSLLI, June 2014), as well as in-house on more than one occasion at The University of Texas at Austin.

Epistemicism

Ivan J Hu, Ph.D.
The University of Texas at Austin, 2015
SUPERVISORS: Richard M. Sainsbury, Johan A. Kamp

I propose a new theory of vagueness centered around the epistemology and normativity of vagueness. The theory is a version of epistemicism—the view that vagueness is a fundamentally epistemic phenomenon—that improves upon existing epistemicist accounts by accommodating both the alleged tolerance and open texture of vague predicates, while foregoing excessive metaphysical commitments. I offer a novel solution to the infamous Sorites paradox, one that outrivals alternative contextualist theories in their ability to explain the phenomenology of vagueness as well as its deontic consequences.

Table of Contents

Introduction

Must vague statements be either true or false, indeterminate or truth-valueless? Are there hidden truths about borderline cases that elude clear classification, or are there no facts of the matter for vague matters? Which of these characterizations captures how vagueness is experienced? Do vague terms have sharp cutoffs at which they stop applying? Does vagueness prevent us from knowing how far our ordinary categories and concepts extend? How far can or should we extend these? What norms govern how we reason using vague expressions? Is the range of borderline cases itself vague? Are there higher orders of vagueness? Is vagueness a phenomenon of language, within the mind, or in the world?

Addressing these issues, typically conceived as falling within the purview of philosophical logic or philosophy of language, carries considerable import not only for traditional philosophical questions within metaphysics and epistemology, the philosophy of science, and the philosophy of mind, but also for fundamental issues in linguistics, artificial intelligence, and cognitive science. Examining the normative and broadly deontic consequences of vagueness poses wide-ranging ramifications in applied ethics, rational choice theory, behavioral economics, and legal theory.

The theory known as epistemicism offers distinctive answers to these central questions surrounding the nature of vagueness. On this type of account, vagueness is fundamentally an epistemic phenomenon. This means there is an unknown fact of the matter for every vague matter. It may, for instance, be unclear whether a borderline bald should count as bald, but according to the epistemic theory, there is a fact of the matter, it is simply unknown. Epistemicists, most notably Timothy Williamson (1994), have defended the theory on account of two merits: among the entire catalogue of theories on market, it alone is able to preserve classical logic and semantics in their entirety, and it offers the most promising treatment of the phenomenon of higher-order vagueness (i.e. the existence of borderline definite cases, definite borderline cases, borderline borderline cases, etc.). Yet the epistemicist theory has the counterintuitive consequence that each vague term has its own (albeit unknown) cutoff point. One can cease to be young after one day of aging; losing a single hair can make you bald; the sky can turn dark with the disappearance of a beam of sunlight. Furthermore, existing epistemicist theories fail to account for the alleged tolerance of vague predicates (roughly, the rule that vague predicates should never apply selectively to sufficiently similar things) and claims of open texture (roughly, saying it is permissible to judge borderline cases either way). This prevents them from accommodating the full range of facts about vagueness, including its phenomenology and its ethical implications.

Epistemicism defends a novel version of the theory that meets these challenges by rethinking how vagueness guides the interaction of logic, epistemology, and phenomenology. I show how epistemicism, contrary to criticism, is able to accommodate both tolerance and open texture after all, while retaining a fully epistemic conception of vagueness. I argue that these two features of an epistemicist theory are not threatened by the existence of cutoffs.

The thesis is organized as a compilation of six essays, each independently standing, with no cross references. The first half examines traditional epistemicist accounts. The second half gives the positive proposal for a new brand of epistemicism.

Epistemicists contend that there is an unknown fact of the matter for every vague matter. Thus they defend bivalence (roughly, the principle that any meaningful claim, even if vague, is either true or false) but posit epistemic gaps within the range of borderline cases. Most theorists balk at the commitment to sharp cutoffs. However, recent work by David Barnett, Cian Dorr, and Crispin Wright has challenged the latter, seemingly innocuous half of epistemicism: what I call the UNKNOWN principle, which claims that vagueness entails ignorance. I defend the principle in "Epistemic Gaps" and "Vague Unknowns".

In "Epistemic Gaps" (Chapter 1), I reply to Barnett (2010). Barnett rejects UNKNOWN on the grounds that unknowability is not truly an *essential* feature of vagueness. To support this, he argues for the possibility of a hypothetical linguistic community in which speaker intuitions about certain vague matters could be *response-dependent* in the sense of enabling knowledge (roughly, if it seems that p then one can know that p). In response, I argue that the mere possibility of such response-dependence poses no threat to epistemicism and that any objection to UNKNOWN based on such alleged counterexamples is question begging at best. I outline how a broadly Williamsonian version of epistemicism can accommodate the possibility of response-dependence and how it has the resources to deliver the intuitively correct predictions about such cases. I also challenge the idea that our intuitions about vague matters might *actually* be response-dependent, arguing that linguistic data favor UNKNOWN in light of the phenomenology of vagueness.

It is here where I develop a general technique for reasoning about vagueness, using minimal logical assumptions (essentially, that any determinacy operator used to formulate claims of vagueness obeys a modal logic at least as strong as **KT**), and exploit it to establish a number of results demonstrating how the various notions of vagueness, higher-order vagueness, vagueness in knowledge, vagueness in intuition, and response-dependence are all interrelated. This same technique gets applied later to various issues in other debates: how vagueness should relate to ignorance (in "Vague Unknowns"), how vagueness interacts with quantification ("'Vague' and Higher-Order Vagueness"), and how vagueness impacts norms and obligations ("Ignorance and Open Texture").

In "Vague Unknowns" (Chapter 2), I reply to the other objections against UNKNOWN advanced by Wright (2001, 2003), Barnett (2009, 2010), Bobzien (2010, 2012), and Dorr (2003).

I argue that each of the substitute principles they propose in place of UNKNOWN is to be found inadequate for reasons independent of vagueness. I also argue that Williamson's (1992, 1994, 1995) own account of how vagueness relates to ignorance is unsatisfactory. I draw some substantive conclusions about the nature of principles like UNKNKOWN that attempt to articulate the connection between vagueness and ignorance, the import of such claims, and how they relate to the prospects for a minimal theory of vagueness in the sense of Greenough (2003). The arguments here draw heavily upon the logical technique developed in "Epistemic Gaps".

In "'Vague' and Higher-Order Vagueness" (Chapter 3), I defend the claim that any vague predicate is higher-order vague because the predicate 'vague' is itself vague. The original proof, based on a clever construction of Sorensen's (1985), is given and defended by Hyde (1994, 2003) and criticized by Tye (1994), Deas (1989), Hull (2005), and Varzi (2003, 2005). A glaring omission in the debate is the general failure to say exactly what the relations are between vague vagueness and higher-order vagueness. To amend this, I derive some results showing exactly how the two notions are related, then explain how this affects the overall dialectical situation.

In "Margins for Error in Meaning" (Chapter 4), I critically examine the prevailing epistemicist theory: Williamson's (1994) margin of error account. A lingering problem for the account is that Williamson's margin for error principles, used to explain our failure to know the underlying status of borderline cases, might also wrongly predict that we fail to know exactly what our vague terms mean. Sainsbury (1997) argues that such semantic knowledge can be preserved on the account, by appealing to the notion of being properly inducted into a linguistic community where competent speakers all use a given vague expression with a single, shared meaning. I first offer an undercutting defeater to this strategy, arguing that the appeal to linguistic induction fails to preclude at least one type of semantic ignorance. I then offer a rebutting defeater, by detailing some counterexamples to Williamson's assumption that being inducted into the relevant linguistic practice guarantees semantic knowledge. Finally, I show why Williamson's claim that we know the meanings of our vague terms *exactly*, turns out to be false on his own account: any semantic knowledge of the meaning of our vague terms must be inexact knowledge.

Having so far defended traditional epistemicist views on vagueness-related ignorance, semantic ignorance, higher-order knowledge, higher-order vagueness, and vague vagueness, while also examining the shortcomings of Williamson's version of the theory, I go on to explore alternative ways to develop the epistemicist theory that, by contrast, can accommodate the tolerance and open texture of vague predicates.

In "Ignorance and Open Texture" (Chapter 5), I examine the ethical or, more broadly, normative implications of vagueness. The epistemicist theory, I argue, is well positioned to address such issues. However, this requires challenging the dominant approach taken by existing versions of epistemicism: that of preserving classical logic and semantics in their entirety. Although considered orthodoxy among epistemicists, such logical conservatism is problematic if epistemicists wish to accommodate the alleged "open texture" of vague predicates. The alternative I propose aims to relax certain logical commitments and thereby account for our permissibility intuitions. Contrary to recent criticism (Shapiro 2006, Soames 1999), I argue that epistemicists *can* endorse limited claims of open texture after all. The discussion here centers around a purported counterexample to epistemicism offered by Shapiro (2006): a case of conditional obligation where the relevant conditions are vague, so that epistemicism appears to predict (wrongly) the existence of unknowable hidden obligations. I argue that there is independent reason to reject any argument against epistemicism based on such problem cases, given the failure of analogous paradoxical reasoning in the Miners Paradox (Kolodny & MacFarlane 2010, Willer 2011). Such objections to epistemicism, I argue, overlook the information-sensitive nature of conditional obligation and presuppose an erroneous logic for deontic conditionals. In fact, I shall argue, it is precisely such a logic that is needed for an adequate treatment of the tolerance of vague predicates.

In "Defeasible Tolerance and the Sorites" (Chapter 6), I present the key proposal: a novel form of epistemicism that, unlike its predecessors, is able to account for the alleged tolerance of vague predicates. The Sorites paradox concerns how repeated use of tolerance conditionals (e.g. If anyone with n hairs is bald, then anyone with n+1 hairs is bald) appears to lead to absurd conclusions (e.g. Anyone with 100,000 hairs is bald). I argue that, in order to avoid paradox, we should see the use of such tolerance conditionals as embodying certain prescriptive norms of judgment, where these exemplify a type of *defeasible* reasoning (i.e. good in many ordinary cases, but not in general), on par with rough-and-ready generic generalizations like "Birds fly". I claim that the defeasibility in question is epistemic in nature: certain conclusions may be suspended pending further information or overturned in light of additional knowledge. I show how this can be done within the Asher-Bonevac-Morreau framework of nonmonotonic deontic logic: nonmonotonic, for modelling the defeasibility of sorites reasoning (Asher & Morreau 1991, Morreau 1997); deontic for capturing the implicit deontic content of tolerance principles (Asher & Bonevac 1996, Bonevac 1998). In this way the proposal ascertains the compatibility of tolerance with the existence of cutoffs. Notably, the account also draws a crucial distinction—gone largely unrecognized in the literature—between obligation-based tolerance and permission-based tolerance. I compare it to existing contextualist treatments of vagueness (Kamp 1981; Raffman 1994, 1996; Soames 1999; Graff-Fara 2000; Shapiro 2006) and argue that it better accounts for the phenomenology of vagueness by capturing key differences between these two kinds of tolerance. I also offer some in-depth comparisons of the defeasible approach to other closely related alternatives that adopt some form of "contextualist" logic (Kamp 1981; Cobreros, Egré, Ripley, and van Rooij 2010; Kolodny & MacFarlane 2010; Willer 2012).

Chapter 1: Epistemic Gaps

Epistemicism identifies vagueness as an inherently epistemic phenomenon. On this conception, vagueness entails ignorance: if it is vague whether *p*, it is unknowable whether *p*. Epistemicists attribute this to human cognitive limitations: according to one familiar account, it is our insensitivity to the minutia of unstable use patterns governing vague expressions that prevents us from knowing how they apply to borderline cases (Williamson 1994). Yet, one may object, these limitations are only contingent. For cognitively enhanced speakers with perfectly reliable and consistent linguistic intuitions, vagueness might fail to have its familiar epistemic consequences, because their intuitions would be, in some important sense, perfectly *response-dependent*. This challenges the epistemicist idea that leaving epistemic gaps is not just an accidental but a *necessary* feature of vagueness. Worse yet, the possibility that ordinary speaker intuitions might actually be response-dependent casts doubt on whether it is even an *actual* feature of vagueness.

This paper offers a two-fold defense of epistemicism. First, I argue that response-dependence is no threat to the *actual* truth of epistemicist claims. I offer some linguistic data that favors epistemicism over its countervailing possibilities: whereas response-dependence fails to be realized among actual ordinary speakers, epistemicist principles comport surprisingly well with the phenomenology of borderline cases.

Second, I argue that response-dependence is no threat to the *necessary* truth of epistemicist claims. I offer a detailed reply to David Barnett (2010), who argues against epistemicism by counterexample. The possibility of a linguistic community exhibiting response-dependence, he claims, shows that epistemicism mischaracterizes the nature of vagueness. I explore various ways of developing this objection—some unanticipated by Barnett—and argue that each version fails: however the alleged counterexample is to be construed, either Barnett's criticisms are question begging against the epistemicist theory, or his envisaged scenario is simply incoherent.

The general lesson that will emerge is that the epistemicist need not deny, but can countenance, the possibility of response-dependence. I outline how a broadly Williamsonian version of epistemicism has the resources to deliver the intuitively correct predictions about such cases regarding both the phenomenology of vagueness as well as the epistemology of higher-order vagueness. I show how to derive these nontrivial results from basic epistemicist principles, using minimal logical assumptions, within a **KT** modal logic for vagueness. In this way, the epistemicist can explain away objections from response-dependence without compromising the ignorance-inducing nature of vagueness.

1.1 Intuitions, ambivalence, response-dependence

Vagueness is the absence of clarity. Say that it is clear that *p* just in case *p* and it is not vague whether *p*. Conversely, it is vague whether *p* just in case it is not clear that *p* and not clear that $\neg p$.[1] The epistemicist claim under question says that vagueness entails ignorance.

> **UNKNOWN**
> If it is vague whether *p*, then it is unknowable that *p* and it is unknowable that $\neg p$

UNKNOWN appears plausible.[2] In *V. for Vendetta* Natalie Portman's character has her head shaved. At the beginning of the buzz cut, she is clearly not bald; by the end, she clearly is. At

---

[1] I am following Barnett's (2009, 2010) decision to avoid the standard terminology, though one may substitute *determinately* or *definitely* for *clearly* without harm in whatever follows.

some time in between, she is borderline bald—it is vague whether or not she is bald. So is she bald then? There's no knowing, according to UNKNOWN.

Why not? A plausible explanation appeals to intuitions. Evidence for someone's baldness comes in the form of intuition—whether or not someone in a certain hair situation *looks* or *seems* bald.[3] To know whether someone is bald would require a certain kind of clarity in intuition: they need to *clearly* look or seem bald, at least if their baldness is to be assessed on the basis of looks alone. Yet the absence of such clarity in intuition is a token characteristic feature of vagueness. Intuitions about borderline cases conflict, whether across different speakers or within the same subject at different or even simultaneous times. We know that Samuel L. Jackson is bald and that Whoopi Goldberg is not. But borderline cases abound in between, whose statuses remain unsettled. Is Danny Devito bald? It's debatable. Opinions differ. Verdicts waver. Our intuitions about baldness are too discordant and unstable to decide the matter once and for all.

Vagueness triggers a feeling of *ambivalence* in intuition. Opposite intuitions clash, each unclear. That is, for either judgment about a borderline case—bald or not—it is unclear whether things appear *that* way.

---

[2] Other authors endorse weaker versions of the principle. Greenough (2003) endorses the existential claim that any vague predicate must have *an* epistemically borderline case (i.e. whose status remains unknown). Wright (2001), on one interpretation, appears to endorse the counterfactual claim that if something were clearly borderline F, its status *would* be unknowable (i.e. the weaker claim that *clear* vagueness entail ignorance)—although he does not think there exist any such cases. Williamson, surprisingly, nowhere offers any formulation or defense of any principle articulating the epistemic consequences of vagueness, except perhaps: "an interpretation *s* admits an interpretation *t* just in case if *s* were correct then speakers of the language could not know *t* to be incorrect. On this view, 'definitely' means something like 'knowably'. Just one interpretation is correct, but speakers of the language cannot know all others to be incorrect. Vagueness is an epistemic phenomenon." (Williamson 1994:164)

[3] More controversially, Barnett (2010:24-27) claims that "mundane" predicates like 'bald' are devoid of external environmental content, insofar as their extension is not determined by some hidden essence underlying appearances of baldness, but determined strictly on the basis of the appearances themselves. That would suggest that "external" facts such as the number of hairs on one's head are not part of the meaning of 'bald' and that baldness is not a real external property of things. I intend to remain neutral on such controversial issues.

AMBIVALENCE

If it is vague whether *p*, then for any ordinary subject *s* who has any relevant evidence regarding whether *p* and considers whether *p*, it is vague whether it seems to *s* that *p* and it is vague whether it seems to *s* that ¬*p*

AMBIVALENCE is plausible for many ordinary cases of vagueness.[4] For consider how hedging is good evidence of vagueness. Asking "Is Danny Devito bald?" readily invites replies like:

He *sort of* is bald, he *sort of* isn't
He *kind of* is, *kind of* isn't
It's *hard to say* whether he's bald

Witness how such hedging behavior persists when reflecting upon our own responses to the issue at hand, as when asked "Does Danny Devito seem bald *to you*?"

Well, he *sort of* does, he *sort of* doesn't
It *kind of* looks like he's bald, it *kind of* doesn't
It's *hard to say* either way

Hedged self-reports about how things appear suggest that our intuitions are themselves vague in the presence of borderline cases. Under severe ambivalence, intensifying my hedges, as in

He *really only sort of* seems bald, *really only sort of* doesn't
He *just kind of* looks bald and *just kind of* doesn't
It's *very hard to say* which way he appears

would indicate that my intuitions are not just vague, but *clearly* vague.

Ambivalence in intuition generally accompanies vagueness. Yet it prevents our knowing the true status (if any) of borderline cases—F or not F—for which it is natural to plead ignorance:

---

[4] Barnett (2009, 2010, manuscript) defends a version of AMBIVALENCE. Although I doubt it is generally true of *all* cases of vagueness, I shall set aside such worries for dialectical purposes.

Is Danny Devito bald? I don't know—it's vague

He sort of is, but he sort of isn't—I can't tell

I'm not sure—he's kind of bald, kind of not

Importantly, these are denials, not hedges, of epistemic claims. Consider the implausibility of knowledge-related hedges:

*I *sort of* know he's bald, *sort of* know he isn't

*You can *kind of* tell he's bald, *kind of* tell he's not

*I'm *kind of* certain he's bald, *kind of* certain he's not

*It's *hard to say* whether I know he's bald

*It's *hard to tell* whether I'm sure he's bald

Even worse are ignorance-related hedges:

*I *sort of* can't tell if he's bald, *sort of* can't tell if he isn't

*It's *hard to say* whether I'm uncertain about his baldness

*I *kind of* don't know that he's bald, *kind of* don't know that he isn't

It's not as if we *vaguely* don't know the status of typical borderline cases; we really *don't* know.[5]

This vindicates UNKNOWN. It also has implications for a principle of *response-dependence*.

RESPONSE-DEPENDENCE

For any ordinary subject who has some relevant evidence regarding whether p and considers whether p, if it seems to s that p then s can know that p and if it seems to s that $\neg p$ then s can know that $\neg p$

Response-dependent conditions can be known to obtain when they seem to obtain.[6] Baldness is not response-dependent, since our intuitions about baldness do not even provide the possibility of

[5] This asymmetry between denied and hedged epistemic claims distinguishes UNKNOWN from the similar but distinct principle UNCLEAR: If it is vague whether p, it is *vaguely* unknowable that p. The data, it appears, tells against UNCLEAR but supports UNKNOWN.

knowledge about borderline cases. We might otherwise have all the relevant evidence for assessing whether Danny Devito is bald. That includes facts about his total hair situation—including hair count, head shape, hair length, hair follicle distribution, hair thickness, and the like—and how these contribute to his overall appearance—looks taken from different viewpoints, up close or from afar. It may also involve facts about overall use patterns for 'bald'—including actual linguistic behavior, dispositions of individuals to assent to or dissent from various uses of 'bald', reactions of speakers, both normal and idiosyncratic, to the application of 'bald' in a wider range of cases, paradigm, borderline, or unusual, opinions of both expert theorists and lay folk alike, whether explicit or tacit, whether in consensus or disagreement, about the semantic properties of 'bald' and its inferential roles in reasoning, and the like. The availability of such evidence, no matter how complete, does nothing to sharpen our intuitions. Possessing all this information gets us no closer to determining the overall status of Danny Devito's hair situation: bald or not?

---

[6] On another understanding, "response-dependent" conditions obtain just when they seem to obtain under normal circumstances. Call this *response-dependence**. The paradigm case is color. Redness is purportedly response-dependent* insofar as: an object *x* is red if and only if for any subject s, if s is properly situated under normal viewing conditions with respect to x, then x looks red to s (see Johnston 1989, Pettit 1991).

The two notions, often misleadingly swept under the single label of "response-dependence", are connected by the notion of luminosity, where a condition is *luminous* if it can be known to obtain whenever it obtains (see Williamson 2000:95). Consider a subject s who is properly situated toward an object *x* under normal viewing conditions. Assume that looking red is luminous, such that: (i) if *x* looks red to s, then s can know that *x* looks red. Assume moreover that redness is response-dependent* and that s, who has carefully reflected over the nature of redness as well as s's own normal viewing conditions of *x*, is able to know this, such that s can know the truth of the response-dependence* conditional: if *x* looks red, then *x* is red. It follows given the closure of knowability over knowable entailment that: (ii) if s can know that *x looks* red then s can know that *x is* red. Combining (i) and (ii) yields the relevant response-dependence conditional: if *x* looks red to s, then s can know that *x* is red.

Such reasoning, if successful, might be used to support premise R1a of the response-dependent objection (below). Fortunately, luminosity is objectionable for independent reasons (see Williamson 2000 ch.4). I shall set aside the notion of response-dependence* in what follows.

The ambivalence of our actual intuitions about applying vague terms in the presence of borderline cases disqualifies their application from the ideal of response-dependence. Ordinary subjects who know perfectly well what Danny Devito looks like are still unable to know whether he's bald, since their intuitions about the matter are vague. Hence response-dependence fails: it is *not* true that if it seems to them that he's (not) bald then they can know that he's (not) bald.[7]

It would appear then that the phenomenology of vagueness speaks in favor of epistemicism. Critics of epistemicism, however, may dissent. Failure of response-dependence, they will say, is a contingent feature of actual linguistic practice. It therefore reflects nothing about the nature of vagueness. The next section looks more closely at this objection.

1.2 The argument from response-dependence

David Barnett (2010) contests UNKNOWN on grounds of the possibility of a hypothetical linguistic community where speaker intuitions are both ambivalent and response-dependent.

---

[7] More precisely: Let p mean *Danny Devito is bald* and s be some ordinary subject who knows what Danny Devito looks like. Given the impossibility of knowledge, both response-dependence conditionals—*If it seems to s that p then s can know that p* and *If it seems to s that ¬p then s can know that ¬p*—will have a false consequent. Given the ambivalence of s's intuitions over the question of Danny Devito's baldness, both antecedents will be unclear in their truth-value. This is enough to make at least one of the conditionals untrue in some circumstance, on any standard analysis of the indicative.

On non-bivalent treatments, both antecedents—*It seems to s that p* and *It seems to s that ¬p*—will either be indeterminate in truth-value (on a trivalent logic), have some intermediate truth-value (on a many-valued logic), or have no truth-value, i.e. no fixed truth-value across all admissible precisifications (on a supervaluationist logic). Thus both conditionals will themselves either be indeterminate in truth-value (on a Kleene or Lukasiewicz trivalent logic), have intermediate truth-value (on a many-valued logic), lack truth-value (on a supervaluationist logic), or simply get evaluated as false (on a Gödel trivalent logic).

On a bivalent treatment, so long as one of the antecedents is true, the corresponding conditional will be false. This has to be so in *some* circumstances, assuming that both kinds of borderline-seemings exist—cases where it seems that p but only vaguely and cases where it seems that $\neg p$ but only vaguely. Otherwise, the only way to guarantee that both conditionals turn out vacuously true is to maintain that in *all* cases of vague intuition, it neither seems that p nor seems that $\neg p$—which is implausible. I wish to set aside such "falsehood-entailing" conceptions of borderlineness, on which to be borderline F is to be *not*-F (although see Raffman 2005 for such an account).

Imagine a hypothetical linguistic community, one just like ours, except whose speakers have enhanced cognitive faculties. Like ours, their language Zenglish contains vague predicates. However, their linguistic intuitions do not suffer the sort of unreliability and inconsistency that plague ours. Application conditions for vague predicates like ‘bald*’—the Zenglish analogue of our ‘bald’—are determined entirely by the collective dispositions among Zenglish speakers concerning use. For predicates as these, things just are as they seem. Linguistic intuitions straightforwardly fix their application conditions. Since Zenglish speaker intuitions are perfectly reliable and consistent, given their cognitive superiority, their linguistic intuitions are perfectly accurate indications about the application conditions for terms like ‘bald*’—in fact, so perfect as to offer knowledge of those conditions. Any Zenglish speaker who has the intuition that ‘bald*’ applies to someone and believes this under normal conditions—she has no reason to think her intuitions are misleading, and so on—comes to know this fact.

In this way, baldness* is response-dependent for Zenglish subjects. Zenglish intuitions make the application conditions of ‘bald*’ completely knowable: under normal conditions, ‘bald*’ can be known by anyone to apply just as it seems to apply. Such response-dependence is not dependent upon the reliability of perceptual faculties. A Zenglish speaker may consider, apart from any observation of some particular individual *x*’s hair situation, whether ‘bald*’ would apply to anyone in the same hair situation as *x*, call it *h*. Indeed, she need not actually view anyone with that hair situation; reflecting upon the application of ‘bald*’ may simply be done from the armchair. Her perceptual faculties may be terribly unreliable. Even so, her intuitions will nevertheless be knowledge-conferring: if it seems to her that anyone in *h* would/wouldn’t be bald*, then anyone in *h* would/wouldn’t be bald*.

Zenglish intuitions represent an idealized version of our own intuitions, without all the instability and inconsistency. However, Barnett claims, this is not a way of imagining away the vagueness of terms like 'bald*' (contrary to Williamson 1994:232). Such predicates in Zenglish might still admit of borderline cases. Vagueness in the application conditions for 'bald*' simply corresponds to vagueness in the speaker intuitions that determine those very conditions. Some hair situations will be borderline bald*, and for any competent Zenglish speaker it will be vague whether it seems to her that 'bald*' would/wouldn't apply to these. Yet her intuitions about the matter (if any) remain knowledge-conferring. It remains true that *if* it seems to her that anyone in some such situation would/wouldn't be bald*, then she can know this. Response-dependence makes a claim about knowability that is conditional upon the presence of certain intuitions. It is simply vague whether these conditions are satisfied in borderline cases. But then it follows (as we shall see, on certain minimal assumptions) that whether things are knowable is itself vague. It is hence no longer clear that the status of borderline cases of 'bald*', although vague, must be unknown.

It follows that UNKNOWN, however extensionally adequate, is not clearly true in all worlds. For the clear truth of UNKNOWN would predict that clear vagueness entails a gap, not vagueness, in knowability. Zengland, however, is a world where vagueness is shielded from its familiar epistemic consequences. The possibility of Zengland is at odds with epistemicism or any other theory of vagueness claiming that vagueness *by its very nature* leaves a trail of epistemic gaps—whatever the world, community, or language.

Hence, the inconsistent triad: AMBIVALENCE, UNKNOWN, RESPONSE-DEPENDENCE. Among these, according to Barnett, AMBIVALENCE is the only necessary truth about vagueness. AMBIVALENCE and UNKNOWN, but not RESPONSE-DEPENDENCE, correctly describe the situation of vague terms in English. Whereas AMBIVALENCE and RESPONSE-DEPENDENCE, but not UNKNOWN, correctly describe the situation of vague terms in Zenglish. Both UNKNOWN and RESPONSE-DEPENDENCE are contingent truths about vagueness: mere reflections of cognitive circumstance.

I have deliberately left the details of the Zenglish scenario *roughly* fleshed out for now. We shall see how there exist more than one way to spell out Barnett's objection, depending on how these details are filled in. Give minimal assumptions, Zengland poses a *weak* counterexample contesting the necessary clear truth of UNKNOWN. Given additional assumptions, it poses a *strong* counterexample purporting to show UNKNOWN to be clearly false. We begin with the weaker variant, call it the **response-dependence argument**. In schematic form:

R1. Possibly, for some p:

   a. it is clearly vague whether p

   b. linguistic intuitions are clearly ambivalent over p, and

   c. whether p is a clearly response-dependent matter.

R2. Necessarily, if intuitions about clearly response-dependent p are clearly ambivalent, then it is vague whether p can be known and it is vague whether $\neg p$ can be known.

R3. Therefore, possibly for some p, it is clearly vague whether p, it is vague whether p can be known, and it is vague whether $\neg p$ can be known; so UNKNOWN is not necessarily clearly true.

I shall eventually argue that the joint truth of claims R1a–R1c is untenable. Before examining the argument, however, let us introduce some notation for convenience.

Let Cp mean it is clear that p, Ip mean it is vague whether p (such that $\mathrm{I}p =_{\mathrm{def}} \neg \mathrm{C}p \,\&\, \neg \mathrm{C}\neg p$), N$p$ mean it seems to one that p, and Kp mean one can know that p. We assume that both the clarity operator C and the knowability operator K obey the modal principles **K** and **T**.[8]

$\mathbf{K_C}$ $\quad \mathrm{C}(p \rightarrow q) \rightarrow (\mathrm{C}p \rightarrow \mathrm{C}q)$

$\mathbf{T_C}$ $\quad \mathrm{C}p \rightarrow p$

These say that clarity is both closed over implication and factive. These are not unreasonable assumptions to make, given that epistemicists already identify vagueness as a source of epistemic uncertainty and hold relatively uncontroversial commitments to a **KT** epistemic logic.[9] Since my aim is to defend a broadly Williamsonian version of epistemicism, I shall assume classical logic throughout.

A couple of immediate results are in order.

Result 1: The response-dependence argument is validated by basic principles of modal propositional logic together with the **KT** assumptions governing clarity and necessity.

> *Proof.* Assume premises R1 and R2 are true. Suppose p is the verifying instance of R1. By **K**, R1 and R2 jointly entail $\Diamond(\mathrm{CI}p \,\&\, \mathrm{IK}p \,\&\, \mathrm{IK}\neg p)$. Assume for reductio that UNKNOWN is necessarily clearly true (i.e. conclusion R3 is false). Instantiating with p yields $\Box\mathrm{C}(\mathrm{I}p \rightarrow (\neg\mathrm{K}p \,\&\neg\mathrm{K}\neg p))$. By **K** and $\mathbf{K_C}$, we get $\Box(\mathrm{CI}p \rightarrow \mathrm{C}(\neg\mathrm{K}p \,\&\neg\mathrm{K}\neg p))$. Given $\mathbf{K_C}$ and &-introduction, C distributes over conjunction, so $\Box(\mathrm{CI}p \rightarrow (\mathrm{C}\neg\mathrm{K}p \,\&\, \mathrm{C}\neg\mathrm{K}\neg p))$; by **K**, **N** and propositional logic, this is equivalent to $\Box\neg(\mathrm{CI}p \,\&\, (\neg\mathrm{C}\neg\mathrm{K}p \vee \neg\mathrm{C}\neg\mathrm{K}\neg p))$—which contradicts the initial possibility claim.

---

[8] Williamson's (1994 ch. 8) epistemic logic of clarity is KTB. Bacon (2011) argues against the B(rouwerian) principle. The defense of epistemicism offered here is free of any commitment to the controversial B principle.

[9] For knowability: $\mathbf{K_E}$ $\quad \mathrm{K}(p \rightarrow q) \rightarrow (\mathrm{K}p \rightarrow \mathrm{K}q)$
$\mathbf{T_E}$ $\quad \mathrm{K}p \rightarrow p$

Result 2: Premise R2 can be shown using contraposition and the closure of clarity (principle $\mathbf{K_C}$)

> *Proof.* Assume that intuitions over clearly response-dependent *p* are clearly ambivalent. By $\mathbf{K_C}$, contraposing the response-dependence conditionals within the scope of the clarity operator in $C(Np \rightarrow Kp)$ and $C(N\neg p \rightarrow K\neg p)$, then distributing C, yields $C\neg Kp \rightarrow C\neg Np$ and $C\neg K\neg p \rightarrow C\neg N\neg p$; contraposing again gives: (i) $\neg C\neg Np \rightarrow \neg C\neg Kp$ and (ii) $\neg C\neg N\neg p \rightarrow \neg C\neg K\neg p$. By $\mathbf{T_E}$, both $C(Kp \rightarrow \neg K\neg p)$ and $C(K\neg p \rightarrow \neg Kp)$ are true. By $\mathbf{K_C}$, it follows from the clear truth of the response-dependence conditionals that $C(Np \rightarrow \neg K\neg p)$ and $C(N\neg p \rightarrow \neg Kp)$. Applying to these the same sequence using $\mathbf{K_C}$ and contraposition delivers: (iii) $\neg C\neg Np \rightarrow \neg CK\neg p$ and (iv) $\neg C\neg N\neg p \rightarrow \neg CKp$. Results (i)–(iv) reduce to: $(\neg C\neg Np \ \& \neg C\neg N\neg p) \rightarrow (IKp \ \& \ IK\neg p)$. Strengthening the antecedent by $\neg CNp$ and $\neg CN\neg p$ allows us to state things more compactly: $(INp \ \& \ IN\neg p) \rightarrow (IKp \ \& \ IK\neg p)$. By the standard rule of Necessitation **N**, we can strengthen this to be necessarily true.

Since minimal assumptions suffice to secure validity and premise R2, the response-dependence objection, thus formulated, arguably hinges on premise R1—the possibility claim. It is worth examining in detail how Barnett argues for its truth.

1.3 Back to Zengland

How exactly is Barnett's hypothetical Zengland supposed to illustrate the compossibility of vagueness, response-dependence and ambivalence? The matter deserves closer scrutiny.

1. Zenglish intuitions vs. actual intuitions. Zenglish speakers' intuitions are supposed to straightforwardly determine the meanings of predicates like 'bald*', in a way that our intuitions do not when it comes to determining the meanings of vague English predicates like 'bald'. Physiological differences account for why Zenglish intuitions exhibit a level of reliability and consistency ours do not. Consider a Zenglish speaker considering under normal conditions whether 'bald*' applies to some individual *x*—that is, whether *x* is bald*, assuming 'bald*' means *bald**. His intuitions are *reliable* insofar as if it were to seem to him that *x* is/isn't bald* then *x* would/wouldn't be bald*. They are *consistent* insofar as if it were to seem to him that *x* is/isn't bald* then it wouldn't seem to him that *x* isn't/is bald*.

These are *robust* features of his intuitions, which remain reliable and consistent even under changes in meaning. The requisite notions may be formulated metalinguistically. They are *robustly reliable* insofar as if it were to seem to him that 'bald*' does/doesn't apply to *x* then 'bald*' does/doesn't apply to *x*. They are *robustly consistent* insofar as if it were to seem to him that 'bald*' does/doesn't apply to *x* then it wouldn't seem to him that 'bald*' doesn't/does apply to *x*. Barnett therefore grants that Zenglish intuitions clearly demonstrate RESPONSE-DEPENDENCE. So it is clearly true of vague terms like 'bald*' that if it seems to one that 'bald*' does/doesn't apply to *x* then 'bald*' does/doesn't apply to *x*.

In contrast, baldness is not a clearly response-dependent matter. For no single English speaker's intuitions are reliable or consistent about the application of 'bald', nevermind robustly so. So no actual ordinary individual's intuitions are clearly knowledge-conferring in the way a Zenglishman's are.

2. Zenglish intuitions and actual use. In actual practice, what plausibly determines the meaning of 'bald' is not how a single individual uses that term but rather community-wide patterns of use. Meaning supervenes on overall, not individual, use. Such meaning-use supervenience is reflected by the perfect response-dependence exhibited by terms in Zenglish. Just as actual overall patterns of use determine the meaning of 'bald', so do Zenglish intuitions about application determine the meaning of 'bald*'. Individual intuition in Zengland is able to mimic what collective use realizes in actual linguistic practice. This is because intuitions do not vary across Zenglish speakers. Given their physiological differences, they are able to achieve perfect consensus in their linguistic intuitions, resulting in Borg-like uniformity: one man's intuition is every man's.

'Bald' and 'bald*' converge in application conditions. The relevant differences between Zenglish and English hold rather at the level of the supervenience base for application, whether in use or in intuition. Thus, it would seem to any Zenglish speaker under normal conditions that 'bald*' does/doesn't apply to anyone in hair condition *h* if and only if (roughly) overall patterns of actual use would determine that 'bald' does/doesn't apply to anyone in *h*. This is no grounds for suspecting that collective use can after all achieve what individual intuition cannot. Although it may be unknown (or unclear) exactly *how* actual use determines meaning, it is nevertheless known (and clear) *that* actual use determines meaning. Accordingly, it is known (and it is clear) *that* Zenglish intuitions determine meaning; exactly *how* is (for all that's been said) another matter.

Similarly for the determination of extension. Although overall use facts suffice to completely determine application conditions, whether 'bald' applies to certain borderline hair situations may be inscrutable if the totality of our actual use patterns is humanly unknowable—perhaps because overall use dispositions over borderline cases are themselves vague (and hence unknowable). Likewise, although any single Zenglish speaker's totality of dispositional intuitions about using 'bald*' is enough to completely determine the application conditions for 'bald*', she won't clearly know the status of borderline hair situations if she doesn't clearly know the state of her own intuitions about those cases. Even Zenglish intuitions—despite their response-dependence—

are ambivalent over borderline cases. We may therefore suppose that Zenglish intuitions clearly demonstrate AMBIVALENCE.[10]

3. Zenglish intuitions vs. actual use. Zenglish intuitions might therefore be seen as encapsulating actual community-wide patterns of use, where each individual Zenglish speaker is, so to speak, the embodiment of an entire linguistic community. Yet differences still exist. Plausibly, even if we were somehow able to know all the actual facts about our overall patterns of use that determine the meaning of 'bald', we would still be unable to know the application conditions for 'bald' completely, because exactly *how* use determines meaning and extension remains inscrutable.[11]

'Bald*' is importantly different. Suppose Zenglish speakers were somehow able to learn exactly what their 'bald*' intuitions are even in borderline cases. (Perhaps they undergo further physiological upgrades of the sort that made their intuitions robustly reliable and consistent in the first place.) Then they *would* thereby know what the exact application conditions are for 'bald*'. For there is nothing unclear or unknown about how these would be determined by Zenglish intuitions. Given perfect response-dependence, such determination is straightforward: if it seems to one (under normal conditions) that anyone in *h* is/isn't bald* then anyone in *h* is/isn't bald*.

The relevant difference is arguably a difference in stability. Overall use patterns for vague English terms are highly sensitive to slight changes in individual use patterns. If just a few individuals used 'bald' slightly differently, the overall pattern of how 'bald' is used—and hence its meaning and application conditions—could be slightly different. Such slight changes are easy to come by, given the rampant unreliability and inconsistency of our intuitions about 'bald'. The

---

[10] We may at least suppose this for now. I shall later give reasons why we should not think Zenglish intuitions are *clearly* ambivalent, given that they are already clearly response-dependent. See §5.

[11] See for instance Williamson 1994 §8.4.

same cannot be said for ‘bald*’. Granted, Zenglish intuitions are in one sense highly sensitive to slight changes in individual use patterns, insofar as if just one individual in the community were to think that ‘bald*’ had slightly different application conditions, so would they all. But local changes in individual use patterns are hard to come by, for the very reason that intuitions of Zenglish speakers for ‘bald*’ always remain in perfect global alignment, given their robust reliability and consistency.

4. Zenglish intuitions vs. Zenglish use. The stability of Zenglish intuitions singles them out as the true supervenience base for the meanings of vague Zenglish terms like ‘bald*’. It is speaker intuitions, not overall use patterns, that determine the application conditions for ‘bald*’.

To see why this is so, consider what happens if, following Barnett, we assume the principle of excluded middle. Given excluded middle, there will be (assuming, for simplicity, that baldness* supervenes on hair number) a cutoff in hair number for ‘bald*’—some n such that anyone with n hairs is bald* but anyone with n+1 hairs is not bald*. Barnett defends the idea that excluded middle is compatible with vagueness: expressions may still be vague, despite having cutoffs, because it will nonetheless be vague where those cutoffs are.[12] On this picture, ‘bald*’ is still vague despite having a cutoff, since it is simply vague where that cutoff is. More generally, every borderline case for ‘bald*’ turns out to have some underlying status—either he is bald* or he is not bald*—though it is simply vague which he is. The guarantee of a fact of the matter is no threat to vagueness, if the facts themselves are unclear.

The cutoff for ‘bald*’—whose existence is guaranteed by excluded middle—cannot be fixed by use. For suppose, moreover, that Zenglish speakers are robustly guided entirely by their

[12] Assuming excluded middle is therefore dialectially harmless. See Barnett 2009:§2.1.1, 2010:§3. See also §6 below.

intuitions when using ‘bald*’: Zenglish speakers are disposed to judge that an individual *x* is/isn’t bald*’ only insofar as it seems to them under normal conditions that *x* is/isn’t bald*. In that case, use patterns must be at least as stable as the intuitions they supervene on.[13] If the underlying intuitions show robust reliability and consistency, then so must the ensuing use patterns. Individual use patterns for ‘bald*’ will then be uniform across the entire Zenglish speaking community. This leaves no room for discrepancies in use, even over borderline cases. Rather, when confronted with a borderline case, Zenglish speakers will either uniformly hedge or uniformly refrain from giving any verdict. Hence patterns of use will fail to specify a cutoff for ‘bald*’.[14]

Given how the use of ‘bald*’ is constitutively dependent upon Zenglish speaker intuitions, it must be shared linguistic intuitions that fix its application conditions. By extension, it must be intuition, rather than use, that fixes a cutoff for ‘bald*’. Yet, whatever intuitions are responsible for fixing the cutoff for ‘bald*’, these can’t be used by Zenglish speakers themselves to locate exactly where it is. Not all changes in intuition—namely, switches from it seeming that anyone with *n* hairs is bald* to it seeming that anyone with *n*+1 hairs is not bald*—are epistemically accessible. This needn’t jeopardize the alleged response-dependence of cutoff intuitions. For they

---

[13] Otherwise, someone could easily apply ‘bald*’ differently than they actually do (given the instability of his use patterns) while still sharing the same intuitions (given the stability of his intuitions), in which case he would cease to let his judgments be entirely guided by intuition—contrary to the *robust* reliance of use upon intuition.

[14] One may question if the hedging or refraining behavior of Zenglish speakers does not in fact suffice to fix cutoffs. Perhaps they hedge or refrain just slightly less before the cutoff, just slightly more after it? To think that use patterns might fix a cutoff for ‘bald*’ in this way is misguided. How a Zenglish speaker uses ‘bald*’ will be determined entirely by his own intuitions about its application, for his intuitions are the same as everyone else’s and equally stable. If he hedges or refrains slightly less at *n* but hedges or refrains slightly more at *n*+1, it is because he recognizes that his intuition that ‘bald*’ applies is slightly stronger at *n* and slightly weaker at *n*+1. But such a cutoff in intuition, by assumption, is not clearly scrutable. For though cutoff intuitions exist, it remains vague where exactly they are (just as: though cutoffs exist, it remains vague where exactly they are).

needn't themselves be clear or known in order to be clearly knowledge-conferring. Intuition may illuminate without being "luminous" (in the sense of Williamson 2000).

Consider a thermometer in a pool. It is carefully engineered so as to reliably and consistently indicate the temperature of a pool. The pool is always as hot or cold as reported, it tracks changes in temperature, it never gives mixed readings, and so on. Yet the thermometer is covered by algae. So its readings are obscured from plain view. They are, in a word, *unclear*. Zenglish intuitions are like thermometers covered in algae: they are good indicators of what is or isn't bald*, but their indications, however accurate, are not always clear. Their epistemic inaccessibility in borderline cases may be due to cognitive limitations. For Zenglish speakers may be cognitively advanced just enough to possess perfect faculties of intuition, but not so advanced as to possess perfect faculties of introspection.

In this way, Zenglish speakers are not *clearly knowledgeable* about the full application conditions for their vague terms, including whether 'bald*' applies to borderline hair situations. Yet neither are Zenglish speakers *clearly ignorant* of the status of such borderline cases. Rather, it is *vague* whether any competent Zenglish speaker who considers the matter knows the truth. That so much is a consequence of the response-dependence argument for *clearly* borderline cases of 'bald*'. This challenges the epistemicist claim to the necessary clear truth of UNKNOWN.

1.4 Reply to Barnett

Recall the possibility premise R1 of the response-dependence argument, which says it is possible for some clearly ambivalent intuitions over some clearly vague matter to exhibit clear response-dependence. Does Barnett's Zengland scenario really illustrate this possibility? I am doubtful. Below, I offer two arguments for why there can be no clear cases of vagueness in a linguistic community like Zengland. The core idea underlying them both is that the nature of Zenglish intuitions precludes Zenglish terms from being clearly vague—a fact that *is* accounted for after all by the epistemicist theory. If correct, this undercuts any attempt to use Zenglish or Zenglish-like examples to vindicate premise R1 of the response-dependence argument.[15]

Argument 1: precisifications as eligible candidates. Zenglish intuitions, as we saw, are perfectly stable. Given such uniformity in intuition across individuals, slight changes in individual use patterns for 'bald*', and hence its meaning, are hard to come by. But if 'bald*' could not have easily meant something different, it is hard to see how 'bald*' could, in some important sense, admit of various precisifications—without which, there arguably is no vagueness (at least of the sort Barnett claims exists in Zenglish). To count as clearly admitting of vagueness, a predicate must clearly have multiple admissible precisifications—which 'bald*' does not. As such, it is unclear whether Zenglish terms can be vague. That is, it is unclear whether borderline cases really exist for 'bald*'.

The argument just given—admittedly Williamsonian in nature—may be summarized roughly as follows.

V1. Any vague predicate T must admit of multiple precisifications.

V2. Having multiple precisifications requires meaning instability: it must be easy for the meaning of T to have been different.

V3. Meaning instability requires use instability: it must be easy for the linguistic use patterns for T to have been different.

---

[15] It also demonstrates how an epistemicist such as Williamson can dodge Barnett's objection, without having to—as Barnett (2010:27) thinks Williamson *must*—deny the very possibility of Zenglish-like scenarios.

V4. Use instability requires intuition instability: it must be easy for speaker intuitions about the application of T to have been different.

V5. Therefore, vagueness requires intuition instability.

Assuming premises V1–V4 are all clear, so is the conclusion V5, which by $\mathbf{K}_\mathbf{C}$ then entails that clear vagueness requires clear intuition instability. But intuitions for 'bald*' are stable; at least they are not clearly unstable. Hence bald*' is not clearly vague.[16]

Argument 2: precisifications as local use patterns. On another understanding, to precisify a term is to adopt a local use pattern for that term, where this does not require seriously altering or otherwise disrupting the overall global use patterns (which perhaps are what determine the range of admissible precisifications for the term in the first place). Local use changes may be conventionalized for practical purposes, as in the court of law: the legal definition of 'adult' applies to anyone eighteen years of age or older. Yet no such understanding of precisification is readily available for the alleged vague terms of Zenglish. What would be the analogue of adopting a local use pattern for intuitions? The very idea that one may be allowed to hold intuitions about a certain matter within a restricted local context is psychologically dubious. Intuitions are not the sort of thing within one's power to change at will. Even if they were, the conventionalization of such an act serves no clear function. How should the permissibility, say, of

[16] A full defense of V1–V4 goes beyond the scope of this paper, as these all make substantive (indeed, distinctively Williamsonian) assumptions about the nature of vagueness—e.g. that meaning determines extension and that meaning is highly sensitive to changes in overall use patterns (see Williamson 1994). For discussion of the notion of "easy possibilities" at work, see Sainsbury (1997). At any rate, I shall later offer independent reason for doubting the possibility of clear ambivalence—and hence of clear vagueness—in Zenglish (see Result 5).

intuiting that 'adult' applies to just all those above a certain threshold in age or maturity serve a clear legal purpose?

Summarized:

V1.* Any vague predicate T must admit of multiple precisifications.

V2*. Having multiple precisifications requires meaning flexibility: it must be easy for the meaning of T to adapt to a restricted local context of use.

V3*. Meaning flexibility requires use flexibility: it must be easy for the linguistic use patterns for T to adapt to a restricted local context of use.

V4*. Use flexibility requires intuition flexibility: it must be easy for speaker intuitions about the application of T to adapt to a restricted local context of use.

V5*. Therefore, vagueness requires intuition instability.

Assuming premises V1*–V4* are all clear, so is the conclusion V5*, which by $\mathbf{K}_C$ then entails that clear vagueness requires clear intuition flexibility—which fails for 'bald*'.

Just as the conception of precisifications as eligible candidates is not available to the opponent of epistemicism who wishes to establish the clear vagueness of Zenglish terms, neither is this alternative conception of precisifications as local use patterns. Whether any substitute conception of precisification will do is unclear—Barnett at least has provided no reason for believing so. The burden of proof falls upon objectors to provide a compelling story for how 'bald*' can admit of various precisifications after all. Absent such an account, we have no reason to believe 'bald*' is clearly vague.

The nature of Zenglish intuitions thus induces a special sort of higher-order vagueness for Zenglish terms. I am not denying here that Zenglish terms can be vague. I am only denying that they can be clearly vague. Either they are *not* vague, or they are only *vaguely* vague, if vague at all. But then the threat to epistemicism vanishes. For the assumption that things are *clearly* vague is essential to the response-dependence argument.[17] Merely saying *p* is vague simpliciter entails, according to epistemicism, that it is unknowable whether *p*. This is compatible with saying it is vaguely knowable whether *p*. For then it is unknowable whether *p*, but only unknowably so.[18]

Furthermore, it can be shown on epistemicist principles that whatever is vaguely vague is vaguely knowable. Consider:

> **KNOWN**
> It is clear whether *p* only if for any subject *s* who is suitably situated with respect to some relevant, non-trivial evidence for assessing whether *p*, where vagueness is the only potential source of ignorance preventing *s* from knowing whether *p*: if *p* were true then *s* could know that *p* and if *p* were false then *s* could know that $\neg p$

A competent Zenglish speaker is "suitably situated" with respect to whether *x* is bald* when she is positioned to see (or envision, if deliberating from the armchair) *x*'s hair situation under normal viewing conditions (no occlusions, illusions, or the like). She need possess only *some*, rather than *every*, available piece of perceptual (or intuition-based) evidence to be in a position to assess the hair situation. There may be multiple viewpoints of *x*'s hair, each independently sufficient to warrant judging that *p* (either that *x is bald** or that *x is not bald**), but clarity of baldness* requires only that seeing things from one of these be enough to warrant a judgment that *p*. The non-triviality constraint ensures that one's knowing *p* (or anything else entailing *p*) does not itself

---

[17] Thus, response-dependence is unlikely to pose any threat to epistemicists like Bobzien (2010) who already deny the existence of any *clearly* borderline cases. However, I find such accounts independently implausible, give that we often appear to be clearly ignorant over the status of borderline cases (see §1)—contrary to prediction (see Result 3).

[18] There exist independent reasons to think vagueness does not in general preclude truth or falsity (see Wright 2003).

serve as evidence for *p*.[19] An intuition that *p* counts as "evidence" in this sense; knowledge that *p* does not.[20] The restriction to only possibilities of vagueness-related ignorance serves to rule out extraneous, unintended obstacles to knowledge that are independent of vagueness (future contingents, Gettier cases, quantum indeterminacy, and so on).

Together, UNKNOWN and KNOWN ensure the proper coordination between vagueness and ignorance—and also between clarity and knowability under epistemically favorable conditions—as required by the epistemicist theory of vagueness. These are not meant to aspire toward any reductive treatment of vagueness. Open to further refinement, they simply serve as a first-pass articulation of its epistemic consequences—both the obstacles to knowledge it presents and the open epistemic possibilities in the absence of vagueness. These links between clarity and its epistemic ramifications persist through higher orders of vagueness. In particular:

Result 3: Given the clear truth of UNKNOWN and KNOWN, vague vagueness entails vague knowledge.

> *Proof.* Assume that UNKNOWN and KNOWN are both clearly true. The clear truth of UNKNOWN predicts $C(Ip \rightarrow \neg(Kp \vee K\neg p))$. Repeated applications of contraposition and $\mathbf{K_C}$ yield $C((Kp \vee K\neg p) \rightarrow \neg Ip)$; so $C(Kp \vee K\neg p) \rightarrow C\neg Ip$; so then $\neg C\neg Ip \rightarrow \neg C(Kp \vee K\neg p)$. Distributing unclarity over disjunction gives (i) $\neg C\neg Ip \rightarrow (\neg CKp \;\&\; \neg CK\neg p)$.[21] Next, the clear truth of

[19] I leave it an open issue as to whether another's testimony that *p* should count as non-trivial evidence that *p*. When considering whether *p*, less attuned subjects may defer to more experienced judgers whose intuitions are more refined, trusting their judgment and believing likewise, and thereby come to know whether *p* on the latter's say-so, rather than their own intuitions. This way of gaining otherwise unavailable knowledge by another's testimony, at any rate, is not possible for Zenglish speakers, since by assumption their intuitions are perfectly in sync.

[20] E = K theorists, along with others who reject this assumption, may substitute some knowledge-neutral notion (e.g. *grounds*) for "evidence" in the formulation of KNOWN.

[21] *Proof.* Suppose $\neg C(A\vee B)$. Assume $\neg(\neg CA \& \neg CB)$ for reductio, which is $\neg\neg CA \vee \neg\neg CB$ by one of de Morgan's equivalences. Assume $\neg\neg CA$. Then CA by double negation elimination; hence $C(A\vee B)$ by ∨-

> KNOWN predicts both $C(Cp \rightarrow Kp)$ and $C(C\neg p \rightarrow K\neg p)$. By repeated applications of contraposition and $\mathbf{K_C}$ these become (ii) $\neg C\neg Cp \rightarrow \neg C\neg Kp$ and (iii) $\neg C\neg C\neg p \rightarrow \neg C\neg K\neg p$. Combining (i)–(iii) gives $IIp \rightarrow (IKp \ \& \ IK\neg p)$.

This predicts that for any competent Zenglish speaker considering a borderline borderline case *h* of 'bald*', it is both *vaguely* known whether *h* is bald* and *vaguely* known whether *h* is not bald*. Consequently, the envisaged possibility of vague knowledge is not only compatible with, but is even accounted for, by the very theory it was designed to refute. The alleged vagueness of 'bald*' turns out to be a mere disguised case of higher-order vagueness. This salvages the necessary clear truth of UNKNOWN. Epistemicism is able to countenance the possibility of perfect response-dependence after all.

1.5 The revenge argument from response-dependence

I have argued that the assumed existence of clear borderline cases is untenable for Zenglish predicates like 'bald*'. Simply dropping the premise of clear vagueness, however, is not enough to disarm the response-dependence argument. For this assumption can be gotten back by strengthening the other premises of the argument. A revenge argument lurks.

R1*. Possibly, for some *p*:

a. linguistic intuitions are *clearly clearly* ambivalent over *p*, and

b. whether *p* is a *clearly clearly* response-dependent matter.

---

introduction and $\mathbf{K_C}$. Likewise, $C(A \vee B)$ is derivable from an assumption that $\neg\neg CB$. $C(A \vee B)$ follows by proof by cases. Contradiction.

R2*. Necessarily, if intuitions about *clearly clearly* response-dependent *p* are *clearly clearly* ambivalent, then:

a. it is *clearly* vague whether *p*

b. it is *clearly* vague whether *p* can be known, and

c. it is *clearly* vague whether ¬*p* can be known.

R3*. Therefore, possibly for some *p*, it is *clearly* vague whether *p*, it is *clearly* vague whether *p* can be known, and it is *clearly* vague whether ¬*p* can be known; so UNKNOWN is not necessarily clearly true.

The same principles used to validate the original response-dependence argument R1–R3 validate the **revenge response-dependence argument** R1*–R3*. The key difference is that clear vagueness (R2*a) is now not assumed but derived.

Result 4: Premise R2* can be shown using contraposition and principle $\mathbf{K_C}$

> *Proof of R2*.* Assume that intuitions over clearly response-dependent *p* are clearly ambivalent. The clear factivity of intuitions, expressed by $C(Nr \rightarrow r)$ and $C(N\neg r \rightarrow \neg r)$, follows from $\mathbf{T_E}$ and $\mathbf{K_C}$. Applying again the familiar sequence using $\mathbf{K_C}$ and contraposition gets us: (i) $\neg C\neg Nr \rightarrow \neg C\neg r$ and (ii) $\neg C\neg N\neg r \rightarrow \neg Cr$. These reduce to $(\neg C\neg Nr \& \neg C\neg N\neg r) \rightarrow (\neg C\neg r \& \neg Cr)$. Simplifying the consequent and then strengthening the antecedent by $\neg CNr$ and $\neg CN\neg r$ yields the more compact $(INr \& IN\neg r) \rightarrow Ir$, which can be turned into a strict conditional by the rule of Necessitation. This secures R2*a. The others R2*b–c are easily gotten by repeating the proof for R2, strengthening the premises by C, then distributing C via $\mathbf{K_C}$ to the conclusion.

Let us say states Φ and Ψ are *correlated* just in case $\Phi p \rightarrow \Psi p$ and $\Phi\neg p \rightarrow \Psi\neg p$. The revenge response-dependence argument then exemplifies this general pattern of reasoning:

CORRELATEcc

Given that being in state Φ is clearly correlated with being in clearly factive state Ψ, then clear vagueness in Φ entails both clear vagueness in Ψ and clear vagueness simpliciter: $\{C(\Psi p \rightarrow p), CC(\Phi p \rightarrow \Psi p), CC(\Phi\neg p \rightarrow \Psi\neg p), CI\Phi p, CI\Phi\neg p\} \vdash CI\Psi p, CI\Psi\neg p, CIp$.

Generalizing from R1*–R3* confirms that any notion of vagueness obeying $\mathbf{K_C}$ and $\mathbf{T_C}$ will verify CORRELATEcc. The same considerations motivating the original possibility premise R1 involving CI*p* appear to also motivate the truncated possibility premise R1* without CI*p*.

The revenge argument relocates the tension in the possibility premise as standing between *clearly clear* response-dependence and *clearly clear* ambivalence in Zenglish intuitions. The epistemicist is forced to choose between the two. Independently plausible principles suggest dispensing with (clearly) clear ambivalence. Consider:

GROUNDS

For any subject *s* who considers whether *p* on the basis of some relevant non-testimonial evidence, where vagueness is the only potential source of ignorance preventing *s* from knowing whether *p*: if *s* can know that *p* then it seems to *s* that *p* and if *s* can know that ¬*p* then it seems to *s* that ¬*p*

Vague matters in Zenglish verify GROUNDS. Speaker intuitions are the only relevant source of evidence when considering whether an individual is bald* on the basis of looks or imagination alone (i.e. apart from another's testimony). Obviously, not all knowledge is accompanied by intuition. One can know that two Müller-Lyer lines are equal in length without any perceptual appearance of them being so. One can reproduce proofs for complex mathematical theorems by rote, without any corresponding intellectual seemings. Yet competence with vague expressions requires that linguistic use be backed by intuition. One is only in a position to know based on looks alone that someone is bald* provided she has the relevant knowledge-grounding intuitions.

Result 5: Given the clear truth of RESPONSE-DEPENDENCE, UNKNOWN and GROUNDS, clear ambivalence is impossible.

> *Proof.* By the clear truth of RESPONSE-DEPENDENCE and GROUNDS, intuition and knowability are clearly equivalent. Since vagueness is closed over clear equivalence[22], we have $INp \leftrightarrow IKp$ and $IN\neg p \leftrightarrow IK\neg p$. Assume that the truth of RESPONSE-DEPENDENCE and GROUNDS are not just clear but *clearly clear*. This strengthens the biconditionals to being *clearly* true. It follows that clear ambivalence ($CINp$ & $CIN\neg p$) implies clear vagueness in knowledge ($CIKp$ & $CIK\neg p$), from which two mutually incompatible consequences follow. One is CIp, given $\mathbf{T_E}$ (the factivity of K) and CORRELATEcc, from which it follows that $C\neg Cp$ by $\mathbf{K_C}$.[23] The other is $CIKp$, by &-elimination; or IKp, by $\mathbf{T_C}$. But IKp implies $\neg C\neg Kp$. Note that the clear truth of UNKNOWN predicts $C(Kp \rightarrow Cp)$, by $\mathbf{K_C}$ and contraposition. It follows that $\neg C\neg Cp$, since vague "possibility" is closed over clear implication.[24] Contradiction.

It turns out that Zenglish intuitions cannot be *clearly* (hence, nor *clearly clearly*) ambivalent. This should be unsurprising, provided that the possibility of clear first-order vagueness in Zenglish was already dismissed, since vagueness in matters of baldness* merely reflect what vagueness is already there in speaker intuitions. The impossibility of clearly borderline cases corresponds to the impossibility of clearly ambivalent intuitions.

---

[22] **Closure of vagueness over clear equivalence.** *Proof.* Assume A and B are clearly equivalent. By $\mathbf{K_C}$ and contraposition: $C(A\rightarrow B)$ implies $C(\neg B\rightarrow\neg A)$; so $C\neg B\rightarrow C\neg A$; so $\neg C\neg A\rightarrow\neg C\neg B$. By the same principles, distributing over $C(B\rightarrow A)$, then contraposing, delivers $\neg CA\rightarrow\neg CB$. Combining yields $IA \rightarrow IB$, as desired.

[23] The substitution instance is: $C(\Psi p \rightarrow p)$, $CC(Np \rightarrow Kp)$, $CC(N\neg p \rightarrow K\neg p)$, $CINp$, $CIN\neg p \vdash CIKp$, $CIK\neg p$, CIp.

[24] That is, $\neg C\neg A\rightarrow\neg C\neg B$ given $C(A\rightarrow B)$. See the proof above (n.22) for the closure of vagueness over clear equivalence.

1.6 The strengthened argument from response-dependence

One final variant of the response-dependence argument must be considered. Barnett (2010:§4) develops the Zengland scenario into a "strong counterexample" against UNKNOWN:

> "Given LEM and that it is now vague whether Harry is bald*, clearly, either Harry is bald* or Harry is not bald*, even though it is vague which. Clearly, if Harry is bald*, then it seems to [some competent Zenglish speaker] Sophie that Harry is bald*. Clearly, if Harry is not bald*, then it seems to Sophie that Harry is not bald*. Thus, clearly, either it seems to Sophie that Harry is bald* or it seems to Sophie that Harry is not bald*, even though it is vague which. Given the clear truth of [the relevant response-dependent conditionals], clearly, either Sophie knows that Harry is bald* or she knows that Harry is not bald*, even though it is vague which. Because a thinker knows *whether p* just in case she knows *p* or she knows not-*p*, clearly, Sophie knows *whether* Harry is bald*. (Barnett 2010:39)

This is essentially the original Zenglish argument supplemented with some extra assumptions. First, the law of excluded middle clearly holds, such that $p \vee \neg p$ is clearly true for any *p*. Second, Zenglish intuitions are *complete* in the following sense.

> **INTUITION-COMPLETENESS**
>
> For any subject *s* who considers whether *p* on the basis of some relevant non-testimonial evidence, if *p* then it seems to *s* that *p* and if $\neg p$ then it seems to *s* that $\neg p$

Barnett thinks we are left in the following situation: it is vague *and* knowable whether *p*, although it is both vague whether *p* is knowable and vague whether $\neg p$ is knowable. That situation, however, is impossible—at least, the epistemicist has grounds for thinking so, given his commitment to UNKNOWN.[25]

---

[25] In fact, Barnett's remarks elsewhere (2009:§2.2 last paragraph) indicate that he thinks this is not just a possibility, but *actually*(!) the case for our own vague statements. If so, this bodes ill for Barnett's overall views on vagueness as a sui generis phenomenon. At least Result 6 shows how the epistemicist should resist that conception.

Result 6: Given UNKNOWN, clear knowledge-whether excludes vague knowledge-that.

> *Proof.* Assume for reductio both (i) C(Kp ∨ K¬p) and (ii) IKp & IK¬p. Knowledge is clearly factive; so C(Kp → p) and C(K¬p → ¬p), which by $\mathbf{K_C}$ and contraposition yield C(¬p → ¬Kp) and C(p → ¬K¬p); hence (iii) C¬p → C¬Kp and (iv) Cp → C¬K¬p. But (the clear validity of) disjunctive syllogism gives: C(Kp ∨ K¬p), C¬Kp |- CK¬p and C(Kp ∨ K¬p), C¬K¬p |- CKp. From (i), (iii) and (iv), we can thus derive C¬p → CK¬p and Cp → CKp. Assuming UNKNOWN is true, we have ¬Cp → ¬Kp and ¬C¬p → ¬K¬p; or K¬p → C¬p and Kp → Cp. Therefore K¬p → CK¬p and Kp → CKp. Contraposing gives IK¬p → ¬K¬p and IKp → ¬Kp. Hence (ii) implies ¬Kp & ¬K¬p. But (i) implies Kp ∨ K¬p by $\mathbf{T_C}$, which by one of de Morgan's laws is equivalent to ¬(¬Kp & ¬K¬p). Contradiction.

It thus appears that Barnett's intended "strong counterexample" presentation of the Zengland scenario is incoherent. He cannot simultaneously claim that *whether p* is clearly knowable while *that p* and *that* ¬*p* are both vaguely knowable. Perhaps the latter claims about vague knowledge should be dropped. Doing so, however, only partially amends the situation. There remains a **strengthened response-dependence argument** that does not assume any vagueness in knowability.

S1. Possibly, for some *p*:

a. it is vague whether *p*

b. intuitions about *p* are *clearly* complete

c. intuitions about *p* are *clearly* response-dependent

S2. Necessarily, if intuitions about vague yet *clearly* response-dependent *p* are *clearly* complete, then clearly it can be known whether *p*.

S3. Therefore, possibly: for some *p*, it is vague *and* knowable whether *p*, and so UNKNOWN is false.

If this argument is clearly sound, then UNKNOWN—a purported *necessary* truth about vagueness—is clearly false, and so is the epistemicist theory of vagueness.

Alas, the argument is unsound. But why? Observe that validity is guaranteed by the uncontroversial modal principle that possibility is closed over strict implication. Moreover, premise S2 is verified on minimal classical assumptions.

Result 7: Premise S2 is verified on classical assumptions for modal propositional logic

> *Proof.* Assume that the law of excluded middle clearly holds: for any p, $C(p \vee \neg p)$. Assume for conditional proof that intuitions are clearly complete for some vague, clearly response-dependent matter. Thus for some p: Ip, $C(p \rightarrow Np)$, $C(\neg p \rightarrow N\neg p)$, $C(Np \rightarrow Kp)$, $C(N\neg p \rightarrow K\neg p)$. By clear excluded middle, we get $C(p \vee \neg p)$. Then $C(Kp \vee K\neg p)$ is derivable given $\mathbf{K_C}$ and basic propositional logic. Strengthening by **N** yields our desired result.

This leaves premise S1. Can intuitions be complete when matters are both vague *and* response-dependent? It appears not. The epistemicist who acknowledges the possibility of clearly response-dependent matters in Zengland (premise S1c)—given his commitment to the classical principle of excluded middle—must either deny the clear completeness of Zenglish intuitions (premise S1b) or deny the possibility of vagueness in such a situation (premise S1a).

Either strategy is viable. There is independent reason to deny that Zenglish intuitions can be complete. To constitute knowledge, belief arguably need *not* be truth-tracking in the sense of verifying counterfactuals of the form *If p one would believe p* and *If ¬p one would believe ¬p* (see Williamson 2000 ch.7). Analogously, to be knowledge-conferring, intuition need not be complete, in the sense of verifying "sensitivity" conditionals of the form *If p it would seem to one that p* and *If ¬p it would seem to one that ¬p*. Sensitivity in this sense would require that the intuitions completely cover the facts. Thus, for any Zenglish speaker considering under normal circumstances whether 'bald*' applies to some individual x, his intuitions would be *sensitive* insofar as: if x were/weren't bald* then it would seem to him that x is/isn't bald*. (They would be *robustly sensitive* insofar as: if 'bald*' were/weren't to apply to x then it would seem to him that 'bald*' does/doesn't apply to x.) Yet, despite their stability and accuracy, Zenglish intuitions need *not* be sensitive. Whereas they do exhibit (robust) reliability and consistency, these features do not entail (robust) sensitivity.[26] In this way Zenglish intuitions are able to meet the ideal of RESPONSE-DEPENDENCE without satisfying INTUITION-COMPLETENESS.

Alternatively, Barnett's own considerations provide reason to deny the very possibility of vagueness, given that response-dependence and intuition-completeness are already in place. Barnett thinks that ambivalence in intuition is an essential feature of vagueness. But this is precisely what the situation rules out.

Result 8: AMBIVALENCE is jointly inconsistent with INTUITION-COMPLETENESS, RESPONSE-DEPENDENCE, and UNKNOWN.

> *Proof.* Suppose that intuitions over some vague matter clearly satisfy all the relevant conditions for INTUITION-COMPLETENESS and RESPONSE-DEPENDENCE—for some p: Ip, C($p \rightarrow$ Np), C($\neg p \rightarrow$ N$\neg p$), C(N$p \rightarrow$ Kp), C(N$\neg p \rightarrow$ K$\neg p$). Assume for reductio that the relevant conditions of AMBIVALENCE are also satisfied, such that INp and IN$\neg p$. By clear excluded middle, we get C($p \vee \neg p$). Given INTUITION-COMPLETENESS and RESPONSE-DEPENDENCE, it follows by $\mathbf{K_C}$ that clearly it is known whether p: C(K$p \vee$ K$\neg p$). However, from clear factivity and INTUITION-COMPLETENESS we can derive C(K$p \rightarrow$ Np) and C(K$\neg p \rightarrow$ N$\neg p$); hence the biconditionals C(K$p \leftrightarrow$ Np) and C(K$\neg p \leftrightarrow$ N$\neg p$), given RESPONSE-DEPENDENCE. Assuming closure of vagueness over clear equivalence (see n.22), AMBIVALENCE implies vagueness in knowledge:

[26] Compare: safety requirements on knowledge are distinct from, and do not entail, their corresponding sensitivity requirements; knowledge-conferring belief can fail sensitivity while satisfying safety (see Williamson 2000 ch.7).

IKp and IK$\neg p$. Given Result 6 (which uses UNKNOWN), these two consequences—clarity in knowledge-whether and vagueness in knowledge-that—are inconsistent.

I have already showed how epistemicist assumptions rule out the possibility of Zenglish intuitions being *clearly* ambivalent (see Result 5). Result 8 is stronger: it shows why they cannot be ambivalent *at all*. Yet without ambivalence, vagueness—as Barnett conceives of it—is impossible in Zengland. But then, as before, the threat to epistemicism vanishes. I submit that the epistemicist, in reply to Barnett's "strong counterexample", ought to, as it were, *baldly* assert that the example is incoherent, so conceived.[27]

I hope to have demonstrated how both weak and strong versions of Barnett's objection fail to provide any non-question-begging reasons to reject epistemicism. Conceived weakly as an objection to the *clear* truth of epistemicism, it fails to differentiate distinct orders of vagueness:

---

[27] Yet another alternative option is to deny the rule of Necessitation for clarity ($\mathbf{N_C}$: $\vdash \phi \Rightarrow \vdash C\phi$). This allows us to accept the truth of any instance of the law of excluded middle ($\vdash p \vee \neg p$), without thereby committing to its *clear* truth ($\nvdash C(p \vee \neg p)$). That is enough to block the derivation of Result 7. Though perhaps more radical than any reply offered so far, this move is not entirely unmotivated, since it also blocks the following derivation of an S4-like consequence for vague Zenglish terms. (The proof resembles that of Result 6.)

Result 9: Any clarity operator obeying rules $\mathbf{K_C}$, $\mathbf{T_C}$, $\mathbf{N_C}$ iterates, given INTUITION-COMPLETENESS, RESPONSE-DEPENDENCE and KNOWN.

> *Proof.* Suppose that intuitions over some potentially vague matter clearly satisfy all the relevant conditions for INTUITION-COMPLETENESS, RESPONSE-DEPENDENCE, and KNOWN. Thus for some p: $C(p \rightarrow Np)$, $C(\neg p \rightarrow N\neg p)$, $C(Np \rightarrow Kp)$, $C(N\neg p \rightarrow K\neg p)$, $C(Kp \rightarrow Np)$, $C(K\neg p \rightarrow N\neg p)$. Assuming excluded middle, we have $p \vee \neg p$. Strengthening by $\mathbf{N_C}$ yields $C(p \vee \neg p)$. By repeated applications of $\mathbf{K_C}$ we get $C(Np \vee N\neg p)$ given INTUITION-COMPLETENESS; hence $C(Kp \vee K\neg p)$ given RESPONSE-DEPENDENCE; hence $C(Cp \vee C\neg p)$ given KNOWN. Given the clear factivity of clarity (i.e. $\mathbf{T_C}$ strengthend by $\mathbf{N_C}$), we get $C(Cp \rightarrow p)$ and $C(C\neg p \rightarrow \neg p)$, which by $\mathbf{K_C}$ and contraposition yield $C(\neg p \rightarrow \neg Cp)$ and $C(p \rightarrow \neg C\neg p)$; hence (i) $C\neg p \rightarrow C\neg Cp$ and (ii) $Cp \rightarrow C\neg C\neg p$. By (the clear validity of) disjunctive syllogism, we have: $C(Cp \vee C\neg p), C\neg Cp \vdash CC\neg p$ and $C(Cp \vee C\neg p), C\neg C\neg p \vdash CCp$. Given our earlier result $C(Cp \vee C\neg p)$, from (i) and (ii) we can therefore derive $C\neg p \rightarrow CC\neg p$ and $Cp \rightarrow CCp$.

The guaranteed iteration of clarity is fatal for any possibility of higher-order vagueness. This is plainly undesirable, if one wishes to treat Zenglish terms as exhibiting unclear clarity: perhaps nothing is vague in Zenglish, though this itself is unclear, such that it is unclear whether anything is vague in Zenglish. (A coherent possibility—since it may be verified that $\{Cp \vee C\neg p, ICp, \neg CC\neg p, IIp\}$ are jointly consistent.)

things may be vague, just never *clearly* vague; intuitions about those matters may be ambivalent, just never *clearly* ambivalent. Conceived more strongly as an objection to the *bare* truth of epistemicism, clear or not, the example is incoherent: vagueness and ambivalence are entirely absent from the scenario. Either way, Zengland fails to present any reasons for rejecting the epistemicist theory of vagueness.

Chapter 2: Vague Unknowns

Presumably, any vague term or concept F will admit of cases that are borderline in the sense that we do not know whether they are F or not-F. Patrick Greenough (2003) has argued that such failures to know the underlying status of borderline cases, if any, should figure into a "minimal theory" of vagueness: a theory-neutral characterization of vagueness agreeable to all parties in the debate over the nature of the phenomenon, regardless of their theoretical differences. Along these lines, Timothy Williamson suggests a way to ostensively define the notion of vagueness: give examples of borderline cases, where it is unclear whether something is F, then posit that "an expression or concept is vague if and only if it can resulting unclarity of the kind just exemplified" (1994:2). Mark Sainsbury agrees that "a certain kind of ignorance is a sign of vagueness" (1995:64).

Minimal or not, the idea that ignorance is intimately tied to vagueness is pervasive. It is intuitive enough to think that ignorance, in some sense, should be a characteristic feature of vagueness. Any such claims of ignorance as being a product of vagueness will, of course, not presuppose that there is any underling fact about a given borderline F—either that it is F or that it is not F—to be known in the first place. Let us assume that vague predicates admit of borderline cases, where something is a *borderline case of Fness* (or more simply, *borderline F*) just in case it is indeterminate whether it is F, i.e. it is neither determinately F nor determinately not-F. Then we might articulate the epistemic consequences of vagueness as follows:

UNKNOWN If it is indeterminate whether *p*, it is unknowable that *p*

Plausible as it may be, UNKNOWN has its dissenters. David Barnett objects to UNKNOWN on metaphysical grounds, arguing that even if is true, its truth is neither clear (2009:§§3-4) nor necessary (2010). Crispin Wright refuses to accept it on epistemic grounds, claiming that UNKNOWN is itself unknown (2001:§5.2). Cian Dorr positively rejects UNKNOWN on pragmatic grounds, claiming that it licenses consequences which, if uttered, would be pragmatically infelicitous.

Yet even these deniers of UNKNOWN offer substitutes of their own that are close cousins to the original. I argue that these various candidate ways of articulating how vagueness relates to ignorance are each problematic in their own right, in ways to which the original formulation is impervious.

2.1 Wright I

Wright (2001) argues that borderlineness—and by extension, vagueness, if indeed this is the paradigmatic manifestation of vagueness—should be conceived of in epistemic terms, since borderline cases constitute a subclass of the more general epistemic category of quandary-inducing phenomena.[1] According to Wright's conditions for quandaryhood, a subject S is in a *quandary* over proposition p at time t iff at t:

(i) S does not know whether *p* (i.e. does not know that *p* and does not know that ¬*p*)

(ii) S does not know any way of knowing whether *p*

(iii) S does not know that there is any way of knowing whether *p*

[1] Wright concedes "that the proponents of the Epistemic Conception of vagueness have the matter half right: that indeterminacy *is* an epistemic matter, that borderline cases should be characterized as cases of (a complicated kind of) ignorance." (Wright 2001:93)

(iv) S does not know that it is metaphysically possible to know whether *p*[2]

(v) (for vague *p*) S does not know that it is impossible to know whether *p*[3]

As Greenough (2008:§4) points out, if (i)-(v) are meant to be truly definitional and not just necessary constraints on quandaryhood, this analysis problematically predicts that anyone who has never thought about some claim *p*, and hence is ignorant in all the relevant ways, is thereby in a state of quandary over *p*. The unreflective masses are hence in constant quandary by default! To extend some charity to the folk, we may, following Greenough's suggestion, patch the analysis by strengthening the ignorance conditions. Thus, S is in a *quandary* over *p* at t iff at t:

(i′) S is not in a position to know whether *p*

(ii′) S is not in a position to know any way of knowing whether *p*

(iii′) S is not in a position to know that there is any way of knowing whether *p*

(iv′) S is not in a position to know that it is metaphysically possible to know whether *p*

(v′) S is not in a position to know that: it is impossible to know whether *p*[4]

The key idea is that those in a quandary over *p* are not just unable to know whether *p*, but unable to know that they are unable to know whether *p*. The unknowability of quandary-producing *p* is itself a source of quandary. As Wright puts it: "a quandary is uncertain through and through"

---

[2] Wright's own remarks that "the region of quandary for *F* [just is] the region of cases where we do not know that knowledge of the truth of an *F*-predication is so much as metaphysically possible" (2003a:465) suggest that (i)-(iii) all reduce to condition (iv). Indeed, they are (assuming epistemic closure) all entailed by a strengthened variant of the latter: (iv*) S does not know that it is metaphysically possible to know there is some way of knowing whether *p*.

[3] Wright (2001:92) excludes (v) from the general definition of quandary: it applies only where the discourse is subject to some principle of Evidential Constraint (if *p*, it is feasible to know that *p*)—as with vague statements.

[4] This follows from Wright's claim that it is "unwarranted" to think that any quandary-presenting *p* is feasibly knowable (2003a:§V), since "it is impossible to know whether *p*" entails "it is not feasibly knowable whether *p*", and so the former is unknowable if the latter is (by contraposition and the closure of knowability over entailment).

(2001:92). The iterative character of such ignorance suggests a natural way to spell out the epistemic consequences of vagueness, conceived as a source of quandary:

UNCLEAR If it is indeterminate whether *p*, it is indeterminate whether it is knowable that *p*

Of course, in the absence of such obstacles to epistemic clarity, whether due to vagueness or any other source of quandary, things can, all else held equal, be clearly known.

CLEAR If it is determinate that *p* (and vagueness is the only potential source of ignorance preventing one from knowing whether *p*), it is determinately knowable that *p*

(The requirement that vagueness be the only relevant potential obstacle to knowledge is meant to exclude other quandary-inducing but determinate claims (e.g. Goldbach's Conjecture), as well as other knowledge defeaters unrelated to vagueness (e.g. perceptual illusions, Gettier cases), from the scope of CLEAR. In what follows, for ease of exposition, I will assume this requirement is satisfied for the relevant discourse.)

For ease of exposition, I shall often abbreviate 'knowably *p*' and 'definitely *p*' as 'K*p*' and 'D*p*'. I take the resulting logic of determinacy to be at minimum **KT**, since determinacy is arguably factive and closed under implication.[5] Indeterminacy is interdefinable with determinacy ($Ip =_{def} \neg Dp \;\&\; \neg D\neg p$). I shall assume throughout that indeterminacy is invariant under negation:[6]

---

[5] That is, determinacy operator 'D' obeys the modal principles $\mathbf{K_D}$: $D(p \rightarrow q) \rightarrow (Dp \rightarrow Dq)$ and $\mathbf{T_D}$: $Dp \rightarrow p$. And given epistemic closure and factivity, knowability operator 'K' obeys $\mathbf{K_E}$: $K(p \rightarrow q) \rightarrow (Kp \rightarrow Kq)$ and $\mathbf{T_E}$: $Kp \rightarrow p$.

[6] SYMMETRY is plausibly a basic datum for "it is vague whether…" claims. It is arguably a general feature of such *wh*-constructions that "it is F (vague, un/known, un/decided, etc.) whether *p*" entails "it is F

SYMMETRY If it is indeterminate whether *p*, it is indeterminate whether ¬*p*

Instances of UNCLEAR and CLEAR are independently derivable given Wright's principle of Evidential Constraint governing potentially vague atomic statements.[7]

EC If Fa, then it is feasibly knowable that Fa

Anything feasibly knowable is knowable, so EC reduces to:[8]

(1) If Fa, then it is knowable that Fa

Wright endorses the *a priori* knowability of EC on antirealist grounds. Presumably, those same grounds justify the *determinate* truth of (1). Assuming determinacy is closed over implication (**K_D**), we get:

(2) If determinately Fa, it is determinately knowable that Fa

This validates CLEAR for atomic *p* (assuming vagueness is the only potential source of ignorance). Since "knowable…" is determinately factive, (1) can be strengthened to be a determinately true biconditional, from which it follows:[9]

---

whether or not *p*", where this in turn (perhaps because it is equivalent to "it is F whether *p* or whether not-*p*", and hence to "it is F whether not-*p* or whether *p*") entails "it is F whether not-*p*". See also n.46.

[7] Wright cautiously restricts EC to atomic predications of form Fa or ¬Fa, expressible within any alleged antirealist discourse where cognitive command is assured (i.e. any difference of opinion will concern a knowable matter), including vague discourses (i.e. those that admit of vague statements). See Wright 2001:59-60 and fn.15-17,41,44.

[8] For Wright, 'it is feasible to know *p*' is weaker than 'one knows *p*', yet is factive—presumably because it entails 'it is knowable that *p*', which is factive—and hence stronger than 'it is logically (conceptually) possible to know *p*'. I shall from here on drop the qualifier "feasibly" for "knowable" (a cumbersome nuisance—what does it *really* add?). At any rate, any inference from 'S can feasibly know *p*' to 'S is positioned to know *p*' is licensed by Wright himself, who explains that for something to be feasibly knowable just is for it to be (humanly) knowable by someone "appropriately placed to recognise" its truth (2001:n.17).

(3) It is indeterminate whether Fa iff it is indeterminate whether it is knowable that Fa

which validates UNCLEAR for atomic p.

Whether UNCLEAR and CLEAR are truly representative of Wright's official views of vagueness as a type of quandary, there is independent reason to find these principles problematic. Observe that the notion of indeterminacy appears in both antecedent and consequent for (any instance of) UNCLEAR. It says that anything indeterminate is *indeterminately* knowable. Reapplying the principle to this claim of indeterminate knowability gets:

(4) If it is indeterminate whether it is knowable that p, it is indeterminate whether it is knowable that it is knowable that p

Thus anything indeterminate is indeterminately knowably knowable. We can then apply the principle to *that* claim of indeterminate knowable knowability. And so on, up through higher orders of knowability. Iterating the procedure results in the sequence of substitution instances: $Ip \rightarrow IKp$, $IKp \rightarrow IKKp$, $IKKp \rightarrow IKKKp$, etc. Chaining together any finite number of substitutions allows us to derive for any arbitrary n:

(5) If it is indeterminate whether p, it is indeterminate whether it is knowable^n that p

(where being knowable^n is being knowably knowably… [n-1 times] knowable) A similar result holds for CLEAR, which too can reapply without end to its own predictions, thereby producing

[9] Strengthening $\mathbf{K_D}$ gets us $D(Kp \rightarrow p)$ for atomic p. The determinate truth of (1) gives $D(p \rightarrow Kp)$. Combining yields $D(p \leftrightarrow Kp)$. But from any $D(A \leftrightarrow B)$ we can derive the consequence $IA \leftrightarrow IB$. *Proof.* First consider $D(A \rightarrow B)$. By **K** and contraposition (i.e. distributing 'D' over $D((A \rightarrow B) \rightarrow (\neg B \rightarrow \neg A))$), we can infer $D(\neg B \rightarrow \neg A)$. Distributing again gets $D\neg B \rightarrow D\neg A$, which by contraposition is (i) $\neg D\neg A \rightarrow \neg D\neg B$. Now consider $D(B \rightarrow A)$. By the same principles, distributing then contraposing delivers (ii) $\neg DA \rightarrow \neg DB$. Combining (i) and (ii) yields $IA \rightarrow IB$. $IB \rightarrow IA$ follows by symmetry of reasoning; hence $IA \leftrightarrow IB$. But assumption $D(A \leftrightarrow B)$ just is $D(A \rightarrow B)$ and $D(B \rightarrow A)$.

the sequence of substitution instances: D*p* → DK*p*, DK*p* → DKK*p*, DKK*p* → DKKK*p*, etc. Or for any n:

(6) If it is determinate that *p*, it is determinate that it is knowable^n that *p*

In this way, UNCLEAR licenses unrestricted iteration of the knowability operator within contexts of indeterminacy, while CLEAR licenses its unrestricted iteration within contexts of determinacy. Together, they license unrestricted iteration for claims of knowability *everywhere*. That is to say, they validate the notorious—and widely rejected—KK Principle (see Williamson 1994:§8.2 and 2000:§5).

KK If it is knowable that *p*, it is knowable that it is knowable that *p*

It appears that UNCLEAR and CLEAR lead to unpalatable epistemic consequences. Is there an alternative way to conceive of the epistemic consequences of vagueness available to the quandary account?

2.2 Wright II

Consider the following variants of UNCLEAR and CLEAR:

UNCLEAR* If it is determinately indeterminate whether *p*, it is unknowable whether *p*

CLEAR* If it is determinately determinate that *p* (and vagueness is the only potential source of ignorance over whether *p*), it is knowable that *p*

UNCLEAR* and CLEAR* lack the self-iterating character had by UNCLEAR and CLEAR, and so pose no (at least, immediate) danger of validating KK.

Yet they appear to fit Wright's conception of quandary. The idea is to stipulate that ignorance ensues only for a narrow, limited range of borderline cases (namely, the definitely borderline cases), but not fully specify whether ignorance is mandated in *all* borderline cases. Partially stating the epistemic consequences of vagueness in this way leaves open the possibility that we are *never* in a position to know that it is unknowable whether *p* for some vague *p*, so long as we cannot know whether any vague *p* falls within that select range. This would vindicate the anti-undecidability condition (v) for quandaryhood.

They can also be seen as verifying, in conjunction with other principles endorsed by Wright, Wright's other views on vagueness, apart from the general discussion concerning quandary. Consider Wright's DEF Principle for any suitable determinacy/definiteness operator 'D'.

DEF — If $\Sigma \models \phi$ and every atomic sentence in Σ is in the scope of a 'D', then $\Sigma \models D\phi$

This effectively says that any true consequence of any set of propositions each definitely true must be a *definitely* true consequence. DEF immediately validates the S4 Axiom for determinacy: $Dp \rightarrow DDp$. An important consequence of the resulting iterativity of the determinacy operator is that it collapses vagueness of higher orders into first-order vagueness:[10]

COLLAPSE — If it is indeterminate whether it is indeterminate whether *p*, it is indeterminate whether *p*

We can then reason as follows:

1	(1)	$DIIp$	A (for reductio)
2	(2)	$DIIp \rightarrow \neg KIp \,\&\, \neg K\neg Ip$	UNCLEAR*
1,2	(3)	$\neg KIp \,\&\, \neg K\neg Ip$	modus ponens,1,2
1,2	(4)	$\neg KIp$	&-elim, 3

[10] *Proof.* Assume IIp. Thus $I\neg Ip$ by SYMMETRY, or $I\neg(\neg Dp \,\&\, \neg D\neg p)$ by definition. By the determinacy of the de Morgan equivalences, $I(Dp \vee D\neg p)$. Hence $\neg D(Dp \vee D\neg p)$ by definition, which entails both $\neg DDp$ and $\neg DD\neg p$ (because DDp would imply Dp by $\mathbf{T_D}$, or $D(Dp \vee D\neg p)$ by $\mathbf{K_D}$ and $\vee$-intro, so $\neg DDp$ by reductio; likewise for $DD\neg p$). It follows, by reductio, that both $\neg Dp$ and $\neg D\neg p$, given DEF. Therefore Ip. (Note: This is essentially the missing proof for Bobzien's (2011) unexplained principle (UU/U).)

5	(5)	$DDIp \rightarrow KIp$	CLEAR*
6	(6)	DIp	A (for reductio)
6	(7)	$DDIp$	DEF, 6
5,6	(8)	KIp	modus ponens,5,7
1,2,5	(9)	$\neg DIp$	reductio,4,6,8
1	(10)	IIp	$\mathbf{T_D}$, 1
11	(11)	$IIp \rightarrow Ip$	$\mathrm{COLLAPSE}_D$,10[11]
1,11	(12)	Ip	modus ponens,10,11
1,11	(13)	DIp	DEF, 12
2,5,11	(14)	$\neg DIIp$	reductio,1,9,13

This essentially rules out the existence of definite higher-order vagueness—specifically, of definite borderline borderlineness. Wright would surely approve.

However, in order for quandary condition (v) to be satisfied in the borderlines, we must suppose that nothing can be *known* to be definitely borderline. Otherwise, by UNCLEAR* (assuming we know this to be true), knowing for some *p* that it is determinately indeterminate whether *p* would mean knowing it is unknowable whether *p*—which violates the anti-undecidability constraint on quandaryhood. Consequently, given CLEAR* and DEF, nothing can be definitely borderline, knowably or otherwise.[12] Therefore, anything borderline will be borderline borderline.[13] Hence the iterativity of vagueness-related indeterminacy:[14]

ITERATE — If it is indeterminate whether *p*, it is indeterminate whether it is indeterminate whether *p*

---

[11] The 'D' subscript indicates that the principle in question is assumed to determinately hold, from which the relevant substitution instance is derived via $\mathbf{T_D}$, so as to license use of DEF later on.

[12] Otherwise, DIp would imply $DDDIp$ by (two applications of) DEF, hence $KDIp$ by CLEAR*, from which it follows by $\mathbf{K_E}$ that $K(\neg Kp \ \&\ \neg K\neg p)$, assuming UNCLEAR* is known to hold (at least for definitely borderline *p*).

[13] Assume Ip. Suppose $D\neg Ip$ (for reductio). Then $\neg Ip$ by $\mathbf{T_D}$—contradiction. Hence $\neg D\neg Ip$. But $\neg DIp$, given CLEAR* and DEF. Therefore IIp.

[14] The "vagueness-related" qualifier is inherited from CLEAR*. DEF, by contrast, is presumably a principle that applies to *all* types of determinacy, vagueness-related or otherwise.

This is the converse of COLLAPSE (where vagueness is concerned). Together, they say that to be borderline borderline (or second-order borderline, as it were) *just is* to be (unqualifiedly) borderline. Insofar as vagueness of any order finds expression through iterated claims of borderlineness using a 'definitely' operator, any second-order vagueness (should it exist) can only be one of two things: borderline borderlineness or borderline definiteness. DEF already rules out the possibility of anything being borderline definite—since if definiteness iterates, nothing definite can be borderline definite. So any case of second-order vagueness must be a borderline borderline case, and therefore a borderline case by COLLAPSE. Thus, second-order vagueness reduces to first-order vagueness.

Any vagueness of higher order $n > 2$ is also reducible. For any *n*th-order vagueness (should it exist) must find expression through iterated claims of borderlineness, where these are either ID(…) or II(…).[15] The first sort of claim is ruled out by DEF, leaving only the second possibility. Any case of *n*th-order vagueness must then be a borderline^n case (i.e. a borderline borderline… [n times] case), and therefore a borderline case by COLLAPSE. What ITERATE guarantees is that the range of borderline cases covers all borderline^n cases for any *n*. Therefore, no borderline case is *merely* borderline, in the sense of being first-order borderline without being higher-order borderline. Together, COLLAPSE and ITERATE might therefore be read reductively, as saying that all higher-order vagueness reduces to first-order vagueness.

Wright (1987, 1992) has used the DEF Principle to argue against the existence of higher-order vagueness as something *distinct* from—and, given the expressive resources afforded by a definitely operator and its iterated variations when combined with negation, *expressively*

---

[15] For any claim ϕ there are only four permutations for stacking Ds and Is: DDϕ, DIϕ, IDϕ, and IIϕ. But only the last two of these are claims of borderlineness.

distinguishable from—first-order vagueness. This combination of COLLAPSE and ITERATE, on their reductive reading, falls in step—if not in letter then surely in spirit—with Wright's denial of (what is *distinctively*) higher-order vagueness, since it shows that there is really only one type of vagueness after all: first-order vagueness.

It is worth noting that, in contrast with most discussions in the literature (including Wright's), I have made no mention of any "gap principles" typically associated with higher-order vagueness.[16] The reducibility proof presented here simply rests upon certain independently motivated principles concerning the epistemic status of individual borderline cases. Nothing is assumed about how adjacent cases in a Sorites series are related—specifically, whether a gap of borderline cases must serve as a buffering zone for every pair of contrary determinate categories at any alleged order of vagueness, or how the subsequent lack of sharp boundaries for any vague category lends to the appearance, illusory or otherwise, of it achieving a smooth transition across a Sorites series without paradox.

Unfortunately, the view contains some surprising epistemic consequences. To streamline our discussion, let us make some initial observations. First, failure of determinate falsity is closed over determinate implication:[17]

DEBATABLE   If it is determinate that if *p* then *q*, if *p* is debatable then *q* is debatable

---

[16] See e.g. Fara (2000), Wright (1987, 1992, 2011), Edgington (1992), Heck (1992), Sainsbury (1991).

[17] Assume $D(A\rightarrow B)$. Given $D((A\rightarrow B)\rightarrow(\neg B\rightarrow\neg A))$ (i.e. assuming contraposition is 'determinately' valid), by $\mathbf{K_D}$ and modus ponens we have $D(\neg B\rightarrow\neg A)$. By $\mathbf{K_D}$ again, $D\neg B\rightarrow D\neg A$. By contraposition again, $\neg D\neg A\rightarrow\neg D\neg B$.

where "…is debatable" abbreviates "it is not determinate that it is not the case that…" ('$\neg D \neg \ldots$'). Second, higher-order vagueness makes it debatable that first-order vagueness is absent, such that the determinacy of either p or $\neg p$ remains debatable:[18]

DEBATABLE* If p is indeterminately indeterminate, either the determinacy of p is debatable or the determinacy of $\neg p$ is debatable

Third, vagueness excludes determinate knowability.[19]

EXCLUSION If it is indeterminate whether p, it is not determinately knowable that p

Now consider an intuitionistically friendly version of what I shall call ***The IK Argument***.[20]

1	(1)	Ip	A
1	(2)	IIp	ITERATE, 1
1	(3)	$\neg D \neg Dp \vee \neg D \neg D \neg p$	DEBATABLE*, 2
4	(4)	$Dp \rightarrow DDp$	DEF (from $Dp \vdash Dp$)
5	(5)	$DDp \rightarrow Kp$	CLEAR*$_D$
4,5	(6)	$Dp \rightarrow Kp$	transitivity, 4,5
4,5	(7)	$D(Dp \rightarrow Kp)$	DEF (from 4,5 $\vdash$ 6)
8	(8)	$D(Dp \rightarrow Kp) \rightarrow (\neg D \neg Dp \rightarrow \neg D \neg Kp)$	DEBATABLE

---

[18] Assume IIp. So $\neg DIp$, or $\neg D(\neg Dp \,\&\, \neg D \neg p)$. Then $\neg D \neg Dp \vee \neg D \neg D \neg p$, given $\neg D(A \& B) \vdash \neg DA \vee \neg DB$. This last inference is justified by classical reasoning. *Proof.* Assume DA and DB. Since &-intro is "determinately" valid, we have $D(A \rightarrow (B \rightarrow (A\&B)))$. By $\mathbf{K_D}$ this becomes $DA \rightarrow D(B \rightarrow (A\&B))$. By modus ponens on DA, we obtain $D(B \rightarrow (A\&B))$. Applying $\mathbf{K_D}$, &-elim, and modus ponens once more, now for DB, yields $D(A\&B)$. In this way determinacy collects over conjunction: $DA\&DB \vdash D(A\&B)$. Contraposing gets $\neg D(A\&B) \vdash \neg(DA\&DB)$. By de Morgan ($\neg\&$), $\neg D(A\&B) \vdash \neg DA \vee \neg DB$. To be sure, intuitionistic logic does not recognize the validity of the de Morgan ($\neg\&$) transformation. However, deriving $\neg DA \vee \neg DB$ does not appear objectionable for any of the standard intuitionist considerations (against excluded middle etc.). Moreover, DEBATABLE* remains an intuitively plausible truth about second-order vagueness, despite failing to be independently derivable within an intuitionistic system (see also n.46 on SYMMETRY).

[19] Assume Ip, so $\neg Dp$ and $\neg D \neg p$. Suppose (for reductio) DKp. By $\mathbf{T_D}$, Kp. By $\mathbf{T_E}$, p. By DEF (using $DKp \vdash p$), Dp. Contradiction. Therefore $\neg DKp$. (Similarly, supposing $DK\neg p$ would contradict $\neg D \neg p$. Therefore $\neg DK \neg p$.)

[20] Engineered to be intuitionistically acceptable, so not even intuitionism-sympathizers like Wright should object.

4,5,8	(9)	$\neg D\neg Dp \rightarrow \neg D\neg Kp$	modus ponens, 7,8
10	(10)	$Ip \rightarrow \neg DKp$	EXCLUSION
1,10	(11)	$\neg DKp$	modus ponens, 1,10
12	(12)	$\neg D\neg Kp$	A (for →-proof)
1,10,12	(13)	$\neg D\neg Kp \,\&\, \neg DKp$	&I, 11,12
1,10,12	(14)	IKp	def(I), 13
1,10,12	(15)	$IKp \vee IK\neg p$	∨-intro, 14
1,10	(16)	$\neg D\neg Kp \rightarrow (IKp \vee IK\neg p)$	→-elim, 12,15
1,4,5,8,10	(17)	$\neg D\neg Dp \rightarrow (IKp \vee IK\neg p)$	transitivity, 9,16
18	(18)	$D\neg p \rightarrow DD\neg p$	DEF (from $D\neg p \vdash D\neg p$)
19	(19)	$DD\neg p \rightarrow K\neg p$	CLEAR*$_D$
18,19	(20)	$D\neg p \rightarrow K\neg p$	transitivity, 18,19
18,19	(21)	$D(D\neg p \rightarrow K\neg p)$	DEF (from 18,19 |- 20)
22	(22)	$D(D\neg p \rightarrow K\neg p) \rightarrow$ $(\neg D\neg D\neg p \rightarrow \neg D\neg K\neg p)$	DEBATABLE
18,19,22	(23)	$\neg D\neg D\neg p \rightarrow \neg D\neg K\neg p$	modus ponens, 21,22
24	(24)	$Ip \rightarrow I\neg p$	SYMMETRY
1,24	(25)	$I\neg p$	modus ponens, 1,24
25	(26)	$I\neg p \rightarrow \neg DK\neg p$	EXCLUSION
1,24,25	(27)	$\neg DK\neg p$	modus ponens, 25,26
28	(28)	$\neg D\neg K\neg p$	A (for →-proof)
1,24,25,28	(29)	$\neg D\neg K\neg p \,\&\, \neg DK\neg p$	&I, 27,28
1,24,25,28	(30)	$IK\neg p$	def(I), 29
1,24,25,28	(31)	$IKp \vee IK\neg p$	∨-intro, 30
1,24,25	(32)	$\neg D\neg K\neg p \rightarrow (IKp \vee IK\neg p)$	→-elim, 28,31
1,18,19,22,24,25	(33)	$\neg D\neg D\neg p \rightarrow (IKp \vee IK\neg p)$	transitivity, 23,32
1,4,5,8,10,18 ,19,22,24,25	(34)	$IKp \vee IK\neg p$	proof by cases, 3,17,33

In this way, vagueness can be seen to generate vagueness in knowledge. The disjunctive result here—that if *p* is indeterminate then either the knowability of *p* or the knowability of ¬*p* is also indeterminate[21]—is strictly speaking weaker than the old principle UNCLEAR, which predicted

[21] The intended result is of form $Ip \rightarrow (\ldots IKp \ldots)$. Yet those who for broadly supervaluationist reasons reject conditional proof might question whether deriving …IK*p*… from a non-(super)true premise I*p* guarantees the (super)truth of $Ip \rightarrow (\ldots IKp \ldots)$). Fortunately, there is a way to reconceive the proof without the deduction theorem. Given ITERATE, we have $\vdash Ip \rightarrow I^2p$. Within a fully classical setting, one could derive $\neg D\neg Dp \vee \neg D\neg D\neg p$ from I^2p. This can be codified in the supervaluationist setting as a validity $\vdash Ip \rightarrow (\neg D\neg Dp \vee \neg D\neg D\neg p)$. (This is *not* derived, of course, by deriving the consequent from the antecedent, but rather justified, as it were, in its own right.) But we have $\vdash Dp \rightarrow DDp$ by DEF, and $\vdash DDp \rightarrow Kp$ by

that both p and $\neg p$ are indeterminately knowable, provided the indeterminacy of p (and hence[22] of $\neg p$).[23] Yet things are arguably just as bad with UNCLEAR* and CLEAR*, as they were with UNCLEAR and CLEAR.

Observe that because of DEF, CLEAR* reduces to CLEAR.[24] Hence it too validates the KK Principle within the scope of determinacy:

KK_{DET} If it is determinate that p, it is determinate that it is knowable^{n} that p

Given DEF, ITERATE, KK^{D} (and hence CLEAR*$_{D}$), *reductio ad absurdum*, and other forms of (what are by Wright's standards) intuitionistically acceptable reasoning, we can reproduce the reasoning exemplified above in the IK Argument, except now applied to generalized higher-order knowability claims, and derive the conclusion $(\neg DK^{n}\neg p \,\&\, IK^{n}p) \vee (\neg DK^{n}p \,\&\, IK^{n}\neg p)$.[25] Thus,

KK_{INDET} If it is indeterminate whether p, then either it is indeterminate whether it is knowable^{n} that p or it is indeterminate whether it is knowable^{n} that $\neg p$

CLEAR*, so $\vdash Dp \rightarrow Kp$. But then $\vdash \neg Kp \rightarrow \neg Dp$ (since contrapositives of valid conditionals are valid). Definitizing by DEF gets $\vdash D(\neg Kp \rightarrow \neg Dp)$, so $\vdash D\neg Kp \rightarrow D\neg Dp$ by $\mathbf{K_D}$. Contraposing again gets $\vdash \neg D\neg Dp \rightarrow \neg D\neg Kp$. By parallel reasoning, $\vdash \neg D\neg D\neg p \rightarrow \neg D\neg K\neg p$. Therefore, $\vdash Ip \rightarrow (\neg D\neg Kp \vee \neg D\neg K\neg p)$, by conditionalized proof by cases: If $\vdash \phi\rightarrow(\varphi \vee \psi)$, $\vdash \varphi \rightarrow \zeta$, $\vdash \psi \rightarrow \zeta$ then $\vdash \phi \rightarrow \zeta$.

[22] Although intuitionism only recognizes $I\neg p \Rightarrow Ip$ (but *not* $Ip \Rightarrow I\neg p$). *Proof.* Assume $I\neg p$; so $\neg D\neg p$ and $\neg D\neg\neg p$. Suppose (for reductio) Dp. By $\mathbf{T_D}$, p. By DNI, $\neg\neg p$. By DEF, $D\neg\neg p$. Contradiction. So $\neg Dp$. But $\neg D\neg p$. Hence Ip.

[23] Unfortunately, the result (line 42) must remain disjunctive, without further simplication. To derive $IKp \,\&\, IK\neg p$, we need *both* $\neg D\neg K\neg p$ and $\neg D\neg Kp$. Given DEF and CLEAR*, these *would* follow from $\neg D\neg D\neg p$ and $\neg D\neg Dp$. But IIp only guarantees that one of these is true.

[24] *Proof.* Assume Dp. By DEF(x2), $DDDp$. Suppose CLEAR* is determinately true, so $D(DDp \rightarrow Kp)$. By $\mathbf{K_D}$, $DDDp \rightarrow DKp$. So DKp by modus ponens.

[25] The proof is essentially the same, except now a series of reductios (provided Ip) of claims $D\neg K^{n}p \,\&\, D\neg K^{n}\neg p$, $DK^{n}p$, and $DK^{n}\neg p$, instead of reductios of $D\neg Kp \,\&\, D\neg K\neg p$ (line 5), DKp (line 26), and $DK\neg p$ (line 32).

Given Ip, either p or $\neg p$ will turn out to be indeterminately knowable^n: which one (if not both), depends on whether it is $\neg D \neg Dp$ or $\neg D \neg D \neg p$ that holds—although one of these must, given IIp (by ITERATE).

Hence, the following tetralemma. Either

1. p is determinately true (Dp), in which case p is determinately knowable^n for any arbitrary n,
2. p is indeterminate but does not determinately fail to be determinately true (Ip & $\neg D \neg Dp$), in which case p is indeterminately knowable^n for any arbitrary n,
3. p is indeterminate but does not determinate fail to be determinately false (Ip & $\neg D \neg D \neg p$), in which case $\neg p$ is indeterminately knowable^n for any arbitrary n, or
4. p is determinately false (D$\neg p$), in which case $\neg p$ is determinately knowable^n for any arbitrary n

where, whatever the outcome, what is knowable is knowably knowable, determinately so or otherwise.[26]

We saw earlier that UNCLEAR and CLEAR validate the objectionable KK Principle. It appears that UNCLEAR* and CLEAR* suffer the same problem. Evidently, the suggested recourse to a partial specification of the epistemic ramifications of vagueness offers no escape from that unpalatable result.

---

[26] Di(tri, etc.)lemma arguments often appeal to some metarule of proof by cases: If A |- C and B |- C then A∨B |- C. Those with supervaluationist reservations (see n.21) may simply reconceive such reasoning by cases as proceeding on the (super)truth of conditionals, rather than inferential relations—i.e. as validating A → C, B → C, A∨B |- C.

2.3 Wright III

Is there no viable way then to codify Wright's conception of quandary in epistemic principles for vagueness? Perhaps it will do to simply write the anti-undecidability condition (v) directly into a principle articulating the epistemic effects of quandary-inducing forms of indeterminacy like vagueness, as in:

UNKNOWN* If it is indeterminate whether *p*, it is unknowable whether *p* and unknowable whether it is unknowable whether *p*

and its inverse principle:

KNOWN* If it is determinate that *p* (and vagueness is the only potential source of ignorance over whether *p*), it is knowable that *p*

Like the previous account, UNKNOWN* and KNOWN* appear to jointly predict that nothing is definitely borderline. For anything determinately indeterminate will be knowably indeterminate, given KNOWN*. But by UNKNOWN* (assuming we can know this to be true), this means it is both knowably unknowable and (knowably) unknowably unknowable—which is impossible.[27] Therefore, any indeterminacy exhibited by a vague claim is itself indeterminate. This is consonant with Wright's idea that any sense of quandary generated by a vague claim is itself a source of quandary. As Wright insists, a state of quandary is by nature not something one can know oneself to be in, since one's own ignorance about the matter at hand must itself remain unknowable.

---

[27] *Pf.* Assume DI*p*. Suppose UNKNOWN* determinately holds: $D(Ip \rightarrow Up \;\&\; UUp)$, where $Up =_{def} \neg Kp \;\&\; \neg K\neg p$. So $DIp \rightarrow D(Up \;\&\; UUp)$ by $\mathbf{K_D}$. By modus ponens, $D(Up \;\&\; UUp)$. By &-elim and $\mathbf{K_D}$, we get both DUp and $DUUp$. The former implies KUp, given KNOWN*. But the latter simplifies to UUp by $\mathbf{T_D}$; so $\neg KUp$. Contradiction.

However, the IK Argument returns to cause trouble. We saw earlier that, provided DEF, *reductio ad absurdum*, and other basic (yet intuitionistically valid) inference rules, we can use principles ITERATE and CLEAR*$_D$ to derive the result that indeterminacy produces indeterminate (un)knowability. The combination of UNKNOWN* and KNOWN*, assuming these are both determinately true, validates those same critical principles exploited in the IK Argument—specifically, ITERATE and CLEAR*$_D$. The IK Argument effectively shows that knowability claims iterate within contexts of indeterminacy. And because of DEF, KNOWN* guarantees that knowability claims iterate within contexts of determinacy.[28]

One may object: why accept KNOWN*, when this together with DEF already licenses KK, at least for determinately true claims? Implicit in the very enterprise of trying to characterize vagueness in terms of its epistemic effects is, I take it, the aim of advancing principles that identify some characteristic feature of vagueness that is epistemic. Such principles need not be reductive (recall UNCLEAR, which spells out the epistemic effects of indeterminacy in terms themselves indeterminate), nor need they be complete (recall UNCLEAR*, which specifies the epistemic features of only a limited range of borderline cases). They must however identify the epistemic markers of vagueness—that is, identify both markers that indicate the presence of vagueness *as well as* markers that indicate its absence. As such, these epistemic principles must come in pairs: one to address the effects of indeterminacy, another to address those of determinacy. If vague matters mandate ignorance (perhaps of some special sort awaiting further explication), surely non-vague matters—where neither vagueness nor any other potential obstacle

[28] From DKp, we can show DK$^n p$ using induction on n. The base step is trivial. For the inductive step, assume DKp. By the inductive hypothesis, DK$^n p$ is derivable. By DEF, DDK$^n p$. By KNOWN*$_D$, D(DK$^n p \rightarrow$ KK$^n p$). By $\mathbf{K_D}$, DDK$^n p \rightarrow$ DKK$^n p$. By modus ponens, DKK$^n p$, or (rewritten) DK$^{n+1} p$.

to knowledge is present—enable knowledge (perhaps of another special sort). Why *shouldn't* perfectly determinate truths be knowable?

Relatedly, Bobzien (2010) has argued that borderlineness should be understood in epistemic terms, as producing a sort of *radical unclarity*, in the sense that it is possible, for any vague statement *p*, that some truthtelling speaker, who is maximally competent (with respect to evaluating whether *p*) and maximally informed (with respect to non-trivial evidence for evaluating whether *p*), is nonetheless unable to tell whether *p*, where this very claim of unclarity is itself subject to the conditions of unclarity (i.e. possibly some truthtelling speaker who is maximally competent and informed with respect to evaluating whether it is unclear whether *p* cannot tell whether it is unclear whether *p*), and *that* claim of unclear unclarity is itself unclear in the same manner, and so on up through all higher orders of unclarity. The iterative character of such unclarity eliminates the possibility of definite borderline cases. On Bobzien's view, the absence of vagueness also results in a sort of self-iterating clarity, where for any determinate *p*, necessarily, all truthtelling subjects who are competent and informed with respect to evaluating *p* will be able to tell that *p*, and all higher-order claims about the clarity of (the clarity of…) *p* are themselves clear.

The current view would appear to encapsulate Bobzien's idea that, where vagueness is concerned, the epistemic aspects of both indeterminacy and determinacy take on a self-iterative character: UNKNOWN* ensures that vagueness produces unclarity at all higher orders, so that nothing is definitely borderline, while KNOWN*, in conjunction with DEF, ensures that determinacy produces clarity at all higher orders. Insofar as UNKNOWN*, KNOWN*, and DEF accurately represent the overall structure of Bobzien's view, the objection raised here against this combination of principles applies equally to the latter: the conception of vagueness as a source of radical unclarity has untenable epistemic consequences.

2.4 Barnett

Barnett (2009) has argued that vagueness is an irreducible, sui generis phenomenon. Nonetheless, it comes with epistemic ramifications, albeit not of the sort envisioned by the epistemicist. Consider some borderline case of baldness. According to Barnett,

> "[…] it is vague whether you believe that [the claim that e.g. Harry is bald, where Harry is borderline bald] is true and vague whether you believe that it is false. While you are in no position to assert either that it is true or that it is false, your state of mind is consistent with the absence of a hidden fact of the matter, for it is consistent with the clear truth of the following biconditionals:
>
> (16) you know that Harry is bald iff Harry is bald.
> (17) you know that Harry is not bald iff Harry is not bald.
>
> If (16) and (17) are clearly true, then vagueness in their right-hand sides entails vagueness in their left- hand sides. So, you need not be *ignorant* of what is going on. You simply do not *clearly* know what is going on, in the sense that you do not clearly know that it is true that Harry is bald and you do not clearly know that it is false that Harry is bald. This is no surprise, for nothing *is* clearly going on: there is no clear fact of the matter. Vagueness as to whether Harry is bald is, without any further analysis, sufficient to explain why you cannot clearly know either proposition: vagueness as to whether *p* entails vagueness as to whether a certain necessary condition on knowing *p* obtains, namely, *p*'s truth." (Barnett 2009:§2.2)

The idea is this: it may be indeterminate whether *p* is knowable, but this is compatible with there being a fact of the matter about whether *p*. Our epistemic situation toward *p* is indeterminate: we are neither clearly knowledgeable about whether *p* nor clearly ignorant about whether *p*. Rather, both *p* and $\neg p$ are indeterminately (un)knowable. Given excluded middle, exactly one of these *is* in fact knowable (the other must be unknowable, given that 'knowable' is factive), it is just unclear which.

This appears, in part, to be an endorsement of UNCLEAR. That so much is a common feature shared by both Barnett's view of vagueness as sui generis and (one reading of) Wright's view of vagueness as quandary. One might think that Barnett's account is nonetheless safe from the complications facing Wright's view, since it is free of any commitment to CLEAR, given that Barnett nowhere avows any determinacy-introducing principle such as Wright's DEF that would license the iteration of determinacy. However, Barnett's appeal to excluded middle proves to be uniquely problematic.[29] The idea of us either knowing p or knowing $\neg p$, it being simply unclear which it is we know when p is vague, already grates harshly against all natural intuition.[30] But setting intuitions aside, classical reasoning demonstrates that this ends up committing Barnett's proposal to the truth of CLEAR. For Barnett assumes that the law of excluded middle not only holds but *determinately* holds for any vague p, so that it is *determinately* true that either p is knowable or $\neg p$ is knowable (i.e. $D(Kp \vee K\neg p)$), albeit vague which. Thus,

1	(1)	$D(Kp \vee K\neg p)$	A
2	(2)	$D((Kp \vee K\neg p) \rightarrow (\neg K\neg p \rightarrow Kp))$	$\vee$-syllogism$_D$
2	(3)	$D(Kp \vee K\neg p) \rightarrow D(\neg K\neg p \rightarrow Kp)$	$\mathbf{K_D}$, 2
1,2	(4)	$D(\neg K\neg p \rightarrow Kp)$	modus ponens, 1,3
1,2	(5)	$D\neg K\neg p \rightarrow DK\neg p$	$\mathbf{K_D}$, 4
6	(6)	$D(K\neg p \rightarrow \neg p)$	$(\mathbf{T_E})_D$
7	(7)	$D((K\neg p \rightarrow \neg p) \rightarrow (\neg\neg p \rightarrow \neg K\neg p))$	contraposition$_D$
7	(8)	$D(K\neg p \rightarrow \neg p) \rightarrow D(\neg\neg p \rightarrow \neg K\neg p)$	$\mathbf{K_D}$, 7
6,7	(9)	$D(\neg\neg p \rightarrow \neg K\neg p)$	modus ponens, 6,8
6,7	(10)	$D\neg\neg p \rightarrow D\neg K\neg p$	$\mathbf{K_D}$, 9
11	(11)	$D(p \rightarrow \neg\neg p)$	$\neg\neg$-intro$_D$

[29] Barnett's biconditionals (16)-(17) are reminiscent of Wright's Evidential Constraint. Indeed, Barnett's argument turns out to be a simplified version of Wright's "Basic Revisionary Argument" against classical logic. Whereas Barnett thinks knowing $p \vee \neg p$ allows us to conclude that either p or $\neg p$ is knowable (and presumably, feasibly so) for vague p, Wright takes this conclusion to be a reductio against supposing excluded middle to ever be known.

[30] Consider the gross infelicity of hedging in one's knowledge reports simultaneously for both p and $\neg p$: *"I *kind of* know he's bald, *kind of* know he isn't", *"It's *sort of true* that I can know he's bald, *sort of true* that I can know he isn't", *"It's *roughly the case* that you can know he's bald, but also *roughly the case* that you can know he isn't", *"You can *sort of* tell he's bald, but you can also *sort of* tell he isn't", *"It's *hard to tell* if I know he's bald, yet equally *hard to tell* if I know he's not bald", *"It's *hard to say* if I know he's bald, also *hard to say* if I know he isn't", *"I'm *roughly* certain he's bald, *roughly* certain he isn't."

11	(12)	$Dp \rightarrow D\neg\neg p$	$\mathbf{K_D}$, 11
6,7,11	(13)	$Dp \rightarrow D\neg K\neg p$	transitivity, 10,12
1,2,6,7,11	(14)	$Dp \rightarrow DKp$	transitivity, 5,13

Given its commitment to UNCLEAR and CLEAR, Barnett's view inherits all the problems faced by Wright's quandary view of vagueness.

2.5 Dorr

Dorr has argued that the principle UNKNOWN has pragmatically objectionable consequences. Suppose you are asked "Is the glass pretty full?" of a borderline pretty full glass. Saying "I don't know" is apt to mislead your hearer into thinking you can't see the glass very well. In this way, pleading ignorance about the underlying status of borderline cases looks infelicitous.[31] Dorr concludes that principles claiming that vagueness entails ignorance, such as UNKNOWN, should be rejected on pragmatic grounds in the sense of having unassertible consequences.

Dorr's official substitute principle is:[32]

EXPORTATION	If it is knowable that it is not determinate that not-*p*, it is not determinate that it is not knowable that *p*

which entails the more concise:[33]

---

[31] The data is remarkably fragile. Notice that any appearance of infelicity with saying "I don't know", in response to being asked "*p*?" for some vague *p*, immediately vanishes, once followed up with some qualifier or hedge—"I don't know, it's vague" sounds perfectly fine. Oddly, this phenomenon is left entirely unaddressed in Dorr's discussion.

[32] I have modified 'unknown' and 'known' in Dorr's own formulations of UNKNOWN and EXPORTATION to be 'unknowable' and 'knowable', respectively, for reasons touched upon earlier (see also §6 below; see however n.39).

[33] Assume KIp. Given $\mathbf{K_E}$, we can derive $K(\neg Dp \,\&\, \neg D\neg p)$ by def(I), so both (i) $K\neg Dp$ and (ii) $K\neg D\neg p$ by &-elim. From (i) and $\mathbf{T_E}$ we get $\neg Dp$. Suppose DKp for reductio. By (the relevant determinate instance of) $\mathbf{T_D}$, we have $D(Kp \rightarrow p)$. Applying $\mathbf{K_D}$ gets $DKp \rightarrow Dp$. So Dp by modus ponens. Contradiction. Hence

EXPORT — If it is knowable that it is indeterminate whether *p*, it is indeterminate whether it is knowable that *p*

Dorr motivates EXPORTATION as follows.

> "What does it take to make oneself a counterexample to [UNKNOWN]? The case of Respondent [i.e. being asked of some borderline pretty full glass whether it is pretty full] suggests that for many substitutions for *P*, it is sufficient if one knows as much about the precise facts upon which the question whether *P* supervenes as any normal human being could know, has a normal grasp of the meaning of the English sentence '*P*', and meets a certain threshold of rationality and reflectiveness. But one doesn't have to know as much as this about the underlying precise facts to be a counterexample to [UNKNOWN]. The following conditional looks determinately true: if Respondent, who knows that the glass is between 60% and 70% full, knows that the glass is pretty full, so does a less opinionated counterpart of Respondent who knows only that the glass is between 60% and 90% full. If so, then since it is indeterminate whether Respondent knows that the glass is pretty full, it is also indeterminate whether her less opinionated counterpart knows that the glass is pretty full." (2003:§7)

The offered motivation for EXPORTATION is dubious. Even supposing it is determinately true that if Respondent knows *p* then her less informed counterpart also knows *p*, it does not follow from this alone that if it is indeterminate whether Respondent knows *p* then it is indeterminate whether her less informed counterpart knows *p*. All that immediately follows is the weaker claim: if it is not determinate that Respondent fails to know *p* then it is not determinate that her less informed counterpart fails to know *p*. For recall DEBATABLE: only failure of determinate falsity, or remaining *debatable*, as it were—and not full-blown indeterminacy—is closed over determinate implication.[34] To derive Dorr's conditional claim of shared indeterminate knowledge, one needs in addition to his conditional claim its determinate converse: it is

---

$\neg DKp$. Now EXPORTATION gives $K\neg D\neg p \rightarrow \neg D\neg Kp$. So $\neg D\neg Kp$ by modus ponens on (ii). Hence $\neg DKp \,\&\, \neg DK\neg p$, or IKp.

[34] Assume $D(A\rightarrow B)$. Given $D((A\rightarrow B)\rightarrow(\neg B\rightarrow\neg A))$ (i.e. assuming contraposition is 'determinately' valid), by $\mathbf{K_D}$ and modus ponens we have $D(\neg B\rightarrow\neg A)$. By $\mathbf{K_D}$ again, $D\neg B\rightarrow D\neg A$. By contraposition again, $\neg D\neg A\rightarrow\neg D\neg B$.

determinately true that if her less informed counterpart knows *p* then Respondent also knows *p*.[35] Granted, this converse claim is just as plausible, and equally (i.e. determinately) true. But then the conclusion of (conditionally) shared indeterminate knowledge is drawn entirely from independent principles of classical reasoning governing determinacy and knowability claims, without any reliance on principles linking the two such as EXPORTATION.[36]

Dorr's objection to UNKNOWN on grounds that responding to borderline questions (any "*p*?" for vague *p*) with "I don't know" appears to conflict with the very analysis of semantic indeterminacy which Dorr himself applies toward vague statements.[37] On that analysis, *p* is semantically indeterminate in context *c* for population *l* iff there is some truth *q* about *c* such that asserting *p* and denying *p* would each be individually permitted for any speaker in *c* who knew *q* (intuitively, all the facts relevant to evaluating whether *p*), according to the conventions of language use prevailing in *l*.[38] Yet Dorr's principle EXPORT predicts that any case where it is

---

[35] We already have $\neg D\neg A \rightarrow \neg D\neg B$, given $D(A\rightarrow B)$ (see n.34). Suppose we also have $D(B\rightarrow A)$. By $\mathbf{K_D}$, $DB\rightarrow DA$. Contraposing gets $\neg DA\rightarrow\neg DB$. Combining results gets us $IA\rightarrow IB$ (by the validity of $\phi\rightarrow\varphi$, $\psi\rightarrow\zeta \Rightarrow \phi\&\psi \rightarrow \varphi\&\zeta$).

[36] Nor is it obvious at all how Dorr's conclusion, which has the form $IK_Rp \rightarrow IK_Cp$, is meant to support instances of either EXPORTATION or EXPORT, whose predictions, of forms $K\neg D\neg p \rightarrow \neg D\neg Kp$ and $KIp \rightarrow IKp$ respectively, concern what a single individual knows (i.e. include only one knowledge operator relativized to a single subject), based on an antecedent condition of knowing (in)non-determinate facts, *not* (in)non-determinately knowing facts. In general it's unclear how any consideration about the supervenience base for indeterminacy facts, counterparts with weakened evidential bases, or the like is supposed to support an epistemic principle like EXPORT(ATION).

[37] Dorr (2003:§3) mentions another notion of semantic indeterminacy, one that *forbids*, rather than permits, both asserting *p* and asserting ¬*p*—which he claims plausibly applies to partially defined predicates with incompletely specified extensions (like Soames' (1999) 'smidget'; also Fine's (1975) 'nice_1', Foster's (1975) 'pearl', Sainsbury's (1991) 'child*'). By contrast, the permissive notion, he maintains, is found "arguably in actual vague languages".

[38] Left as such, the analysis is clearly inadequate. Having to know *all* the relevant facts is surely too strong: knowing merely *some* of the relevant facts should suffice for permission to assert *p*. Nothing in the analysis rules out trivial interpretations of *q*: for all it says, the relevant evidence *q* might consist in *p* itself or the permission to assert *p*, in which case anyone who, *per impossibile* (arguably if *p* is vague), knew *that* fact *would* be forbidden from asserting ¬*p* after all—thereby voiding the concept of semantic indeterminacy of

knowably indeterminate whether *p* will be such that it is indeterminate whether *p* is knowable. Given the symmetry of indeterminacy (i.e. indeterminacy whether *p* entails indeterminacy whether ¬*p*), in such a case it will therefore be indeterminate whether *p* is *un*knowable. Presumably, this means that where the indeterminacy of *p* is known by someone (and is therefore knowable), it will be indeterminate whether that person does *not* know *p*.[39] But according to Dorr's favored analysis of indeterminacy, so long as the speaker knows the relevant facts, she is perfectly entitled to assert "I don't know".[40] For although she fails to be determinately ignorant,

---

any application. A suitable notion of context might help rule out such unintended interpretations, but Dorr gives no exposition of such a notion. As such, the attempt to paraphrase away any loose talk of "all the relevant facts" ("weasel words", he claims, that are unfit for an acceptable analysis) in terms of truths about the context is to be found wanting. Indeed, the most obvious relevant facts will presumably not be about the context *at all*—at least, not in any sense proposed by contextualists in the philosophical literature on vagueness (Dorr cites Lewis, Kamp, Raffman, Soames and Graff in fn.6), for whom "context" denotes something broadly linguistic or psychological that can affect standards of evaluation when applying vague predicates. Someone's physical height must surely count (if anything does) among the "relevant facts" for evaluating tallness. But that has nothing to do with the context of judgment in any intended sense: whether record-keeping conversational scoreboards (Lewis 1979), salient background presuppositions (Kamp 1981), internal psychological states (Raffman 1994, 1996), mutually agreed upon stipulations (Soames 1999), or speaker-relative interests (Graff 2000). For further discussion on various construals of "context", see Shapiro (2006: ch.1§5,7).

[39] This only provably follows provided it is determinately true that if *p* is knowable then *p* is known (by the subject under consideration, i.e. who knows *p* is indeterminate). (Symbolizing 'S knows' as 'K_S': only given $D(K_S p \rightarrow Kp)$ *and* $D(Kp \rightarrow K_S p)$ does $IKp \rightarrow IK_S p$ follow; see n.35.) But this is plausible enough: surely, *some* individual will (be determinately such that they) know what's knowable concerning some (potentially vague) statement *p*. At any rate, this wrinkle only complicates my presentation. Dorr's own formulation of EXPORTATION in terms of 'known' clearly predicts that some claim *p* (known to be vague) is vaguely unknown.

[40] Why not apply the permissive notion of semantic indeterminacy to the claim of vague knowledge, rather than vague ignorance, since indeterminacy of *p* allows for *both* asserting *p* and denying *p* (so, presumably, asserting ¬*p*)? One might question whether a denial of "I know *p*" is enough to justify a full assertion of "I don't know *p*". The use of SYMMETRY ($IKp \rightarrow I\neg Kp$) to secure the claim of vague ignorance ($I\neg Kp$) is meant to circumvent this worry.

neither is she determinately knowledgeable—and where determinacy disappears, permission pervades.[41]

EXPORT(ATION) thus appears both unjustified and inconsistent with Dorr's own views on vagueness-related indeterminacy. Setting aside these worries, how exactly are these considerations supposed to refute UNKNOWN anyhow? Dorr concedes that any falsification of UNKNOWN will be indirect:

> "If it's not determinately the case that [the subject] doesn't know whether the glass is pretty full, counterexamples to [UNKNOWN]—cases where it's indeterminate whether P and also indeterminate whether a certain person knows that P—must be quite common […] I wouldn't want to suggest that the only way to be a counterexample to [UNKNOWN] is to satisfy the antecedent of this principle [EXPORTATION—i.e. to satisfy $K\neg D\neg p$]. It might be enough, for example, if one didn't determinately fail to satisfy the antecedent [i.e. for $\neg D\neg K\neg D\neg p$ to hold]." (2003:§7)

The idea is that principles like EXPORT(ATION) rival UNKNOWN by offering competing epistemic predictions. So long as the consequent of either EXPORTATION or EXPORT is satisfied—that is, $\neg D\neg Kp$ or IKp (given Ip, these are equivalent, and in turn imply $I\neg Kp$)—for some borderline case, there is a counterexample to UNKNOWN. Hence Dorr must be denying the very possibility of *definite* ignorance ($D\neg Kp$) about the underlying status of any borderline case,

[41] If pleading ignorance about vague matters is supposed to be infelicitous for wholly pragmatic reasons, apart from violating any general semantic conventions governing language use, this remains to be spelled out. At the very least, this is not the line of argument advanced by Dorr, who appears to think that answering "Yes" or "No" (without qualification) to borderline questions *does* in fact violate our conventional norms of language, and moreover is misleading only insofar as it does just that: "the fact that the conventions are what they are entails that one can typically avoid misleading one's interlocutors only by conforming to the conventions. (Perhaps this is because the latter fact is partly constitutive of the fact that the conventions are what they are.)." (§3) In that case, simply answering "No" to the borderline question "Do you know if *p*?" would violate linguistic convention (even if a more elaborate answer "No, I don't quite know if *p*" were still admissible). For possible pragmatic solutions, see for instance Brian Weatherson's "Vagueness and Pragmatics" (manuscript) and J.R.G. Williams' "On the Pragmatics of Vagueness" (manuscript).

whatever the epistemic status of that case's borderlineness—whether knowable *tout court* (KI*p*) or only vaguely knowable (IKI*p*). There may be ignorance, it just can't be definite.

Dorr's "no-ignorance" view of vagueness essentially rejects the possibility that for some *p* it is vague whether *p* and definitely unknowable that *p*. This is classically equivalent to saying:[42]

(7) If it is indeterminate whether *p*, it is not determinately unknowable that *p*

And because 'knowable' is definitely factive, we already independently have:[43]

(8) If it is indeterminate whether *p*, it is not determinately knowable that *p*

But now, (7) and (8) together entail UNCLEAR ($Ip \rightarrow IKp$). It appears that Dorr's account of vagueness does not fall far from those of other deniers of UNKNOWN like (certain interpretations of) Wright and Barnett. Yet discrepancies remain. Dorr (2009) explicitly rejects all principles, including Wright's DEF, that would license the unrestricted iteration of determinacy.[44] Without any such iteration principle for determinacy, it may appear that there is no way to validate CLEAR on Dorr's view. Might the no-ignorance theory of vagueness thus avoid the combination of UNCLEAR and CLEAR and its undesirable validation of the KK Principle within all contexts, determinacy and indeterminacy alike? There is reason to be doubtful.

Recall that *being debatable*, in the (admittedly artificial) sense of not being determinately ruled out (abbreviated "¬D¬…"), is closed over determinate implication:

DEBATABLE If it is determinate that if *p* then *q*, if *p* is debatable then *q* is debatable

---

[42] The step is even intuitionistically valid, since it only requires $\exists x(\ldots) \Rightarrow \neg\forall x \neg(\ldots)$ and $\neg(A\&B) \Rightarrow (A \rightarrow \neg B)$.

[43] Suppose for reductio I*p* but DK*p*. The relevant definitized instance of $\mathbf{T_E}$ is $D(Kp \rightarrow p)$. By $\mathbf{K_D}$, $DKp \rightarrow Dp$. So D*p* by modus ponens. But ¬D*p* by assumption (i.e. I*p*). Contradiction. Therefore $Ip \rightarrow \neg DKp$.

[44] His reasons there (2009:§7) concern the denial of anything being ultratrue (i.e. definitely^n true for any *n*).

Then we can reason as follows.

1	(1)	$\mathrm{IK}p \,\&\, \mathrm{IK}\neg p$	A
1	(2)	$\mathrm{IK}p$	&E, 1
1	(3)	$\neg\mathrm{DK}p \,\&\, \neg\mathrm{D}\neg\mathrm{K}p$	def(I), 2
1	(4)	$\neg\mathrm{D}\neg\mathrm{K}p$	&E, 3
5	(5)	$\mathrm{D}(\mathrm{K}p \rightarrow p)$	$(\mathbf{T_D})_{\mathrm{D}}$
6	(6)	$\mathrm{D}(\mathrm{K}p \rightarrow p) \rightarrow (\neg\mathrm{D}\neg\mathrm{K}p \rightarrow \neg\mathrm{D}\neg p)$	DEBATABLE
5,6	(7)	$\neg\mathrm{D}\neg\mathrm{K}p \rightarrow \neg\mathrm{D}\neg p$	modus ponens, 5,6
1,5,6	(8)	$\neg\mathrm{D}\neg p$	modus ponens, 4,7
1	(9)	$\mathrm{IK}\neg p$	&E, 1
1	(10)	$\neg\mathrm{DK}\neg p \,\&\, \neg\mathrm{D}\neg\mathrm{K}\neg p$	def(I), 9
1	(11)	$\neg\mathrm{D}\neg\mathrm{K}\neg p$	&E, 10
12	(12)	$\mathrm{D}(\mathrm{K}\neg p \rightarrow \neg p)$	$(\mathbf{T_D})_{\mathrm{D}}$
13	(13)	$\mathrm{D}(\mathrm{K}\neg p \rightarrow \neg p) \rightarrow (\neg\mathrm{D}\neg\mathrm{K}\neg p \rightarrow \neg\mathrm{D}\neg\neg p)$	DEBATABLE
12,13	(14)	$\neg\mathrm{D}\neg\mathrm{K}\neg p \rightarrow \neg\mathrm{D}\neg\neg p$	modus ponens, 12,13
1,12,13	(15)	$\neg\mathrm{D}\neg\neg p$	modus ponens, 11,14
16	(16)	$\mathrm{D}p$	A (for reductio)
17	(17)	$\mathrm{D}(p \rightarrow \neg\neg p)$	$\neg\neg$-intro$_{\mathrm{D}}$
17	(18)	$\mathrm{D}p \rightarrow \mathrm{D}\neg\neg p$	$\mathbf{K_D}$, 17
16,17	(19)	$\mathrm{D}\neg\neg p$	modus ponens, 16,18
17	(20)	$\neg\mathrm{D}p$	reductio, 15,16,19
1,5,6,17	(21)	$\neg\mathrm{D}p \,\&\, \neg\mathrm{D}\neg p$	&-intro, 8,20
1,5,6,17	(22)	$\mathrm{I}p$	def(I), 21

This shows that when both *p* and ¬*p* are vaguely knowable, it is vague whether *p*. But UNCLEAR and SYMMETRY already guaranteed the converse entailment: I*p* entails both IK*p* and IK¬*p*.[45] In other words, any borderline claim *p* will be marked by the indeterminate knowability of both *p* and ¬*p*. The result here confirms that the indeterminate knowability of both *p* and ¬*p* is an epistemic marker of *only* borderline cases: nothing definite shares this epistemic feature.

[45] Assume I*p*. By UNCLEAR, IK*p*. But by SYMMETRY, it also follows from I*p* that I¬*p*, hence IK¬*p* by UNCLEAR.

Equivalently, vagueness is exclusively marked by the indeterminate *un*knowability of both p and $\neg p$.[46]

Hence the derivable equivalence of Ip and IKp & IK$\neg p$. Now, at the first order of determinacy, there are only three options: either it is determinate that p, it is indeterminate whether p, or it is determinate that $\neg p$. These three possibilities {Dp, Ip, D$\neg p$} partition the various ways p may be or fail to be (first-order) determinate. Given that the option first-order indeterminacy (Ip) is now seen to be equivalent to indeterminate knowability for both p and $\neg p$ (IKp & IK$\neg p$), it follows that either other remaining option of first-order determinacy (Dp or D$\neg p$) is equivalent to the condition of determinate knowability (DKp or DK$\neg p$, respectively). Thus {DKp, IKp & IK$\neg p$, DK$\neg p$} partition the various ways either p or $\neg p$ may be or fail to be determinately knowable.

Hence the equivalence of Dp and DKp. This validates CLEAR, for we can now derive any instance of the form D$p \rightarrow$ DKp.[47] But then Dorr's no-ignorance view of vagueness incurs the same exact problems facing all the previous views committed to UNCLEAR and CLEAR.

---

[46] This follows from the equivalence of Iϕ and I$\neg\phi$ for any ϕ (including Kp). To demonstrate that, it suffices to show the equivalence of $\neg$Dϕ and $\neg$D$\neg\neg\phi$ for any ϕ (since the other conjunct $\neg$D$\neg\phi$ is shared by both Iϕ and I$\neg\phi$). *Pf.* (RtoL) $\neg$Dϕ is easily derivable from $\neg$D$\neg\neg\phi$ using reductio, $\neg\neg$-intro$_D$ (i.e. D($\phi \rightarrow \neg\neg\phi$)), $\mathbf{K_D}$, and modus ponens (in the manner of lines 15-20 above). (LtoR) The converse is classically derivable by similar means. Assume $\neg$Dϕ. By $\neg\neg$-intro$_D$, D($\neg\neg\phi \rightarrow \phi$). By $\mathbf{K_D}$, D$\neg\neg\phi \rightarrow$ Dϕ. By modus tollens (or contraposition and modus ponens), we get $\neg$D$\neg\neg\phi$. Note the essential use of double negation elimination. Since this rule is intuitionistically invalid, the logical intuitionist lacks the resources to derive I$\phi \rightarrow$ I$\neg\phi$ independently of any separate rule such as SYMMETRY. So much the worse for intuitionism, given the independent plausibility of SYMMETRY. Granted, some theorists do choose to reject SYMMETRY, but do so for reasons independent of intuitionism (e.g. Raffman 2005).

[47] Assume Dp. Then $\neg$Ip and $\neg$D$\neg p$. By IKp & IK$\neg p \Rightarrow$ Ip, from $\neg$Ip we have $\neg$(IKp & IK$\neg p$). From $\neg$D$\neg p$ we have $\neg$DK$\neg p$ (otherwise DK$\neg p$ implies D$\neg p$ by $\mathbf{K_D}$ and $\mathbf{T_E}$, contradicting $\neg$D$\neg p$). But provided a "completeness" constraint DK$p \vee$ (IKp & IK$\neg p$) $\vee$ DK$\neg p$ (i.e. either one of p or $\neg p$ is determinately knowable or both are indeterminately knowable—since, we are assuming, neither is ever determinately *un*knowable), we then have DKp.

2.6 Williamson

One might suppose my defense of UNKNOWN against its detractors constitutes a defense of the premier version of the epistemicist theory—Williamson's. In a way, yes. But not entirely.

Williamson in more than one place appears to take a quasi-eliminativist stance toward the notion of definiteness, claiming 'definitely' has no robust, defensible meaning apart from an epistemic construal:

> "If we cannot grasp the concept of definiteness by means of the concept of truth, can we grasp it at all? No illuminating analysis of 'definitely' is in prospect. Even if we grasp the concept as primitive, why suppose it to be philosophically significant? One can make sense of the supervaluationist apparatus [by which 'definitely' is defined in terms of admissible interpretations] if one assumes that an interpretation *s* admits an interpretation *t* just in case if *s* were correct then speakers of the language could not know *t* to be incorrect. On this view, 'definitely' means something like 'knowably'. Just one interpretation is correct, but speakers of the language cannot know all others to be incorrect. Vagueness is an epistemic phenomenon." (Williamson 1994:164)

and

> "Let it be that [some vague claim] is neither definitely true nor definitely false. In reporting this obvious truth, the philosopher has no right to stipulate a theoretical sense for 'definitely'. Rather, it must be used in a sense expressive of what is obvious. Yet what is *obvious* is just that vague sentences are sometimes neither knowably true nor knowably false. The simplest hypothesis is that this is the *only* sense in which the vague sentences are neither definitely true nor definitely false." (Williamson 1992:150–1)

The suggestion is that standard understandings of 'definitely' or 'determinately' (or any other cognate used to introduce and give expression to the notion of being a borderline case), construed in broadly supervaluationist terms of *semantically admissible* interpretations, should be *replaced* by an epistemic understanding—one perhaps (as in Williamson's own account) given in terms of *epistemically possible* interpretations.[48] Talk of determinacy, as far as vagueness is concerned, is

[48] In this vein, Williamson proposes a logic of clarity, detailed in the Appendix of his (1994), in which the 'clarity' operator obeys a **KTB** modal logic. His (2005:§9) discussion reiterates this idea that epistemicism

on this suggestion to be reinterpreted as, or otherwise understood entirely in terms of, talk of knowability. This falls short of any reductive aspirations for 'determinately', since presumably not everything determinate is knowable.[49] Yet it comes close to saying the notion of (in)determinacy reduces to that of something's being (un)knowable.

What the suggestion is *not*—it is worth pointing out—is to give a bridging principle connecting the notions of definiteness and knowability (or indeterminacy and unknowability), in the manner of UNKNOWN (or UNCLEAR$^{(*)}$, EXPORT, etc.). For this would be to treat the two operators, 'definitely' and 'knowably', as expressing distinct notions—contrary to the task of reinterpreting one in terms of the other. What Williamson wants, by contrast, is for 'definitely' to be completely reinterpreted as meaning "something like 'knowably'".

Although elsewhere, in his (1995) Williamson does appear to leave room for keeping the two notions of determinacy and knowability (at least conceptually) distinct. There, he defends the principle:

DETERMINATE If it is knowable that *p*, it is determinate that *p*

which in turn validates UNKNOWN.[50] However, nowhere in that discussion or elsewhere does he offer any explicit formulation for a converse principle articulating the epistemic consequences for *definiteness*. Instead, Williamson (1995:175–6) appears to be resigned to letting determinacy be only *partially* characterized in terms of epistemic conditions: knowability is a sufficient but not

---

"can take over and reinterpret the formal apparatus of supervaluationism" originally used to define the 'definitely' operator.

[49] The favorite example Williamson and Wright, etc. all return to is: Goldbach's (as of yet, unproven) Conjecture.

[50] By DETERMINATE, $Kp \rightarrow Dp$ and $K\neg p \rightarrow D\neg p$. Contraposing gets $\neg Dp \rightarrow \neg Kp$ and $\neg D\neg p \rightarrow \neg K\neg p$. Hence $\neg Dp \;\&\; \neg D\neg p \rightarrow \neg Kp \;\&\; \neg K\neg p$, or $Ip \rightarrow \neg Kp \;\&\; \neg K\neg p$.

necessary condition for determinacy, while unknowability is a necessary but not sufficient condition for indeterminacy. One might have thought the *epistemic theory* of vagueness could do better. (It is a *theory* after all, is it not?) Can some principle complete the account? Our earlier discussion suggests a natural candidate:

> KNOWN If it is determinate that *p* (and vagueness is the only potential source of ignorance preventing one from knowing whether *p*), it is knowable that *p*

Together, KNOWN and UNKNOWN provide full necessary and sufficient conditions for determining whether any given case is one of (in)determinacy, where these are specified completely in terms of epistemic conditions of (un)knowability, on the assumption that vagueness is the only potential source of unknowability for the claim in question. The fact that this last qualification mentions both *analysans* and *analysandum* (if even alleged as such) disqualifies the account from being a proper, noncircular analysis. But this is as it should be—the epistemic characterization of vagueness was never meant to reduce the notion of (in)determinacy to that of (un)knowability, any more than Williamson's endorsement of DETERMINATE strove toward full reductionism.

Any reductionist epistemic conception about vagueness, at any rate, should be rejected—and arguably along with any attempt to reinterpret the 'definitely' operator as meaning ("something like") 'knowably'. For both ambitions belie the real intent behind the epistemicist theory, since they presuppose that epistemicism somehow conceives of vagueness as a special sort of ignorance. Strictly speaking, it does not. Vagueness is no more a form of ignorance or unknowability, than determinacy is a form of knowledge or knowability. Rather, according to the epistemicist theory (at least Williamson's version), vagueness is a special *source* of ignorance, the absence of which (i.e. determinacy), all else equal, enables knowledge. Vagueness is not so much an epistemic phenomenon, as it is an epistemically *characterized* phenomenon, whose true nature remains to be determined. Indeed, on Williamson's account, the underlying nature of vagueness turns out to be *semantic*. For Williamson claims that for any purportedly vague expression, it is its *meaning instability*—that is, the hyper-sensitivity of its meaning to the humanly undetectable minutiae of the term's overall use patterns among normal competent speakers within the broader linguistic community—that accounts for its ever-shifting, humanly unknowable extension, and also therefore the unknowability of the underlying status belonging to any borderline case.[51]

2.7 Greenough

The foregoing discussion provides some structure to situate the other theories so far considered, all of which agree that vagueness results in some sort of epistemic deficit or another. Dorr essentially agrees with Williamson that the source of epistemic deficiency is semantic, but denies there is any underlying semantic fact of the matter about linguistic conventions to adjudicate the status of borderline cases. Alternatively, the source of epistemic deficiency can itself be seen as *epistemic* in nature. Bobzien claims the epistemic deficiency, although true of ordinary human subjects, really consists in the inability of *idealized* informed, competent, truthtelling speakers to

[51] Some points worth noting: (i) The extension for 'F' need not in principle be complete, in the sense of preserving bivalence—i.e. anything truly satisfies either 'F' or 'not-F'—although Williamson (notoriously) defends this claim. (ii) The explanation for ignorance or unknowability on Williamson's view rests on his margin for error account of knowledge (really, a special case of his safety condition for knowledge; see his (2000)). (iii) Williamson identifies a vague term's overall use patterns with speakers' dispositions to either assent to or dissent from applying it in varying circumstances (although exactly how these aggregate and contribute to the collectively-determined extension of the term will, even at the individual level, be non-algorithmic and irremediably unsurveyable). In this way, the (meta)semantics for vague terms proves ultimately to be a matter of collective psychology. So perhaps the nature of vagueness remains *epistemic* on the account after all (although not distinctively so, given how semantics and psychology are intertwined).

arrive at a knowledgeable verdict about vague matters. Wright claims the deficiency is a broader epistemic phenomenon of quandary. Although he does not specify the nature of this phenomenon, the self-referential character of its undecidable undecidability (i.e. that whether any case of it is what it is, is itself cast in doubt) perhaps explains why its nature would prove rather elusive. Barnett refuses to reduce things at all, whether to epistemic or semantic terms, since the source of the epistemic deficiency is for him *sui generis* in nature. Greenough simply chooses to leave the nature of the phenomenon unspecified (at least within the "minimal theory"), maintaining only that vague predicates must exhibit epistemic deficiency in some cases.

Uncovering the roots of vagueness-related epistemic deficiency, ignorance or otherwise, raises a host of questions. Is it knowable that *p* if it is vague whether *p*? If not, how bad is the epistemic failure—must *p* remain unknown, or vaguely unknown, or unknowably unknown, etc.? What sort of epistemic entitlement is afforded by the determinacy of *p*—knowledge of *p*, determinate knowledge of *p*, knowable knowledge of *p*, etc.? This set of issues concerns the *epistemology* of vagueness.

By contrast, determining whether bivalence, excluded middle, and other features of classical logic are preserved when reasoning about borderline cases arguably poses a separate group of issues, ones that are *metaphysical* in nature. Are there cutoffs to vague predicates, concepts, properties, etc.? Is there an underlying truth or fact of the matter about borderline cases? about borderline borderline cases, or borderline definite cases, etc.? These questions, as we saw, generate their own array of answers. Some believe the truth structure for the underlying facts in vague matters remains entirely classical (Barnett, Williamson) or at least concede that vague matters admit of classical reasoning (Bobzien, Greenough), while others may choose to selectively reject the classical principles of bivalence and excluded middle (Dorr) or otherwise undertake the wholesale rejection of classical logic and semantics altogether (Wright).

The absence of any clear split or alignment along positions in both sets of issues—what to say regarding the truth status of borderline claims vs. how to conceive of the nature of the epistemic deficiencies presented by vagueness—gives some reassurance that the two domains of inquiry, epistemological and metaphysical, should indeed be kept distinct.

What then of the prospects for a minimal theory, in the sense of Greenough (2003)? The quest for seeking any universal agreement on the set of metaphysical issues is widely recognized by now to be utterly hopeless (as any survey of the literature on vagueness will quickly confirm). But maybe hope remains for finding a fixed point of consensus on the epistemology. We may concede that UNKNOWN cannot serve as a minimal theory of vagueness, since it is neither minimal, as evinced by skepticism from theorists like Barnett, Dorr, and Wright, nor a theory, since it states only necessary (but insufficient) epistemic conditions for vagueness. Yet the claim that vagueness has *some epistemic consequence* is incontrovertible—no one denies that. To think that vagueness could somehow be epistemically inconsequential is sheer madness. Perhaps then the dictum that "Ignorance is a consequence of vagueness" may serve as our minimal claim about vagueness. Stated loosely, it simply says that ignorance is present, in some capacity or another, whenever something is a borderline case. (Note that this is stronger in scope than Greenough's own minimal theory, which predicts merely that ignorance will be present in *some*—hence not necessarily every—borderline case.)

How to unpack this notion of "consequence" turns out to be a source of contention. There is some disagreement over whether to construe ignorance as a *de facto* consequence (Greenough, Williamson (dipustably)), as a *vague* consequence (Wright (first reading), Barnett), as a *qualified* consequence only when the case in question is definite borderline (Wright (second reading)), as an *unknowable* consequence (Wright (third reading), Bobzien), or as an *epistemically qualified* consequence only when the case in question is knowably borderline (Dorr). There is no disagreement, however, that some sort of epistemic deficiency is somehow intimately tied to the notion of being a borderline case. That deficiency is presumably gone when things are not vague. Although how exactly to construe this is equally contested: it is dispute whether the knowability is a *de facto*, determinate, qualified, knowable, or otherwise restricted consequence of determinacy.

That said, UNKNOWN and KNOWN, as I have argued, remain the most promising way of explicating the dictum, roughly put, that "Ignorance is a consequence of vagueness". Crucially, the arguments I gave against the other contender methods for developing this idea made no question-begging assumptions about the metaphysics of borderline cases—they did not, for instance, rest on classical rules of reasoning that are intuitionistically unjustified. Those results were established on independent grounds. In that way, the plausibility of concluding that vagueness entails ignorance rests on no questionable metaphysical commitments (to cutoffs, bivalence, or any Williamsonian black-box metasemantics), at least not any that would already decide matters in advance of weighing the epistemic considerations. But if this is right, the original ignorance-entailing conception of vagueness, even if it doesn't itself count as a minimal account of vagueness, nonetheless proves to be the best candidate for realizing what *is* a minimal account of vagueness.[52] And that is reason enough to endorse it.

---

[52] I say "minimal *account*", because the label "theory" is often reserved for *non-circular* analyses and the analysis via UNKNOWN and KNOWN (or their paired variants), while nontrivial, remains circular (albeit of the acceptable kind). The use of "analysis" is less controversial, since both necessary and sufficient conditions are fully provided.

Chapter 3: 'Vague' and Higher-Order Vagueness

3.1 Hyde's argument

Sorensen's (1985) argument exploits the vagueness of ordinary predicates to induce a sort of vagueness in 'vague'. Consider 'small'. It is vague insofar as it admits of border cases: 0 clearly is small and 10^6, say, clearly isn't, but in between there is no clearly last small integer. Hence the sorites argument for 'small':

(1) 1 is small
For any integer *n*, if *n* is small, then *n*+1 is small
Therefore, 10^6 is small

Sorensen's series is constructed by defining predicates of the form 'n-small' for every *n* as follows:

(2) *x* is *n*-small iff *x* is small or $x < n$

Clearly, '1-small' is vague: it shares exactly the same border cases with 'small' above 0. When *n* is clearly no longer small, say $n=10^6$, then 'n-small' is clearly precise because its extension is completely determined by the precise 'less than *n*' condition: anything less than 10^6 is *n*-small, anything greater is not. Since there is no clearly last vague '*n*-small' predicate, we now have a sorites argument for 'vague':

(3) '1-small' is vague
For any integer *n*, if '*n*-small' is vague, then '*n*+1-small' is vague
Therefore, '10^6-small' is vague

Hence the vagueness of 'vague'. Hyde (1994) takes this to show that any vague predicate is higher-order vague, by the following argument.[1] According to the "paradigmatic conception" of vagueness:

(4) '… is vague' means 'there are border cases of …'

But Sorensen's sorites setup appears to demonstrate that:

(5) The predicate '… is vague' is itself vague

From (4) and (5), any vagueness appearing in the analysandum must appear in the analysans:

(6) 'there are border cases of …' is vague

But assuming 'there are …' is not vague, the vagueness of the analysans must appear elsewhere, namely

(7) 'border cases of …' is vague

which, analyzed via (4), means

(8) There are border cases of 'border cases of …'

Thus, any predicate that satisfies '… is vague' will admit of "higher-order" border cases. That is,

(9) Any vague predicate has border border cases

Tye (1994) disputes the validity of the argument. The argument's premises, he claims, do not appear sufficiently general to establish its universal conclusion that *every* vague predicate has

[1] The reconstruction largely follows Varzi's (2003), with minor amendments. I have simplified 'can be analyzed' to 'means' in (4) to exclude other possible analyses (although all (4) really requires is coextension, not synonymy or meaning reduction), included the copula in ' '… is vague' ' in (5) (without which, the phrase would denote an adjective instead of an actual predicate), pluralized 'border case[s] of…' in (7) to match the previous premise, and chosen to stick with Hyde's original conclusion (9), which makes a claim about all, not just some, vague predicates.

border border cases. All Hyde's argument shows, if anything, is that such higher-order border cases appear *somewhere* among the vague predicates. That is because (7) and (8) are naturally read as implying only the existence of higher-order border cases for *some*, rather than *every*, vague predicate. Otherwise, if (7) really were true of every vague predicate, so would its preceding premise (6) universally apply to all vague predicates—which is absurd, since many vague predicates clearly do have border cases. Indeed, as shown by (5), all Sorensen's construction guarantees is the existence of *some* border case for '… is vague'.

Varzi (2003) agrees with Tye that Hyde's original conclusion is too strong. In contrast, I believe the argument can be suitably modified to restore full generality to its conclusion. Part of the worry, I take it, is that Sorensen's setup generates higher-order border cases using only a single vague predicate, 'small'. Without being told how to extend Sorensen's argument to other vague predicates, however, one has little reason to believe this consequence will hold for all vague predicates in general. Fortunately, there is an obvious extension to the Sorensen construction. For any vague predicate 'F' with a sorites series consisting of items arranged in decreasing order of Fness, for any item *n* we may define '*n*-F' as such:

(10) *x* is *n*-F iff *x* is F or *x* is more F than *n*

and construct a corresponding sorites series for 'vague' using the newly defined '*n*-F' predicates. In this way, for many other vague predicates 'F' besides 'small', we can construct a series of Sorensen-style predicates with at least one border case of 'vague'.[2] So perhaps Tye's generality objection is not insurmountable after all.

---

[2] Note that *x* and *n* can range over individuals of any sort, so the result is not limited to just numerical predicates defined over integers (e.g. *x* is Bob-tall iff *x* is tall or *x* is taller than Bob). Nor does it require that the relevant sorites sequence we start with be enumerable (although any resulting Sorensen sequence must still be enumerable, since defining each '*n*-F' predicate requires a name for the item *n* and there can

Or is it? Unfortunately, the argument, left in its current formulation, fails to capture any generality of consequence that Sorensen's sorites might have to offer. The problem lies not in its conclusion (9), as Tye and Varzi suggest, but rather in the inference from (6) to (7). The problem is that premises like (5) and (6) are too unspecific: all they entail is the existence of some border case or another for 'vague', without specifying exactly which vague predicate 'F' was used to construct the Sorensen sequence guaranteeing that existence. One may of course wish to read (7) as claiming that 'border cases of 'F' ' is vague for some particular vague 'F'. But this can't follow from (6), since the existential content of (6) is too weak to recover any particular information about the original predicate.

Besides, there is independent reason to object to Hyde's argument in its current presentation. One obvious major deficiency is how it attempts to apply the paradigmatic conception of vagueness—an analysis presumably meant for claims of *predicate vagueness,* since it is only predicates, strictly speaking, that admit of border cases—to several other types of expressions, like sentences ('there are border cases of …'), noun phrases ('border cases of …'), and quantifiers ('there are …')—expression types for which the notion of a border case either appears incoherent (at worst) or remains undefined (at best).

The difficulty is brought out further by Tye's criticism that Hyde's argument conflates two distinct notions: vague vagueness and higher-order vagueness. A predicate is *vaguely vague* insofar as it is indeterminate whether it has border cases. This is not the same as being *higher-order vague*, that is, having border border (border…) cases. Perhaps, Tye suggests, Sorensen's setup rests solely upon the vague vagueness of ordinary vague predicates like 'small', in which

---

only be denumerably many of these). This is important for handling non-denumerably infinite sorites sequences, like phenomenal sorites (e.g. color continua).

case Hyde's argument does nothing to demonstrate the existence of higher-order vagueness. Tye does not indicate exactly where he thinks this would invalidate Hyde's argument, only that it casts doubt on Hyde's conclusion. The suspect however is unmistakable. Whereas (6) appears to make a claim about vague vagueness (i.e. indeterminacy about *having* a border case), (7) appears to make a claim about higher-order vagueness (i.e. indeterminacy about *being* a border case). As before, the discrepancy between (6) and (7) remains unjustified.[3] The scope difference between claims of vague vagueness ('it is indeterminate whether there exists a border case of…') and claims of higher-order vagueness ('there exists a case such that it is indeterminate whether it is a border case of…') is starkly apparent. Yet no party in the debate—neither Tye in his criticism, nor Varzi in his endorsement, nor Hyde in his (2003) reply—has bothered to address this outstanding issue. Nor, apparently, have any outsiders. Even Crispin Wright is dismissive about sorting out possible connections:

> "It seems obvious enough that there is little connection between [the thesis that 'vague' is vague] and [the existence of higher-order vagueness, as expressed by means of a 'definitely' operator, as in e.g. 'borderline borderline…' or 'borderline definitely…']. It seems quite consistent with holding to the Buffering view [i.e. that there exist borderline borderline cases], or with thinking of "Definitely P" as vagueness-inheriting though precision-increasing when applied to a vague claim P, that the notion of vagueness itself should divide all expressions into two sharply bounded kinds—that there is never any vagueness about the question whether an expression is vague or not. Conversely, one might think of the distinction between vague expressions and others as admitting of borderline cases but hold to a view of the nature of vagueness according to which there are no higher-order borderline cases; and one might simultaneously repudiate any operator of definiteness, or take the view that any legitimate such operator generates only precise claims. At any rate, these are all prima facie compatibilities. If there are deeper tensions, that would be interesting—but they remain to be brought out." (Wright 2010, p.531)

[3] Hyde appeals to Rolf's (1980:§3) *Inheritance Principle*: If all the constituent phrases of a complex phrase are precise then the complex phrase is precise. But if I am correct, the principle is less innocent than one might suspect.

Evidently, things "remain to be brought out". In the rest of this comment, therefore, I wish to clarify the relations between vague vagueness and higher-order vagueness and restate Hyde's argument in a way that avoids the problems so far discussed.

3.2 From vague vagueness to higher-order vagueness

To abbreviate claims about border cases, let us introduce the standard operators, I: 'it is indeterminate whether …' and D: 'it is determinate that …'.[4]

Let us assume the determinate truth of the biconditional defining any predicate '*n*-small':

(11) D(*x* is *n*-small iff *x* is small or $x < n$)

Given the preservation of indeterminacy across determinate equivalence,[5] we can derive the conditions under which something is a border case of '*n*-small':

(12) I(*x* is *n*-small) iff I(*x* is small or $x < n$)

Notice that the right-side condition fails if *x* is (determinately) less than *n*. Hence (12) reduces to:[6]

---

[4] One may still define indeterminacy in terms of determinacy ($\mathrm{I}p :=_{def} \neg\mathrm{D}p \,\&\, \neg\mathrm{D}\neg p$), or vice versa ($\mathrm{D}p :=_{def} p \,\&\, \neg\mathrm{I}p$), so long as the reducing notion is understood in terms of border cases, not vagueness (e.g. letting 'determinately, *x* is F' mean '*x* is F and *x* is not a border case of Fness' or letting 'it is indeterminate whether *x* is F' mean '*x* is a border case of Fness', respectively). Otherwise, the paradigmatic conception of vagueness would be circular.

[5] Let us suppose the determinacy operator 'D' obeys rule **K**: $\mathrm{D}(p \rightarrow q) \rightarrow (\mathrm{D}p \rightarrow \mathrm{D}q)$. First consider $\mathrm{D}(p \rightarrow q)$. By **K** and contraposition (i.e. distributing 'D' over $\mathrm{D}((p \rightarrow q) \rightarrow (\neg q \rightarrow \neg p))$), we can infer $\mathrm{D}(\neg q \rightarrow \neg p)$. Distributing again gets $\mathrm{D}\neg q \rightarrow \mathrm{D}\neg p$, which by contraposition is (i) $\neg\mathrm{D}\neg p \rightarrow \neg\mathrm{D}\neg q$. Now consider $\mathrm{D}(q \rightarrow p)$. By the same principles, distributing then contraposing delivers (ii) $\neg\mathrm{D}p \rightarrow \neg\mathrm{D}q$. Combining (i) and (ii) yields $\mathrm{I}p \rightarrow \mathrm{I}q$. $\mathrm{I}q \rightarrow \mathrm{I}p$ follows by symmetry of reasoning; hence $\mathrm{I}p \leftrightarrow \mathrm{I}q$. But assumptions $\mathrm{D}(p \rightarrow q)$ and $\mathrm{D}(q \rightarrow p)$ are just $\mathrm{D}(p \leftrightarrow q)$.

[6] (LtoR) Assume $\mathrm{I}(\mathrm{A} \vee \mathrm{B})$ and $\mathrm{DB} \vee \mathrm{D}\neg\mathrm{B}$. Suppose DA for reductio. By **K** and $\vee$-introduction, $\mathrm{D}(\mathrm{A} \vee \mathrm{B})$, which contradicts $\mathrm{I}(\mathrm{A} \vee \mathrm{B})$. So $\neg$DA. By parity of reasoning we get $\neg$DB. Hence $\mathrm{D}\neg\mathrm{B}$ by $\vee$-syllogism. Suppose $\mathrm{D}\neg\mathrm{A}$ for reductio. Since determinacy collects over conjunction (because 'D' distributes over

(13) $I(x \text{ is } n\text{-small})$ iff $I(x \text{ is small})$ and $x \geq n$

Therefore, the border cases of '*n*-small' are exactly those border cases of 'small' not less than *n*. That such cases exist just is what it means for '*n*-small' to be vague on the paradigmatic conception. That is,[7]

(14) $\exists x\, I(x \text{ is } n\text{-small})$ iff $\exists x\, (I(x \text{ is small}) \text{ and } x \geq n)$

This says that '*n*-small' is vague iff there is some border case of 'small' that is greater than or equal to *n*.

What about the conditions for '*n*-small' being *vaguely* vague? for 'vague' being vague? Further variations on (13) are derivable by iterating existential quantifiers and indeterminacy operators, such as[8]

(15) $I(\exists x\, I(x \text{ is } n\text{-small}))$ iff $I(\exists x\, (I(x \text{ is small}) \text{ and } x \geq n))$

This says that '*n*-small' is vaguely vague iff it is indeterminate whether there is some border case of 'small' greater than or equal to *n*. Since it is one of these '*n*-small' predicates within the Sorensen sequence that is supposed to account for the vagueness of 'vague', we can now state,

---

$D(A\rightarrow(B\rightarrow(A\&B)))$ via **K**), we get $D(\neg A\&\neg B)$. By **K** and de Morgan, this becomes $D\neg(A\vee B)$, which contradicts $I(A\vee B)$. So $\neg D\neg A$. Hence IA. (RtoL) Assume IA and $D\neg B$. Now $D(A\vee B)$ would imply DA (by ∨-syllogism and **K**), contrary to IA; so $\neg D(A\vee B)$. And $D\neg(A\vee B)$ would imply $D(\neg A\&\neg B)$, thus $D\neg A$ (by de Morgan and **K**), contra IA; so $\neg D\neg(A\vee B)$. So $I(A\vee B)$.

[7] One way to obtain this is to quantify each side of (13). Someone worried about quantifying into indeterminacy contexts may simply read (13) (and each of its preceding premises) as an implicitly universally quantified sentence of the form $\forall x(A(\ldots x\ldots)\leftrightarrow B(\ldots x\ldots))$. The inference to (14) merely weakens this to $\exists xA(\ldots x\ldots)\leftrightarrow\exists xB(\ldots x\ldots)$.

[8] To derive this we assume that (i) we may strengthen (11) to be determinate; (ii) determinacy is closed over entailment (so that (14) is made determinate too); (iii) indeterminacy is preserved across determinate equivalence (as in (12)). Such a strengthening is surely permissible, since the entire deduction rests only upon a definitional truth.

somewhat more precisely than in Hyde's original argument (or even later reconstructions), the exact requisite conditions.

(16) $\exists n$ I($\exists x$ I(x is n-small)) iff $\exists n$ I($\exists x$ (I(x is small) and $x \geq n$))

For 'vague' to be vague, it suffices that for some '*n*-small' predicate, it is indeterminate whether it admits of any border cases—i.e. it is indeterminate whether anything is *n*-small—which by (16) comes to the requirement that it be indeterminate whether anything greater than or equal to some *n* is a border case of 'small'. Note that this is distinct from requiring that something greater than or equal to some *n* is a *border* border case of 'small' (i.e. $\exists n \exists x$ I(I(x is small) and $x \geq n$)). This is because, as we shall see, it is not true in general that the indeterminacy operator commutes with existentials. As such, claims of vague vagueness having the form I∃x(IFx) are not equivalent to claims of higher-order vagueness having the form ∃x(IIFx).

3.3 Quantifiers and indeterminacy

Does Hyde's argument equivocate between I∃x(IFx) and ∃x(IIFx)? Hyde takes himself to address "the problem of higher-order vagueness", by which he means the existence of borderline borderline cases, rather than the borderline existence of borderline cases.[9] But, as (16) shows, it isn't clear that he is entitled to this assumption. Not without further justification, at least. The controversial inference moves from claims of vague vagueness to claims of borderline borderlineness:

---

[9] Following standard parlance, I shall use 'borderline'/'border' and 'definitely'/'determinately' interchangeably.

($I\exists I \Rightarrow \exists II$) If it is indeterminate whether borderline Fs exist, there exist borderline borderline Fs

which is an instance of the more general problematic inference pattern:[10]

($I\exists \Rightarrow \exists I$) ?If it is indeterminate whether Fs exist, there exist borderline Fs

Notice that the converse move of exporting an indeterminacy operator outside the scope of an existential is clearly fallacious, since the mere possibility of borderline borderline cases *plus* definite borderline cases shows the former alone does not logically entail vague vagueness.[11] Hence the failure of:

($\exists II \Rightarrow I\exists I$) *If there exist borderline borderline Fs, it is indeterminate whether borderline Fs exist

And given that most vague predicates do admit of both determinate and borderline cases, this fails too:

($\exists I \Rightarrow I\exists$) *If there exist borderline Fs, it is indeterminate whether Fs exist

---

[10] To appreciate just how hard it is to justify such inferences, consider the following strategy. To show $I\exists xFx \rightarrow \exists xIFx$, it suffices to derive the contrapositives of (i) $\neg D\exists xFx \rightarrow \exists x\neg DFx$ and (ii) $\neg D\neg\exists xFx \rightarrow \exists x\neg D\neg Fx$. For (i): $\neg\exists x\neg DFx$ *iff* (*by duality*) $\forall xDFx$ *only if* (*by K, $\exists$-Intro*) $\forall xD(\exists xFx)$ *only if* (*by vacuous $\forall$*) $D\exists xFx$ *only if* (*by $\neg\neg$-Intro*) $\neg\neg D\exists xFx$. For (ii): $\neg\exists x\neg D\neg Fx$ *iff* (*by duality*) $\forall xD\neg Fx$ *only if* (*by ??*) $D\forall x\neg Fx$ *iff* (*by duality*) $D\neg\exists xFx$ *iff* (*by $\neg\neg$-Intro*) $\neg\neg D\neg\exists xFx$. The critical missing entailment ?? resembles the Barcan Formula in modal logic. It is, however, arguably licensed on a supervaluationist reading of the determinacy operator 'D' as a universal quantifier ranging over admissible precisifications, which then must surely commute with the standard first-order universal quantifier '$\forall$'. (Perhaps this is why Varzi (2005), an avowed supervaluationist, accepts Hyde's argument as sound.)

The reverse "converse Barcan" inference from $D\forall x\neg Fx$ to $\forall xD\neg Fx$, like its modal analogue, is more plausible. *Pf.* Assume for reductio (i) $D\forall x\neg Fx$ but (ii) $\neg\forall xD\neg Fx$. Then $\exists x\neg D\neg Fx$ by (ii). Suppose $\neg D\neg Fa$ (as a witness). Since $\neg Fa$ "determinately" follows from $\forall x\neg Fx$, we have $D(\forall x\neg Fx \rightarrow \neg Fa)$. By **K** and (i), $D\neg Fa$. Contradiction.

[11] Since $\exists xIIFx$ and $\exists xDIFx$ are compatible and the latter already entails $D\exists xIFx$, the former can't entail $I\exists xIFx$.

In contrast, (I∃I ⇒ ∃II) does not admit of trivial refutation. Yet it remains contentious. Just because the very existence of borderline cases is vague, why should there thereby be borderline borderline cases?

Here is one way of providing the missing justification. Consider the *iterative conception* of vagueness, according to which any vagueness-related indeterminacy iterates.[12] According to the paradigmatic conception, the relevant sense of indeterminacy is borderlineness. Together, these predict that anything that is a borderline case of 'F' is a borderline borderline case of 'F'. This is equivalent to saying that there are no definite borderline cases of 'F'.[13] If true, this would guarantee the existence of higher-order borderline cases, assuming there are first-order borderline cases.

Can vagueness be shown to iterate in this way? Perhaps so. Consider any arbitrary predicate 'F'. Suppose 'F' has definite borderline cases. Thus, for some individual *a*, *a* is a definite borderline case of 'F'. That is to say, it is definitely the case that *a* is borderline F. Therefore, it is definitely the case that *something* is borderline F. But in that case, it cannot be vague whether anything is borderline F.

Consequently, the existence of definite borderline cases implies the definite—hence, non-vague—existence of borderline cases.[14] That is, any predicate 'F' that has definite borderline

---

[12] Versions of this thought may be found in (one interpretation of) Wright (2001, 2003) and Bobzien (2010).

[13] That is, assuming 'D' obeys **T**: D*p* → *p*, which is uncontroversial (if anything is, about the logic of vagueness). *Proof.* Since ¬∃xDIFx *iff* ∀x¬DIFx (by duality), it suffices to show IFa → IIFa *iff* ¬DIFa for some arbitrary a. Thus, assume (for reductio) both IFa → IIFa and DIFa. By **T**, IFa. By modus ponens, IIFa, so ¬DIFa. Contradiction.

[14] Formally: Assume ∃xDIFx. Thus, DIFa for some a. Generalizing within the scope of 'D', we get D∃xIFx, hence ¬I∃xIFx. (This last step applies **K** to D(IFa → ∃xIFx)—which is the object-language way

cases is definitely, not vaguely, vague. Contrapositively, anything vaguely vague lacks definite borderline cases. Given our earlier observations, this consequence of vague vagueness—the absence of definite borderline cases—guarantees that borderlineness (relative to 'F') iterates: anything borderline can only be borderline borderline, if its own borderline status fails to be determinate. Assuming there are in fact cases of borderlineness, it immediately follows that there will be cases (indeed, the very same ones!) of borderline borderlineness.

Everything therefore hinges on the existence of—plain, old, unexciting—*first-order* borderlineness. But is that not exactly what is at issue in alleged cases of vague vagueness—whether there are any (first-order) borderline cases to begin with? Although the situation might appear paradoxical, it really isn't. The idea that something can admit borderline cases, but only vaguely so, remains coherent. That borderline cases exist can nonetheless be true, even without being definitely so. One can accept the claim that borderline cases exist without assenting to its determinacy; one can therefore accept the possibility that borderline cases exist, but only vaguely.[15] This is coherent, so long as one does not *determinately* assert both claims at once—i.e. assert their joint determinate truth—which *would* lead to contradiction.[16] Hence, the compatibility of vagueness and vague vagueness.

Similarly, there is nothing paradoxical about the idea of borderline cases existing, but only as borderline borderline (and at that, only as borderline borderline borderline, and at that…). But how, one might ask, can the borderlineness of a borderline case, at any given order, possibly be

---

of expressing the "determinate" validity (i.e. the determinacy-preserving nature) of the rule of ∃-generalization.)

[15] One can easily derive $p \rightarrow \neg \mathrm{D} \neg p$ (using **T** and reductio). Thus, if p and $\neg \mathrm{D}p$ are true (consistent), so are p and $\mathrm{I}p$.

[16] Assume (for reductio) $\mathrm{D}(p \;\&\; \mathrm{I}p)$. By **K** and &-elimination, $\mathrm{D}p$ and $\mathrm{DI}p$. But then $\mathrm{I}p$ by **T**; so $\neg \mathrm{D}p$. Contradiction.

borderline itself? Unusual, perhaps—but there is no paradox here, since borderline Fness is compatible with Fness proper. And we have no reason to believe that being borderline F entails not being F. To insist otherwise is to commit what Wright (2003) calls the "third possibility" fallacy of thinking that vagueness somehow precludes truth or falsity. Indeed, the iterativity of vagueness might be thought to reiterate this very idea that borderline Fness, rather than disqualifying something from being F, is entirely compatible with being F: neither precludes the other. The notion of a borderline case itself is no exception: being borderline F cannot preclude being borderline borderline F (on pain of excluding the very thing it entails, given its iterative character). Hence, the compatibility of borderlineness and borderline borderlineness.

Furthermore, to claim that borderlineness appears somewhere is not to say exactly where. One can maintain the existence of borderline cases without having to identify any particular case as being borderline. Indeed, one can maintain the *definite* existence of borderline cases without having to identify any particular case as being *definitely* borderline. After all, given the iterative nature of borderlineness, it may turn out that *nothing* is definitely borderline. Hence, the converse direction of our earlier dictum—that existing definite borderline cases are borderline cases that definitely do exist—must fail.[17]

(D∃I ⇒ ∃DI) *If there definitely exist borderline Fs, there exist definitely borderline Fs[18]

---

[17] Otherwise, given both directions, D∃xIFx *iff* ∃xDIFx *iff* (by duality) ¬∀x¬DIFx, contrary to iterativity.

[18] Surprisingly, Wright's (1992) **DEF** principle (any consequence of any set of propositions each definitely true must also be definitely true) appears to *validate* (D∃I ⇒ ∃DI), as follows. Assume D∃xIFx. By **T**, ∃xIFx. Thus IFa for some a by ∃-instantiation. Applying **DEF** strengthens this to DIFa. Hence ∃xDIFx by ∃-generalization.□

This spells trouble. Wright's (2001:§8) quandary view already appears to commit him to an iterative conception of vagueness, given his claim that any state of quandary induced by vagueness is itself a source of quandary. If so, he must deny the existence (or possibility) of definite borderline cases. Given (D∃I ⇒ ∃DI), this means denying the definite existence (or possibility) of borderline cases. This consequence is not

In general, importing the determinacy operator inside the scope of the existential is invalid.[19]

(D∃ ⇒ ∃D) *If there definitely exist Fs, there exist definite Fs

Exporting determinacy outside the existential, however, is unproblematic.[20]

(∃D ⇒ D∃) If there exist definite Fs, there definitely exist Fs

Contrapositively, the vague existence of Fs rules out the existence of definite Fs. By parallel reasoning, the vague existence of non-Fs rules out the existence of definite non-Fs. In that case, nothing is definitely F, nor is anything definitely not-F. Thus everything must be borderline F. This licenses

(I∃ ⇒ ∀I) If it is indeterminate whether any Fs exist and indeterminate whether any non-Fs exist, everything is borderline F

---

restricted to only the purported cases of vague vagueness, but applies to any vague predicate whatsoever. Therefore, Wright's own views on vagueness force him to deny that vagueness, in the sense of borderlineness, definitely exists. (Although his later (2010:§VIII) advocates adopting a noncognitivist stance toward the contents of our judgments about borderlineness—which might suggest that he believes it *is* vague, after all, whether any vagueness, so conceived, exists at all.)

Don't Wright's broader views on higher-order vagueness preclude the iterative conception? I don't see how. The use of definite borderline cases in his (1987) paradoxical proof against the possibility of higher-order vagueness is entirely inessential. Elsewhere, Wright (1992:132fn.6) appears to concede this point: the proof does *not*—despite appearances—require the existence of definite borderline cases, but only requires the existence of definitely non-definite Fs, for which definite non-Fs suffice (since D¬Fx entails D¬DFx given **K** and **T**).

[19] Modeling determinacy off necessity makes this clear. D∃xIFx should not entail ∃xDIFx, any more than e.g. "Necessarily, a man on the FBI 100 Most Wanted is to be caught dead or alive" (□∃x(◊Fx&◊¬Fx)) entails "A man on the FBI 100 Most Wanted is necessarily such that *he* is to be caught dead or alive" (∃x□(◊Fx&◊¬Fx)). Such a modal interpretation of the determinacy operator (as quantifying over admissible precisifications instead of worlds) is essentially what allows supervaluationists to distinguish between *de dicto* and *de re* readings of determinacy claims and thereby admit the determinate existence of cutoffs while denying the existence of determinate cutoffs. (Although determinacy-introducing rules like Wright's **DEF** principle threaten to render this incoherent; see n.18.)

[20] Assume ∃xDFx. So DFa for some a (by ∃-instantiation). Since ∃-generalization is determinacy-preserving, we can ∃-generalize within the scope of 'D' (i.e. distribute over D(Fa → ∃xFx) via **K**) and get D∃xFx. Hence, ¬I∃xFx.

which can be equivalently expressed as importing the indeterminacy operator inside a universal:[21]

(I∀ ⇒ ∀I) If it is indeterminate whether everything is F and indeterminate whether everything is not-F, everything is borderline F

To recap: we have seen how to derive claims about the existence of borderline borderline cases from claims about the borderline existence of borderline cases, thus vindicating (I∃I ⇒ ∃II). This establishes the connection we sought between two otherwise disparate notions, the vagueness of 'vague' and higher-order vagueness. Some philosophers have been pessimistic about the prospects of finding any real connection between the two. Our discussion reveals such pessimism to be unfounded.

3.4 Sorensen-Hyde vs. Tye-Deas-Hull-Varzi

Where does this leave Hyde's argument? Recall how its assumption that 'vague' is vague required there to be some '*n*-small' predicate within Sorensen's series that is borderline vague. The truth-conditions for '*n*-small', when properly unpacked, showed exactly what was needed for this to happen.

(15) I(∃*x* I(*x* is *n*-small)) iff I(∃*x* (I(*x* is small) and $x \geq n$))

However, these conditions are given in terms of *vague vagueness* (specifically, that of 'small'). Hyde's intent, by contrast, was to draw a general conclusion about *higher-order vagueness*. What is missing is some justification for thinking that any predicate, so long as it is vaguely vague, is guaranteed to be higher-order vague. Our brief exposition on the interaction of indeterminacy with quantification promises a solution. There, I showed how claims of higher-order vagueness

---

[21] We have I∃xFx *iff* (*by def*) I¬∃xFx *iff* (*by duality*) I∀x¬Fx. Similarly, I∃x¬Fx *iff* I∀xFx.

are derivable from claims of vague vagueness (though not vice versa), thus validating inferences of the form (I∃I ⇒ ∃II).[22] As a result, the following is a derivable consequence of (15).

(17) I(∃*x* I(*x* is *n*-small)) only if ∃*x* I(I(*x* is small) and $x \geq n$))[23]

which, given the determinacy of inequality, simplifies to[24]

(18) I(∃*x* I(*x* is *n*-small)) only if ∃*x* (II(*x* is small) and $x \geq n$)

Hence, any vague vagueness for '*n*-small' will be enough to guarantee that something greater than or equal to *n* is borderline borderline small, in which case 'small' is higher-order vague. The Sorensen construction shows that some '*n*-small' predicate, formed out of the ordinary predicate 'small', will be vaguely vague. Therefore, the ordinary predicate 'small' is in fact higher-order vague. We have already seen how the Sorensen construction can be extended to other vague predicates. So this conclusion is not limited to just 'small'. If the argument is able to demonstrate higher-order vagueness *here*, it will surely demonstrate it elsewhere for *any* ordinary vague predicate that can be "Sorensen'ed".

---

[22] This creates trouble for Tye's own views on vagueness (1990, 1994), since he denies the existence of higher-order vagueness but maintains that ordinary vague predicates are nonetheless vaguely vague.

[23] 'Only if' because of the failure of (∃II ⇒ I∃I). Regardless, (I∃I ⇒ ∃II) is the needed direction of entailment.

[24] (LtoR) Assume I(A&B), where DB∨D¬B. Then (i) ¬D(A&B) and (ii) ¬D¬(A&B). Observe that (i) entails DA→¬DB (otherwise, DA&DB would imply D(A&B) since determinacy collects over conjunction (call this rule D-Collect; for proof see n.6)—contrary to ¬D(A&B)). By contraposition, DB→¬DA. Recall that D($\phi \rightarrow \psi$) entails $\neg D \neg \phi \rightarrow \neg D \neg \psi$ (see n.5). In particular, given that &-elimination is "determinately" valid (i.e. D((A&B)→A)), we have ¬D¬(A&B)→¬D¬A. Thus it follows from (ii) that ¬D¬A. Likewise, we can obtain ¬D¬B, which by ∨-syllogism on our assumption DB∨D¬B implies DB. So ¬DA by modus ponens. Therefore IA (and also DB, or B). (RtoL) Assume IA and DB. Suppose D(A&B) for reductio. By **K** and "determinate" &-elimination, we get DA—contrary to IA. Hence ¬D(A&B). Suppose D¬(A&B) for reductio. By **K** and de Morgan(¬&), we get D(¬A∨¬B). But given DB, we can now derive D¬A by "determinate" ∨-syllogism (i.e. applying **K** to $D(((\phi \vee \psi) \& \neg \psi) \rightarrow \phi)$ to get $D((\phi \vee \psi) \& \neg \psi) \rightarrow D\phi$, which by D-Collect reduces to $(D(\phi \vee \psi) \& D\neg \psi) \rightarrow D\phi$)—contrary to IA. Hence ¬D¬(A&B). Therefore I(A&B).

This bridges the gap in Hyde's argument. It also puts to rest Tye's concerns about possible quantifier conflation. Tye's suggestion was that perhaps all Sorensen's sorites shows is that ordinary vague predicates like 'small' are vaguely vague, not that they are higher-order vague, and therefore Hyde's use of Sorensen-style predicates does nothing to prove the existence of higher-order vagueness. The validity of ($I\exists I \Rightarrow \exists II$) shows this is impossible.[25] The existence of vague vagueness *logically entails* the existence of higher-order vagueness. So Hyde's argument cannot demonstrate the first without the second.

Another issue, discussed by Deas (1989), Hull (2005) and Varzi (2005), concerns the true "source" of the alleged vagueness of 'vague'. Deas and Hull claim that although vagueness must appear somewhere in the Sorites series of '*n*-small' predicates, this can be traced to the vagueness of 'small', rather than to any vagueness of 'vague'.[26] As such, Hyde's argument is unsound because Sorensen's setup fails to demonstrate the existence of any higher-order vagueness.

---

[25] Hyde's (2003) own argument against the coherence of Tye's suggestion is flawed, since it relies on supposing—fallaciously—that being vaguely vague entails not being vague (see fn.1-2). This looks like an example of Wright's "third possibility" fallacy. Tye appears to make no commitment to such an inference anyhow. And if so, for good reason. It simply *cannot* be true in general that being vaguely F entails not being F. Otherwise, being vaguely F, since equivalent to being vaguely not-F (by the symmetry of vagueness), would entail not being not-F, hence being F (assuming anything not not-F is F), whereby nothing would be vaguely F for any F (indeed, nothing would be vague at all!). (Logically: If Ip entails $\neg p$, but is equivalent to $I\neg p$, it entails $\neg\neg p$, or p; therefore $\neg Ip$.)

[26] They argue for this on the basis that '*n*-small' gets mentioned in subject position, rather than used in predicate position, in Sorensen's argument. I fail to see the relevance of this observation. Instead of " '*n*-small' is vague", one could equally have said e.g. "some *x* is a borderline case of '*n*-small' ", "being *n*-small admits of borderline cases", or "some *x* is borderline *n*-small". These are all truth-conditionally equivalent, intertranslatable variants of the same claim, only '*n*-small' occurs differently in each (mentioned as subject, mentioned as predicate, used as subject, or used as predicate, respectively). Not even the observation itself is quite correct either, since the inductive premise in (3)—a universally quantified conditional—strictly speaking, doesn't mention *any* specific predicate '*n*-small' (for any *n*). Rather, these are mentioned *by proxy* through " '*n*-small' " and " '*n*+1-small' ", which stand in for individual predicate names (not the predicates themselves), such that substituting in appropriate values (i.e. predicate names) for " '*n*-small' " and " '*n*+1-small' " results in a statement that mentions, rather than uses, predicates '*k*-small' and '*k*+1-small' for some *k*.

The truth-conditions for claims of vagueness derived earlier suggest that this objection is misguided. For we might interpret each derived condition as implicitly expressing what depends on what, as follows. By (14), any vagueness of ‘*n*-small’ derives from some corresponding vagueness of ‘small’. By (15), any vague vagueness of ‘*n*-small’ derives from some corresponding vague vagueness of ‘small’. But the alleged vagueness of ‘vague’ is given by the vague vagueness of some ‘*n*-small’ in Sorensen’s setup. Thus, any vagueness of ‘vague’ will ultimately derive from the higher-order vagueness of some ordinary vague predicate like ‘small’. In that way, the true source of vagueness does lie in ‘small’—but not, as Deas and Hull think, within its vagueness *per se*, rather only in its *higher-order* vagueness. And this is all Hyde’s argument intended to show: that ordinary vague predicates like ‘small’ are higher-order vague. Any reliance on the disputed idea that ‘vague’ itself is vague was inessential for reaching that result.

Finally, there is the charge of circularity. Varzi (2003) criticizes Hyde’s argument for being question begging, since it presupposes the very thing it seeks to prove: the existence of higher-order vagueness. To that end, he offers a proof that any vagueness for ‘*n*-small’ reduces to that of ‘small’, then attempts to argue on this basis that any vague vagueness exhibited by some ‘*n*-small’ also reduces to the higher-order vagueness of ‘small’.[27] In response, Hyde (2003) concedes the entailment relation between his premises and conclusion, but insists that the argument, while

[27] Varzi’s proof that ‘small’ is higher-order vague relies upon assuming Sorensen’s sorites inductive premise to be true. Surely this is mistaken. All the soriticality of ‘vague’ requires is the premise’s *apparent*, not actual, truth.

deductively valid, is nevertheless able to support noncircular reasoning in favor of a nontrivial conclusion.[28]

This paper offers some means to ward off the threat of circularity. Observe that even if the vagueness of 'vague' reduces to (and hence presupposes) the higher-order vagueness of 'small', the reduction relation is not one of simple equivalence. For claims of vague vagueness only entail, and are not entailed by, claims of higher-order vagueness. This lends some support (however marginal) to the idea that the logical relations tying the vagueness of 'vague' to that of 'small' are not trivial enough to warrant the charge of circularity.

More importantly, there really does, as Hyde claims, seem to be an epistemic gap in his argument between premises and the conclusion. Sorensen's argument (3) is soritical, because we find ourselves unable to pinpoint a precise stopping point where the '*n*-small' predicates cease to be vague. That is a simple *first-order* intuition about the soriticality of 'vague'—specifically, about the lack of a sharp boundary for 'vague' along a sequence of Sorensen-style predicates. This is, importantly, *not* the intuition that there is no sharp boundary for 'being less than or equal to some borderline case of 'small' ' along the integers—which is a *higher-order* intuition about the soriticality of a complex predicate, one that embeds other predicates while involving notions like inequality, quantification and borderlineness. Although (as we verified) the contents of these intuitions are logically equivalent, they remain distinct intuitions. (Many surveying Sorensen's setup for the first time, I gather, will likely have the first, not the second.) Likewise, there is the possible intuition (available perhaps to the logically acute) that no sharp boundary exists for

[28] Nowhere, however, does Varzi bother to work out in detail the derivations of these claims. Nor does he address or acknowledge Tye's concerns regarding the distinction between vague vagueness and higher-order vagueness, as evidenced by his—somewhat hasty—endorsement of the validity of (his own reconstruction of) Hyde's argument. Hyde's (2003:304) reply offers little improvement, since he uncritically concedes "the absence of a logical gap".

‘being less than or equal to some borderline borderline case of ‘small’ ’ along the integers. Although (as we saw) the content of this intuition logically follows from that of the first two, it remains a distinct intuition, with arguably different evidential value. Someone who intuited *that* would stand in a different epistemic position than someone who merely intuited the soriticality of ‘vague’ (i.e. saw and appreciated its lack of sharp boundaries, without first working out the truth-conditions for each ‘*n*-small’). To conclude that ‘small’ is higher-order vague on the basis of the first sort of intuition might be criticized, understandably, as constituting some form of circular reasoning. Reaching the same conclusion on the basis of the simpler intuition, however, arguably involves no circularity of inference. Deriving it simply confirms higher-order vagueness to be a real consequence of what we intuit.

3.5 Higher-order vagueness

I now wish to discuss some issues not mentioned in the Sorensen-Hyde vs. Tye-Deas-Hull-Varzi debate: Even if the possibility of some vaguely vague ‘*n*-small’ guarantees that ‘small’ will have borderline borderline cases, this only demonstrates *second-order* vagueness for ordinary predicates like ‘small’. What guarantees that predicates like ‘small’ will have borderline borderline borderline cases? or higher-order borderline cases at higher orders?

If we had some way of showing there were some *vaguely vaguely* vague ‘n-small’ predicate in the Sorensen series, that would be sufficient to derive the result that ‘small’ has *third-order* borderline cases, by repeated applications of (I∃I ⇒ ∃II). But we cannot assume this. All the soriticality of ‘vague’ along the Sorensen series shows—if this intuition is even correct—is that some ‘*n*-small’ is *vaguely* vague. There is no intuition or appearance of soriticality to support the idea that ‘vague’ has *second-order* borderline cases. Even if there were, nothing in Sorensen’s sorites would show that ‘vague’ has borderline cases at yet higher orders. So there would be no way to confirm that ordinary predicates like ‘small’ are indeed *higher-order vague*, in the sense of admitting *n*th-order borderline cases for any *n*.

All Hyde’s argument has shown so far is that any vague predicate is second-order vague. This falls short of proving that any vague predicate is *n*th-order vague for any order *n*. Yet clearly that was Hyde’s intention: to demonstrate higher-order vagueness at *all* orders, by demonstrating how “ “border case” is vague, and “border border case” is vague, etc., resulting in the inadequacy of anything short of an infinite iteration within the characterisation.” (1994:39) Yet Hyde’s considerations never manage to establish that conclusion. Even if he is right in thinking that the self-characterizing, or “homological”, character of the paradigm analysis of ‘vague’ reveals the notion of ‘border case’ to be vague as well, this only guarantees the existence of *border border* cases, not the existence of *border border border* cases.[29] The latter would be guaranteed if ‘border border case’ were vague, but Hyde’s argument does nothing to show this. Indeed, it is doubtful whether any amount of conceptual analysis can show how the vagueness of ‘borderline F’ is supposed to generalize to higher orders, in order for it to be true that ‘borderline borderline F’, ‘borderline borderline borderline F’, etc. are all vague too.

So what *can* guarantee that any vague predicate is higher-order vague? I shall argue that Sorensen’s proof of the vagueness of ‘vague’ really is powerful enough to demonstrate the existence of higher-order vagueness at any order for ordinary vague predicates, assuming it already successfully demonstrates their second-order vagueness. To see this, however, requires

---

[29] There is still the issue of what it even means to say “ ‘border case’ is vague”, as ‘border case’ is not a predicate. Plus, as Tye (1994:44) points out, this only guarantees border border cases for *some*, not all, vague predicates.

going beyond mere conceptual analysis of the terms 'vague' or 'border' and instead employing semantic reasoning.

In our proof for (I∃I ⇒ ∃II) (see n.14), we saw that a consequence of vague vagueness is the iteration of borderlineness: any predicate F that is vaguely vague will admit only of borderline borderline cases, if it admits of borderline cases at all. Hence the result of *restricted iterativity*:

(I∃I ⇒ I/II) If it is indeterminate whether borderline Fs exist, anything borderline F is borderline borderline F

This iterative effect is not limited to just second-order borderlineness: anything second-order borderline will be third-order borderline, anything third-order borderline will be fourth-order borderline, and so on. Notice however that this iterative behavior only holds selectively for those predicates that are *vaguely* vague. Thus, while it may be true that some '*n*-small'—one of the vaguely vague ones in the Sorensen series—will exhibit higher-order borderline cases, it does not immediately follow that this will also be true of 'small', whose vagueness is, we may assume, not itself vague: there *definitely* are borderline small numbers. Nonetheless, (I∃I ⇒ I/II) guarantees that some '*n*-small', because vaguely vague, will be higher-order vague. That is because, given restricted iterativity, anything borderline *n*-small will be borderline borderline *n*-small, and hence borderline borderline borderline *n*-small, etc. And, as I argued earlier (see n.25), there can exist borderline *n*-small numbers, even if only *vaguely* so, since being F is compatible with being vaguely F.

Thus, we may suppose that some ‘*n*-small’ in Sorensen’s sorites will be vague, but only vaguely so, in which case it will admit of higher-order borderline cases of all higher orders. It will be convenient to first recall some previously established results:[30]

(I∨) It is indeterminate whether A or B and it is determinate whether B or ¬B, *if and only if* it is indeterminate whether A and it is determinate that ¬B

(I&) It is indeterminate whether A and B and it is determinate whether B or ¬B, *if and only if* it is indeterminate whether A and it is determinate that B

We can then reason as follows. Consider first-order borderline cases of ‘*n*-small’:

(19) I(x is n-small)
iff I(x is small or $x < n$) by definition(‘*n*-small’)
iff I(x is small) and $x \geq n$ by (I∨)

Thus, whatever is borderline *n*-small is borderline small. Consider then second-order borderline cases of ‘*n*-small’:

(20) II(x is n-small)
iff II(x is small or $x < n$) by definition(‘*n*-small’)
iff I(I(x is small) and $x \geq n$) by (I∨)
iff II(x is small) and $x \geq n$ by (I&)

Thus, whatever is borderline borderline *n*-small is borderline borderline small. Consider now third-order borderline cases of ‘*n*-small’:

(21) III(x is n-small)
iff III(x is small or $x < n$) by definition(‘*n*-small’)
iff II(I(x is small) and $x \geq n$) by (I∨)

[30] That is, **(I∨):** I(A∨B), DB∨D¬B -||- IA, D¬B (see n.6) and **(I&)**: I(A&B), DB∨D¬B -||- IA, DB (see n.24).

iff I(II(x is small) and $x \geq n$) by (I&)

iff III(x is small) and $x \geq n$ by (I&)

Thus, whatever is borderline borderline borderline *n*-small is borderline borderline borderline small. Continuing in this way, we can show that anything higher-order borderline *n*-small is (to the same degree) higher-order borderline small. But (I∃I ⇒ I/II) guarantees the existence of higher-order borderline cases of '*n*-small' at all higher orders. These will therefore be higher-order borderline small at all higher orders. Hence the higher-order vagueness of 'small'.

3.6 'Vague' at higher orders

And what of 'vague' itself—is it higher-order vague as well? The Sorensen sequence contains predicates that are borderline cases of 'vague'. But are there also borderline borderline cases of 'vague'? or borderline borderline borderline (etc.) cases of 'vague'? Perhaps 'vague' exhibits higher-order vagueness in the same way ordinary vague predicates are higher-order vague.

To be sure, the soriticality of 'vague', as exhibited by the Sorensen sorites, does not show 'vague' to be higher-order vague, any more than the soriticality of any ordinary vague predicate like 'small' shows *it* to be higher-order vague. I claim nonetheless that 'vague' *is* higher-order vague. This can be shown by another Sorensen-style argument, except we must now use a modified version of the Sorensen construction. We define the predicate ''*n*-small'-vague' as follows. For positive integers k, n:

(22) '*k*-small' is '*n*-small'-vague iff '*k*-small' is vague or $k < n$

Here, ‘‘n-small’-vague’ is a second-order predicate that takes first-order predicates ‘*k*-small’ as argument.[31]

Consider the predicate ‘‘1-small’-vague’. No *k* is less than 1, so the comparative condition is (definitely) not satisfied by any ‘k-small’. As for the other condition (namely, ‘*k*-small’ is vague), this will be satisfied by those ‘*k*-small’ appearing earlier in the series for smaller values of *k* (since these share the same borderline cases as, and hence are just as vague as, ‘small’ itself), but will not be satisfied by those ‘*k*-small’ appearing later in the series for greater values of *k* (since none of these are vague). Thus, the predicate ‘‘1-small’-vague’ holds true of ‘*k*-small’ when *k* is small, but not when *k* is large. Yet there is no clearly last ‘*k*-small’ predicate that is ‘1-small’-vague, since (as the original Sorensen argument showed) there is no clearly last vague ‘*k*-small’. So ‘‘1-small’-vague’ is soritical and admits of borderline cases. Hence ‘‘1-small’-vague’ is vague.

Now consider the predicate ‘‘10^6-small’-vague’. We are assuming 10^6 is large enough to be *definitely definitely* not-small, so that ‘10^6-small’ is *definitely* precise.[32] Thus, for any $k \geq 10^6$, ‘*k*-small’ will be *definitely* precise, since it too will have definitely lack borderline cases, but it will also be *definitely* false that $k < n$; so ‘*k*-small’ will *definitely* fail both conditions, and hence *definitely* not be ‘10^6-small’-vague. For any $k < 10^6$, it will be *definitely* true that $k < 10^6$, so ‘*k*-small’ will *definitely* be ‘10^6-small’-vague. Thus, for any value of *k*, ‘*k*-small’ will be either

[31] That it is *second*-order should be unproblematic, since ‘is vague’ is already a second-order predicate as well.

[32] By definition, we have: *x* is 10^6-small iff *x* is small or $x < 10^6$. For any $x < 10^6$: $DD(x < 10^6)$, so DD(*x* is 10^6-small) by **K**, ∨-intro; hence D¬I(*x* is 10^6-small). For any $x \geq 10^6$: DD¬(*x* is small) by assumption, and $DD\neg(x < 10^6)$, so DD¬(*x* is 10^6-small) by **K**, D-Collect, de Morgan(¬∨); hence D¬I(*x* is 10^6-small). Thus, we can reason by cases: $\forall x(x < 10^6 \rightarrow D\neg I(x \text{ is } 10^6\text{-small}))$ and $\forall x(x \geq 10^6 \rightarrow D\neg I(x \text{ is } 10^6\text{-small}))$, therefore $\forall x(D\neg I(x \text{ is } 10^6\text{-small}))$.

definitely '10^6-small'-vague or *definitely not* '10^6-small'-vague, but never *borderline* '10^6-small'-vague. So ''10^6-small'-vague' has no borderline instances, and is therefore precise, i.e. not vague.[33]

We now see the makings of a sorites: ''1-small'-vague' is vague while ''10^6-small'-vague' is not vague, yet there is no clearly last vague ''*n*-small'-vague' predicate. The appearance of there being no boundary at all entices us to accept the following sorites argument.

(23) ''1-small'-vague' is vague
For any *n*, if ''*n*-small'-vague' is vague, then ''*n*+1-small'-vague' is vague
Therefore, ''10^6-small'-vague' is vague

This constitutes a *higher-order* Sorensen sorites. Like Sorensen's original sorites, it demonstrates the vagueness of 'vague', except now the soriticality of 'vague' is relative to a sequence of ''*n*-small'-vague' predicates. The absence of any sharp cutoff along that sequence—i.e. any *n* such that ''*n*-small'-vague' is definitely vague but ''*n*+1-small'-vague' is definitely not vague—means there is some borderline vague ''*n*-small'-vague'.

From this fact we can then reason as before.

(24)	I(∃*k* I('*k*-small' is '*n*-small'-vague)	for some fixed *n*
	iff I(∃*k* I('*k*-small' is vague or $k < n$)	by definition(''*n*-small'-vague')
	iff I(∃*k* (I('*k*-small' is vague) and $k \geq n$)	by (I∨)
	only if ∃*k* I(I('*k*-small' is vague) and $k \geq n$))	by (I∃I ⇒ ∃II)
	iff ∃*k* (II('*k*-small' is vague) and $k \geq n$)	by (I&)

Thus, given the soriticality of 'vague' and its failure to draw no sharp boundaries within the higher-order Sorensen sequence of ''*n*-small'-vague' predicates, it follows that 'vague' has borderline borderline cases.

---

[33] The formal argument mirrors that in the previous n.32.

Reasoning as before confirms that this is not limited to second-order borderlineness. Since we know there must exist some ‘*k*-small’ that is borderline ‘*n*-small’-vague (since the original Sorensen series guarantees the existence borderline vague ‘*k*-small’), then:

(25)	I(‘*k*-small’ is ‘*n*-small’-vague)	for some *k*
	…	
	only if II(‘*k*-small’ is ‘*n*-small’-vague)	by (I∃I ⇒ I/II)
	iff II(‘*k*-small’ is vague or $k < n$)	by definition(‘‘*n*-small’-vague’)
	iff I(I(‘*k*-small’ is vague) and $k \geq n$)	by (I∨)
	iff II(‘*k*-small’ is vague) and $k \geq n$	by (I&)
	…	
	only if III(‘*k*-small’ is ‘*n*-small’-vague)	by (I∃I ⇒ I/II)
	iff III(‘*k*-small’ is vague or $k < n$)	by definition(‘‘*n*-small’-vague’)
	iff II(I(‘*k*-small’ is vague) and $k \geq n$)	by (I∨)
	iff I(II(‘*k*-small’ is vague) and $k \geq n$)	by (I&)
	iff III(‘*k*-small’ is vague) and $k \geq n$)	by (I&)
	…	
	etc.	

Therefore, ‘vague’ has higher-order borderline cases at *all* orders.

Sorensen used his original sorites to argue against the idea from Frege and Russell of effecting a clean divide of the precise from the vague in order to demarcate the former as the proper domain of logic. Our higher-order Sorensen sorites shows that any finer classificatory scheme must also be out of reach: we cannot precisely delineate the definitely precise from the borderline vague, given the existence of second-order borderline vague predicates[34]; nor among

[34] Sorensen (1985:136) makes this point, but claims, without argument, that ‘vague’ has borderline borderline cases.

the definitely definitely precise and the borderline borderline vague, given the existence of third-order borderline vague predicates; and so on, at any further order of precision.[35]

3.7 Vaguely vague vagueness

Our discussion has underscored the importance of separating claims of vague vagueness from claims of higher-order vagueness for any vague predicate. 'Vague' is not exempt from this division, since it too is vague. I have argued that it is not just vague but *higher-order* vague. But is it *vaguely* vague? Now, supposing 'vague' *were* vaguely vague, that would serve to show why it is higher-order vague, since the latter is a direct consequence of the former by (I∃I⇒∃II). But the converse entailment fails. So any higher-order vagueness of 'vague' is no guarantee of its vague vagueness. It remains an open question: Is 'vague' not just vague and higher-order vague but vaguely vague?

If 'vague' is vaguely vague, then it must be like one of the '*n*-small' predicates in the original Sorensen sequence. Recall that for 'vague' to exhibit vagueness (i.e. unqualified, first-order vagueness) in Sorensen's series, some '*n*-small' predicate must be borderline vague. Thus, for 'vague' to exhibit *vague* vagueness in the Sorensen series, it must be indeterminate whether *any* '*n*-small' predicate is borderline vague. Sorensen and Hyde think there is in fact some '*n*-small' that is borderline vague. Yet is the existence of such a borderline vague predicate a *determinate*

---

[35] *Proof.* By definition, 'F' is vague *iff* ∃xIFx *iff* ¬∀x¬IFx *iff* ¬∀x(DFx ∨ D¬Fx) *iff* ¬('F' is precise). Definitizing these definitional truths gives D('F' is vague ↔¬('F' is precise)). Given D(A↔B) |- IA↔IB (see n.5), we can derive I('F' is vague) ↔ I¬('F' is precise), which is just I('F' is vague) ↔ I('F' is precise). Definitizing definitions again yields D(I('F' is vague) ↔ I('F' is precise)), hence II('F' is vague) ↔ II('F' is precise). Repeating *n* times for any *n* shows that 'F' is *n*th-order borderline vague iff 'F' is *n*th-order borderline precise.

truth or a *vague* one? That will determine whether 'vague' is *definitely* vague or *borderline* vague, at least within the context of the Sorensen series.

Hyde thinks 'vague' can only be determinately vague, if vague at all. As he puts it: "vagueness determinately is or determinately is not homological." (2003:302) The Sorensen setup must then guarantee the *determinate* existence of borderline vague '*n*-small' predicates. This is, of course, not to guarantee that some '*n*-small' will be *determinately* borderline vague. For it may be true that all borderline vague '*n*-small' turn out to be *higher-order* borderline (and hence not determinately borderline) vague. This is nonetheless compatible with it being determinately true that such borderline vague '*n*-small' exist. (Recall the general failure of importing determinacy inside existentials, i.e. the invalidity of (D∃I ⇒ ∃DI).) Indeed, our discussion of the higher-order vagueness of 'vague' confirmed this to be the case. Given (I∃I ⇒ I/II), any vaguely vague '*n*-small' is not just borderline vague, but borderline borderline (etc.) vague.

In contrast, Tye thinks that 'vague' might be vaguely vague—which, according to his own views on vagueness, would make 'vague' just like all other "common or garden vague predicates." (1994:45) In that case, the existence of some borderline vague '*n*-small' may still be true, just not *determinately* so.[36] This means that anything borderline vague must be borderline

---

[36] One may even think that, given the *vague* vagueness of any purportedly vague predicate, nothing is vague *at all*—not even 'vague' itself! This type of nihilism about 'vague', although deeply implausible, is nonetheless coherent, assuming that being vaguely F is compatible with not being F (so that it is consistent to maintain both I(∃xIFx) and ¬∃xIFx for all F, in which case no predicate ever has borderline cases—though this fact is only *vaguely* true).

Hyde appears to attribute this sort of view to Tye, when he describes Tye's objection as offering "an alternative explanation of the soritical nature of 'vague' as evidenced by Sorensen's argument—it is vaguely vague and thus not vague." (2003:302) However, this is not the only way to take Tye's suggestion. The vague vagueness of 'vague' could be meant to preempt the need for higher-order vagueness *only*—i.e. *not* the existence of first-order vagueness—in explaining the soriticality of 'vague' in the Sorensen sequence. Thus 'vague' could be vaguely vague without being higher-order vague. Unfortunately, given (I∃I ⇒ ∃II), this situation is impossible anyhow (see n.22).

borderline vague.[37] Otherwise, if some '*n*-small' were determinately borderline vague, it would be determinately true *of it* that some borderline vague '*n*-small' exists, hence determinately true *simpliciter* that some borderline vague '*n*-small' exists—contrary to assumption.[38]

The disagreement here reflects a difference in general attitude had toward ordinary vague predicates. One may think, following Hyde, that the soriticality of any ordinary vague predicate 'F' guarantees the determinate existence of borderline Fs. For some predicates ('small' etc.) just are *definitely* vague, and hence, on the paradigmatic conception, *definitely* admit of borderline cases. Alternatively, one might think that all the soriticality of an ordinary vague 'F' guarantees is that there are some borderline Fs, not that there *determinately* are such cases: their existence remains indeterminate. Thus whether any predicate 'F' we would ordinarily count as vague really *is* vague, is itself a vague matter.

It is important to note that Tye's suggestion that no predicate is definitely vague does not invalidate the point that 'vague' is vague. To be sure, Hyde and others commenting on Sorensen all take for granted that '*n*-small' will be definitely vague for definitely small values of *n*, since these will just share the same borderline cases as 'small', and there definitely are borderline cases for 'small'—that is, assuming 'small' is indeed definitely vague. Whereas Tye denies this: no '*n*-

[37] Other aspects of Tye's general views on vagueness must be called into question in light of our discussion (noted earlier in n.22). Tye denies the existence of higher-order vagueness for ordinary vague predicates, on grounds that the vague vagueness of such "garden" vague predicates is enough to defuse the sorites paradox (see his 1990:551-2). Any higher-order vagueness of 'vague' is denied on the same grounds: it is the *vague* vagueness of 'vague', not any higher-order vagueness, that explains away the Sorensen sorites argument for 'vague' (1994:45). And yet, such distinctions are incoherent if, as I have argued, ($I\exists I \Rightarrow \exists II$) is valid. Pending an explanation for why ($I\exists I \Rightarrow \exists II$) should fail, the only way to restore logical consistency to Tye's view and block the iteration of borderlineness (which is what generates higher-order vagueness), given the vague existence of borderline cases, is to deny that there are any borderline cases at all—in which case, everything purportedly vague is vaguely vague, but nothing is higher-order vague, *nor even first-order vague* (see also n.36). No doubt some will find this solution unsatisfactory.

[38] Recall the proof for ($\exists DI \Rightarrow D\exists I$) from n.14.

small' is definitely vague, not even '1-small', because 'small' is not definitely vague in the first place—no vague predicate is. If so, can we still take the Sorensen series to demonstrate the vagueness of 'vague'? Sorensen's setup would appear to require the existence of both definitely vague '*n*-small' and definitely precise (i.e. definitely not vague) '*n*-small', so that the borderline vague '*n*-small' may appear in between these. Instead, Tye envisages simply a range of borderline vague '*n*-small' followed by definitely precise '*n*-small'. Yet Sorensen's conclusion that 'vague' is vague still goes through. For even on Tye's picture, 'vague' remains soritical: it is still true that no sharp boundary exists between the vague '*n*-small' and the non-vague '*n*-small', even if no definite instances of the former exist. That is enough to guarantee the existence of some borderline vague '*n*-small', which is all the Sorensen sorites was meant to show.[39]

Nor would Tye's denial of definite vagueness affect any other results established so far. The derivability of higher-order vagueness from vague vagueness (I∃I ⇒ ∃II), used to show that 'small' is second-order vague, only required the existence of borderline vague '*n*-small', not their *definite* existence. The iterativity of borderlineness (I∃I ⇒ I/II), used to show that 'small' is higher-order vague, only required the existence of borderline vague '*n*-small', not their *definite* existence. The applications of these results to the higher-order Sorensen sorites, used to show that 'vague' is higher-order vague, only required the existence of borderline vague ''*n*-small'-vague',

---

[39] Whether to persist in calling 'vague' *soritical* is largely terminological. One may take the soriticality of F to consist in the absence of sharp boundaries for 'F', i.e. no point at which the definite Fs become definite non-Fs, where this *presupposes* that there are definite Fs to be begin with. On that understanding of "soritical", although 'vague' fails to be soritical since it lacks definite instances, 'non-vague' nonetheless still counts as soritical, because plenty of predicates are *definitely* not-vague, and whatever reason one had for thinking that 'vague' isn't sharply bounded counts equally in favor of thinking that 'non-vague' isn't sharply bounded. Presumably, this means there will be borderline *non-vague* predicates—but these are just the borderline *vague* predicates. On another understanding, soriticality consists in susceptibility to sorites argument. This too is satisfied on Tye's picture, since the sorites inductive premise (if '*n*-small' is vague, then so is '*n*+1-small') still appears (albeit illusorily) to be true, despite the fact that its antecedent is never *definitely* satisfied.

not their *definite* existence. Any assumption that 'small', '*n*-small', ''*n*-small'-vague', or even 'vague' itself is *definitely* vague was inessential. Thus, Tye's suggestion that 'vague' can only be vaguely vague poses no real alternative to either Sorensen's conclusion that 'vague' is vague or Hyde's conclusion that higher-order vagueness exists, nor is it incompatible with the idea defended here that 'vague' may be higher-order vague. Indeed, it confirms Sorensen's conviction that any perfectionist desire to excise all vagueness from the realm of logic by cordoning off the precise portion of our logical languages: for if 'vague' is vaguely vague, so is 'precise' only vaguely precise, with the result that no precise language can precisely describe its own precision.[40]

A theoretical choice must be made. One may construe the soriticality of 'vague' as evidence for the definite existence of borderline cases for 'vague', in which case 'vague' is *definitely* vague. Or alternatively, one may doubt whether any predicate definitely admits of borderline cases and so construe the soriticality of 'vague' as evidence merely for the weaker claim that it has (and not that it *definitely* has) borderline cases. Both proposals are coherent.[41]

Can Sorensen's sorites settle the matter? If it can, I do not see how. The Sorensen construction at best guarantees the existence of borderline vague predicates. It does not guarantee their *determinate* existence. To get that result requires assuming things that go beyond Sorensen's argument: issues that hinge on the nature of vague predicates and their alleged soriticality. Alas, Soresen's sorites cannot prove everything. This arguably demonstrates yet another way in which

---

[40] *Pf.* First, 'vague' is vague *iff* ∃F I('F' is vague) *iff* ∃F I¬('F' is vague) *iff* ∃F I('F' is precise) *iff* 'precise' is vague. Definitizing these definitions gets D('vague' is vague ↔ 'precise' is vague). Given D(A↔B) |- IA↔IB (see n.5), we then have: I('vague' is vague) *iff* I('precise' is vague). But: I('precise' is vague) *iff* I¬('precise' is vague) *iff* I('precise' is precise). Therefore, 'vague' is vaguely vague *iff* 'precise' is vaguely precise.

[41] As noted earlier, Hyde's reply to Tye rests on fallacious reasoning (see n.25).

the phenomenon of higher-order vagueness comes apart from that of vague vagueness: Sorensen-style sorites are able to show that 'vague' exhibits the first, but are unable to decide whether it exhibits the second. Between the two, higher-order vagueness is the easier to demonstrate; the notion of vague vagueness remains elusive.

Chapter 4: Margins for Error in Meaning

Timothy Williamson has sought to defend epistemicism against the objection that our inability to find sharp boundaries for vague terms is evidence for their nonexistence. He argues that, if they were to exist, we would be unable to know exactly where they lay. His explanation for this is in terms of margin for error principles to the effect that if one knows in a given case, one does not falsely believe in sufficiently similar cases. Our vague terms could have undergone slight, undetectable changes in meaning. The fact that we would, despite this difference, continue using them as we do violates margin for error principles, and therefore undermines any claim to knowledge about the boundaries of our vague terms.

Thus, on Williamson's account, it is meaning instability that explains vagueness-related ignorance according to some appropriate margin for error principle. Is our knowledge of what our vague terms mean also undercut by such meaning instability together with margin for error principles? It is, if our meaning beliefs are able to easily deviate from the facts. Williamson insists they do not. He appeals to a principle of linguistic induction on which induction into the appropriate term-using practice suffices for knowledge of its meaning. In this paper I dispute the cogency of that appeal.

Section 1 provides a reconstruction of Williamson's account that margin for error principles undercut claims to knowledge of a vague term's boundaries. Section 2 formulates the objection that these principles also undercut claims to knowledge of a vague term's meaning. Section 3 examines Williamson's response. Sections 4 and 5 raise objections to this response. I first argue that a principle of linguistic induction fails to preclude one type of possibility where one's meaning beliefs deviate from the facts. I go on to argue that there are independent reasons for rejecting the principle itself.

4.1 Knowledge of meaning

Take the vague predicate 'bald'. On epistemicism, this vague predicate has an unknown sharp cutoff point: there exists a unique number n such that anyone with n hairs on his scalp is bald but anyone with n+1 hairs is not bald, only we are ignorant of what that number is.[1] Suppose the cutoff for being bald were at 3,000 hairs. Williamson's preferred explanation for why we are ignorant of this fact is as follows. I may take 3,000 to be the cutoff point: anyone with 3,000 or less hairs is bald and otherwise is not bald. I may believe for instance that borderline Barry, who I know to have 3,000 hairs on his head, is bald. My belief would as a matter of fact be true. For Barry's hair count would qualify him to be bald—just barely, but barely enough. But my belief would not constitute knowledge. For knowledge requires some margin for error. As Williamson has put it: "if one knows in a given case, one does not falsely believe in sufficiently close cases."[2] If I am to know that Barry is bald, it must be the case that I believe truly in all similar cases.[3] Yet

[1] Whether someone is bald ordinarily depends on how many hairs he has on his scalp. In what follows, I shall assume that whether one is bald depends solely on the number of hairs on one's scalp. To be sure, it depends on much more. Actual linguistic practice demands we recognize that other factors are relevant in determining whether one is bald, for instance, the shape and configuration of one's hair, and that certain factors are relevant in determining the appropriate comparison class, for instance, one's age. Our simplifying assumption is harmless though. Everything said in what follows can be generalized to more complicated accounts on which baldness depends on more than just how many hairs one has on his scalp.

[2] Cf. Williamson 2000: 76.

[3] Belief and knowledge are both sensitive to the way we conceptualize things. The object of our belief that p or of our knowledge that p must be conceptualized as such. Thus, "I know that Barry is bald" should be understood as being elliptical for "I know that Barry is bald under the guise of a sentence of the form 'Barry is bald'." Similarly, "I believe that Barry is bald" should be understood as being elliptical for "I believe that Barry is bald under the guise of a sentence of the form 'Barry is bald'." Crucially, if I am to

this is not so. Had linguistic practice been slightly different, the predicate 'bald' could have easily meant something slightly different, such that its cutoff were one removed from what it actually is, say, at 2,999 hairs instead of 3,000 hairs. Despite this difference I might nonetheless be equally adept at using the term 'bald', as manifested by my willingness to apply it to exactly the same individuals that I am actually willing to apply it to, for instance, to Barry. Assuming the relevant facts are held fixed—that he exists, that he has 3,000 hairs, that I know his hair count, that I am just as willing to count Barry as 'bald', etc.—the scenario would be sufficiently similar to the actual world, only the extension of 'bald' has shifted by a hair. Yet my utterance that "Barry is bald" in such a counterfactual scenario would express a false belief about Barry, since he would have one hair more than the cutoff at 2,999 hairs. This violates the relevant margin for error constraint: regarding whether Barry satisfies 'bald', I wouldn't believe truly in all similar cases; thus I can't be said to actually know that Barry is bald.

And what can be said about borderline Barry can be said about all other cutoff borderline cases. I can't actually know that anyone with exactly 3,000 hairs is bald. So I can't actually know that anyone with 3,000 hairs or less is bald and is otherwise not bald. Nor can anyone else. Hence, the actual cutoff point for 'bald' is unknowable, even if by sheer accident or chance someone somehow hits upon the correct number.

Indeed, the same sort of reasoning extends from cutoff borderline cases such as borderline Barry to all borderline cases in general. In this way, our ignorance of where cutoffs for vague terms lie is just a special case of our ignorance in borderline cases. Williamson's explanation for our ignorance is essentially the same across the whole spectrum of borderlineness: for any given

---

count as believing, and hence knowing, that 3,000 is the cutoff mark for 'bald', it is not enough for me to believe, for instance, that anyone is bald if he has that greatest number of hairs one can have and still be bald.

borderline case of F-ness, I cannot know it to be F (or not F) because ‘F’ is unstable in meaning: it could have easily meant something slightly different such that it failed to apply (or did apply) to the thing in question.

Williamson claims such meaning instability is the source of all vagueness-related ignorance.[4] It is this appeal to meaning instability that I find in danger of being incongruent with his claims about knowing what our vague expressions mean. The worry is this: if knowing requires not falsely believing in all similar cases, as a margin for error principle would have it, then I do not count as knowing what I mean by ‘bald’ if I am mistaken about what ‘bald’ means in some nearby counterfactual case. But that sort of nearby possibility is exactly what is suggested by Williamson’s own account when it is said, for instance, that I would be unable to detect slight shifts in the meaning of ‘bald’. Williamson of course is eager to insist that we know perfectly well what our vague expressions mean. However, it is not apparent, the objection goes, whether by the lights of his own theory he is entitled to that claim.

4.2 Linguistic deviance

Suppose that knowledge of what a vague predicate means can be expressed in this kind of way: ‘bald’ means *bald*. Suppose that *bald* is in fact the concept picked out by ‘bald’ given the way the term is actually used in the linguistic community. Consider another concept in the vicinity, *bald**, whose cutoff is just one-removed from that of *bald*. Williamson agrees that ‘bald’ could have easily meant *bald**, had the community-wide linguistic practice been slightly different. The

---

[4] Cf. Williamson 1994: 230-4.

question then is whether it is easily possible for me to have false beliefs about what 'bald' means, even if I currently believe truly that 'bald' means *bald*.

It seems it is. Instead of *bald*, I could have easily possessed the concept *bald**. The possibility intuition in question concerns only my own individual usage of 'bald' and not the rest of the linguistic community's usage. That is, it seems easily possible for the overall community-wide use patterns of 'bald' to be the same as they actually are such that the concept circulating throughout is *bald*, and yet somehow I diverge from the community-wide practice by exercising my own "maverick" concept *bald**. Call this a case of *counterfactual maverick deviance*. In such a case, everyone else might believe truly that 'bald' meant *bald*, whereas I would believe falsely that 'bald' meant *bald**. But a principle of margin for error would then dictate that I *don't* count as actually knowing what 'bald' means.

Thus the claim to knowledge of meaning seems threatened by nearby counterfactual possibilities of divergence between the facts of what a vague term means and what I take it to mean.

And not just in close counterfactual scenarios. This might be what happens actually. Suppose today I believe that 'bald' means *bald*. Then tomorrow I begin exercising my own maverick concept *bald**. In so doing I diverge in my own use of 'bald' from the rest of the linguistic community which continues to exercise the concept *bald** in its use of 'bald'. Accordingly, my belief about what 'bald' means also diverges from the community: everyone else continues to believe that 'bald' means *bald* whereas I now believe that 'bald' means *bald**. Call this a case of *actual maverick deviance*. In such a case, my later belief that 'bald' means *bald** is false, since it still carries its old meaning, despite the new presence of maverick concepts. As before, a principle of margin for error would then dictate that I don't count as actually knowing what 'bald' means.

The possibility of divergence also works in the reverse direction.[5] The meaning of 'bald' fixed by community-wide usage could have been *bald** rather than *bald*. Still, it seems that in such a case I could have easily continued to believe it to mean *bald*. As we shall see later, Williamson denies this possibility on the grounds that the concept I exercise in my use of 'bald' must align with the rest of the community's: if everyone else means *bald** by 'bald', then so must I, assuming I am properly inducted into the communal 'bald'-using practice. Now, this might accurately describe how things are ordinarily, or perhaps (we should like to think) how things are actually. Given that I am as a matter of fact a normal English speaker, what I mean by 'bald' is just whatever any other English speaker means by 'bald'; and what I take myself to mean by 'bald' is just whatever any other English speaker takes himself or someone else to mean by 'bald'.

But all that is contingent on what linguistic practice I happen to be inducted into. What if I were unwittingly inducted—as it were, abducted—into a different linguistic community in which 'bald' meant something different? Call this a case of *counterfactual linguistic abduction*. We may imagine in such a case that the change in meaning from *bald* to *bald** is undetectable. Since I wouldn't be able to detect the difference, it seems I could easily continue exercising the concept *bald* whenever I use 'bald' and thus continuing believing that 'bald' means *bald*. But I would be wrong to believe so, for 'bald' in the new community carries a different meaning. Again, a principle of margin for error would say that I thereby don't actually know what 'bald' means.

---

[5] Sainsbury (1994:916) already points out that the objection under discussion really entertains two types of possibility: cases of (what I am calling) maverick deviance are to be distinguished from cases of (what I am calling) linguistic abduction.

The intuition remains if we were to suppose that this were something that actually happens. Suppose I have all my life been a part of an English-speaking practice in which ‘bald’ means *bald*. On Monday I believe that ‘bald’ means *bald*. Given my induction into this practice, I can be said to believe truly. On Tuesday however I am abducted, taken out of my current linguistic community, and placed in another one. The swap is done without my knowledge or consent. I am completely unaware of any change. In particular I am unable to detect the difference between the meaning ‘bald’ carries in my new surroundings and the meaning it carried in my old surroundings. In my new environment I continue to be exposed to and actively participate in usage of the term ‘bald’; only the difference between *bald* and *bald** is too minute to detect, so I do not suspect any difference in the relevant ‘bald’-speaking practice. Call this a case of *actual linguistic abduction*. It seems that in such a case I could easily continue believing that ‘bald’ meant *bald* after the abduction, in which case I would believe falsely. By a principle of margin for error, I therefore can’t be said to actually know what ‘bald’ means.

In all these cases, the intuition is that it is easily possible for an ordinary English speaker to have false beliefs about what ‘bald’ means. But this is jointly inconsistent with the claim that the ordinary English speaker knows what ‘bald’ means together with a margin for error principle that one knows what ‘bald’ means only if his belief about what ‘bald’ means is true in all similar cases. Williamson’s contention that knowledge requires a margin for error seems to predict that no one in fact knows the meanings of vague expressions. In fact, matters are made even worse if we assume (not implausibly) that metaphysical modality conforms to S5. If all possibilities are necessarily possible, then a principle of margin for error would dictate, not just that no one does in fact know the meanings of vague expressions, but that no one *could possibly* know the meanings of vague expressions—clearly, a bad result.

Generalized to any vague predicate 'F', the inconsistent triad becomes:

P1 Anyone who knows what 'F' means must have a true belief about what 'F' means in all similar cases.

P2 Ordinary users of 'F' know what 'F' means.

P3 It is easily possible for anyone's belief about what 'F' means to be false because it diverges from what 'F' in fact means.

How might one reasonably attempt to restore consistency to Williamson's account? A seemingly promising solution would be to deny its most controversial premise, P3. This is the solution favored and endorsed by Williamson himself (first proposed in Williamson 1994 and later defended in Sainsbury 1997).[6] Indeed, given his prior commitments to P1 and P2, it is natural to conclude that Williamson must deny P3. If so, he owes an explanation for why cases of counterfactual maverick deviance, cases of actual maverick deviance, cases of counterfactual linguistic abduction, and cases of actual linguistic abduction are all impossibilities, or at the very least why none of these are easily possible. Attempted explanations of why this should be so shall be taken up next.[7]

---

[6] Williamson says the view "is quite consistent with the relevant margin for error principles. If 'heap' had meant something slightly different, speakers would have recognized that slightly different meaning. They would not have misidentified it as the present meaning. Whatever the exact details of their disposition to assent and dissent, they would then have been participants in the practice of using 'heap' as it would then have been. The identification even of a vague meaning can manifest a disposition to be reliably right." (1994: 237)

[7] Other options include: revising P1 by giving an alternative formulation of margin for error principles, or revising P2 by qualifying what "knowledge of meaning" amounts to with respect to vague terms in such a way that is jointly consistent with P1 and P3. I take these up in a longer version of this paper. My own preferred solution, on behalf of Williamson, is the second: to retreat to the weaker claim that we know what our vague terms mean, but only inexactly.

4.3 Semantic externalism and linguistic induction

On the envisaged solution, the imagined scenarios where my beliefs about what my vague terms mean easily diverge from the facts about what they do mean, are all deemed impossible. The explanation for why they are impossible may be thought to stem from a commitment to a sort of semantic externalism, on which in no close world, or scenario similar to how things actually are, does an individual's concept diverge from the community's.[8]

In order to fully work, however, the reply must embody more than just the semantic externalist idea that an individual's concepts are invariably aligned with the community's as a whole. It must also assume that our knowledge of those meanings is so invariably aligned. Thus, not only must what I mean by 'bald' align with what the community as a whole means by 'bald', but what I believe 'bald' to mean must also align with what the community believes it to mean.

What guarantees this? Some explanation has to be given if we are to be convinced that my beliefs about meaning cannot easily go wrong.[9] Williamson's own account appeals to the notion of induction into a linguistic practice:

---

[8] As articulated by Mark Sainsbury:

> "I am party to a conceptual and linguistic practice; there are some possible differences in the practice which simply entail that my concept or language is different, regardless of how similar things may seem to me in the different situations. We cannot say in any detail what such differences are like (if we could, we would understand in detail how meaning supervenes on use, which we do not); but they are easily possible in that their obtaining would involve only minute behavioural shifts in the community, not ones with any significant impact on belief-forming mechanisms. The easy possibility is thus a social shift which drags my concept with it. What could not so easily happen is that my concept would get out of line with the community's. The difficulty is not causal, but issues from the externalist view: what counts as the precise extension of a subject's vague concepts is fixed on a community wide basis, so, in close worlds, idiosyncrasies in his usage will not count as manifestation of an idiosyncratic concept. In close worlds, a subject's concepts co-incide with those of his fellows, whether theirs are the same as or different from their (and his) actual ones. (1997: 916-7)

[9] Here is Sainsbury:

> "To know what a word means is to be completely inducted into a practice that does in fact determine a meaning. (1994:211)

Thus I am guaranteed to know the meaning of an expression, vague or not, simply in virtue of being inducted into the practice of using that expression.

Unfortunately, Williamson never exactly says what it is to be inducted into the same linguistic practice as everyone else. The most he offers toward any positive account is a necessary condition: "rough matching" in my own overall use patterns of 'bald' with others' is required if we are to count as being inducted into the same 'bald'-using practice.[10] To be sure, requiring that I exhibit *exactly the same* overall use patterns for 'bald' as everyone else in the 'bald'-using practice is obviously too strong. Conversely, similarity in individual use by itself is by no means sufficient: two 'bald' users in completely separated linguistic communities or worlds cannot possibly count as sharing in the same 'bald'-using practice even if they use 'bald' in

---

> "Externalism is the not the only way to attain the structure of this answer (though it seems to be the way Williamson would prefer). Perhaps I have a mechanism which reliably aligns my concepts with the ambient ones. Its reliability ensures that neither of the two possibilities which seemed to threaten knowledge of meaning [maverick deviance and linguistic abduction] would be easy; so the threat would peter out. The truth, I presume, requires a combination of both kinds of factor. There will be some externalist, thus constitutive, determinations; but these will be possible only if some causal mechanism reliably keeps me in touch with the concepts, language and topics of discourse of my fellows, for this mechanism will be relevant to which other speakers and objects help constitutively to determine my concepts and meanings." (1997:917)

As I understand him, Sainsbury is here drawing attention to the need for an explanation to make good on the externalist reply. The suggestion offered is that "some causal mechanism" or another is what reliably keeps me in touch with the meanings of my terms. The proposal must be worked out in fuller detail, however, if it is to actually be convincing and to avoid sounding unsatisfying and hopelessly programmatic. More pressing perhaps is the point that reliability is exactly what is at issue here. Theory aside, why should I even believe that my judgments about meaning will be reliable in the first place?

[10] "To be inducted into a practice, it is not necessary to acquire dispositions that exactly match those of other insiders. Of two people who understand the word 'thin', one may be willing to apply it in a slightly wider range of cases than the other. Rough matching is enough. Perhaps no two speakers of English match exactly in their dispositions to use 'thin'. It does not follow that no two speakers of English mean exactly the same by 'thin'. For what individual speakers mean by a word can be parasitic on its meaning in a public language. The dispositions of all practitioners collectively determine a sense that is available to each." (1994:211)

roughly the same way; presumably, some necessary causal connection is missing.[11] Beyond these remarks Williamson does not say much more.

I remain pessimistic about the prospect of defining linguistic induction as anything more than just knowing what the relevant terms mean. For our purposes, however, I am willing to grant that the notion of linguistic induction is one which we have an intuitive grasp on. Still, I will argue, the theory is met with problems. Let us assume there is such a thing as being inducted into a linguistic practice—whatever that may mean.[12] The proposal then becomes:

(Induction) Necessarily, one knows what a vague term means just in case he is inducted into the relevant practice using that term.

To make good on the externalist reply, Williamson needs the additional assumptions that:

(i) One can be inducted into only one relevant practice at a time.

(ii) Each term has only one fixed meaning per given practice that is available to all those inducted into that practice.

Otherwise, if I can be simultaneously inducted into distinct practices yielding different meanings for the same term, or if I am inducted into a single practice that simultaneously yields different meanings for the same term, then I shall have different candidate meanings to choose from for some single term, in which case induction into a practice will hardly guarantee that I know *the* meaning of that term.

---

[11] Perhaps the notion of induction into the relevant practice must ultimately be given in causal terms. If so, it risks being as explanatorily premature as the move considered earlier, namely, of positing the existence of a "causal mechanism" to explain how one invariably knows the meanings of his terms. The burden of explanation collapses back onto the task of identifying what that mechanism is.

[12] Williamson speaks of being "completely inducted" into a linguistic practice. But we can ignore for the moment what it is to be *completely* inducted into a practice, as opposed to being only *partially* inducted into it. In a longer version of this paper, I argue that Williamson is allowed to admit the notion of partial induction into his theory, so long as it corresponds to vague knowledge of meaning, and not inexact knowledge of meaning.

Does it follow that I couldn't easily have false beliefs about what 'bald' means? The externalist reply proceeds as follows. Assume (Induction) together with (i) and (ii). Applied to 'bald', these imply that knowing the meaning of 'bald' is necessarily coextensive with being inducted into the relevant 'bald'-using practice (where "the meaning of 'bald'" and "the relevant 'bald'-using practice" are understood to be unrigidified descriptions).

This rules out maverick deviance. It is impossible that if 'bald' had still meant *bald*, I could easily believe that 'bald' meant *bald**. For maverick concepts like *bald** are not easy to acquire. In nearby worlds, if the rest of the community means *bald* by 'bald', then that is what I mean too by 'bald', assuming I am inducted into the relevant practice. Indeed, I wouldn't even possess the concept *bald** because I wouldn't be inducted into any practice in which 'bald' carried that meaning. So I cannot believe that 'bald' means *bald** in nearby worlds where it still means *bald*. I cannot take its meaning to be different where its meaning is no different.

Actual linguistic practice becomes a mere special case: given that I am actually inducted into the 'bald'-using practice, I know that 'bald' means *bald*, and I can't believe it means otherwise because I don't actually possess any similar yet distinct concepts such as *bald**. Thus both types of possibilities are ruled out: counterfactual and actual maverick deviance.

This also appears to rule out linguistic abduction. It is impossible that if 'bald' had meant *bald**, I could easily believe that 'bald' still meant *bald*. In nearby worlds where 'bald' means something different, I am no longer inducted into a practice in which it means *bald*; thus I no longer possess that concept. So I cannot still believe that 'bald' means *bald* in nearby worlds where it means *bald**. I cannot take its meaning to be the same where its meaning is different.

The externalist reply exploits the following feature of (Induction) plus (i) and (ii). It is not just deviant concepts that are difficult to come by: deviant beliefs about what concepts one is exercising are also difficult to come by. It is not just that, in nearby worlds where everyone else in the linguistic community means *bald** by 'bald', that is what I mean too by 'bald' in those worlds. It is that in nearby worlds where everyone else takes 'bald' to mean *bald**, that is what I take 'bald' to mean too in those worlds (if I take it to mean anything). This is not to say that the envisaged deviance from community-wide meaning is impossible, only that it never occurs in close counterfactual circumstances.

Of course, the concepts *bald* and *bald** are different meanings of 'bald', one actual and one counterfactual, precisely insofar as they have different cutoff points. That so much is given by Williamson's own views on the meaning instability inherently found in vague predicates. Williamson would deny however that taking 'bald' to have a slightly different extension from that which it in fact has is thereby taking it to have a slightly different meaning; a supposed difference of extension does not entail a supposed difference of meaning.[13] Whenever I attempt or purport to envisage a nearby scenario where my beliefs about the meaning of a vague term diverge from the facts, I do not succeed: I merely envisage a scenario where I have different beliefs about the extension of 'bald', not a scenario where I have different beliefs about its meaning; similarly, supposed sameness of extension does not entail supposed sameness of meaning.

---

[13] It seems there is at least room for initial worry. Williamson claims that a difference in use does not imply a difference in meaning when the meaning of the term is "stabilized by natural divisions, so that a small difference in use would make no difference in meaning." This is the case with natural kind terms like 'gold'. But, he adds, slight shifts in our use of vague terms do give way to slight shifts in meaning, because unlike natural kind terms, "the meaning of a vague term is not stabilized by natural divisions in this way. A slight shift along one axis of measurement in all our dispositions to use 'thin' would slightly shift the meaning and extension of 'thin'." Of course, the envisaged shift in use is at the level of the entire linguistic community, and not the individual speaker. But a worry can be raised as to why this pattern of meaning shifts being induced by use shifts should not be replicated in the individual case as well. Cf. Williamson 1994: 231.

Thus, in envisaging a case of maverick deviance, I am not really envisaging a case where I believe 'bald' to mean *bald** when in fact it means *bald*. I merely envisage a case where I believe the extension to be different when it is not. Likewise, in envisaging a case of linguistic abduction, it is contested, I am not really envisaging a case where I believe 'bald' to mean *bald* when in fact it means *bald**. I merely envisage a case where I believe the extension to be the same when it is not. Both types of possibility show just how unreliable I am when it comes to judging where the cutoff for a vague predicate like 'bald' is. But such unreliability in my judgment of extension should not be mistaken as unreliability in my judgment of meaning.

One might protest in defense of P3: "Aren't meanings for vague predicates at least in part individuated by extension, in which case, a difference in extension beliefs *does* entail a difference in meaning beliefs?"

To see why this does not follow, consider how two competent users of the same vague term might disagree over whether it applies to a given borderline case, without disagreeing over its meaning. Consider once again borderline Barry. Suppose that Tim assents to "Barry is bald" while Jim dissents from or is unwilling to assent to "Barry is bald". Clearly, they disagree. But over what? First of all, they disagree over whether Barry is bald. To settle the disagreement would be to settle the matter of fact of whether Barry is bald: if it is the case that Barry is bald, then only Tim is correct; if not, then only Jim is correct. Second of all, they disagree over the extension of 'bald'. Tim believes that Barry falls within the extension of 'bald'; Jim does not. It does not follow from this however that Tim and Jim disagree over what 'bald' means. Otherwise, if their disagreement were (even in part) over the meaning of 'bald', then it would be possible to settle their disagreement by settling what 'bald' in fact meant. That would imply that Tim and Jim meant different things by 'bald'. Suppose then for the sake of argument that according to

Tim 'bald' means *bald* and that according to Jim 'bald' means *bald**. Then Tim's utterance of "Barry is bald" expresses his belief *that Barry is bald* and Jim's utterance of "Barry is bald" expresses his belief *that Barry is bald**. These two beliefs clearly vary in truth conditions: Tim's belief is true if and only if Barry has 3,000 hairs or less, whereas Jim's belief is true if and only if Barry has 2,999 hairs or less; so, supposing Barry had 3,000 hairs, Tim's belief would be true but Jim's belief would be false. Now, it is arguably a trademark feature of disagreement in belief that, if the disagreement is mutually recognized, each party must take his own beliefs (relevant to the disagreement) to be mutually exclusive with the other party's beliefs (relevant to the disagreement). Tim and Jim (we may assume) do mutually recognize their disagreement. So Jim, in particular, must take his own belief, as expressed by his own utterance of "Barry is bald", to be mutually exclusive with Tim's belief, as expressed by Tim's utterance of "Barry is bald". But this is something he isn't forced to do in the current situation. He needn't believe that if his own belief is true then Tim's belief is false. For Barry might have 2,999 hairs. In that case, it would be the case both that Barry is bald* and that Barry is bald, making it both true that Barry is bald* and true that Barry is bald, making true both Jim's belief *that Barry is bald** and Tim's belief *that Barry is bald*; so Jim's belief being true would not exclude Tim's belief being true—contrary to the assumption of mutually recognized disagreement. Thus, we should not suppose that in their disagreement, by their respective uses of 'bald', Tim means one thing, *bald*, while Jim means another thing, *bald**.

In general, we should not suppose it is possible that where disagreement occurs over the application of a vague predicate to borderline cases, one party means one thing while the other means another. This is not surprising. Intuitively, the disagreement between Tim and Jim is not over the meaning of 'bald'. (In fact, they might insist that they did mean the same thing by 'bald', and that this is what makes it possible for them to agree in the first place!) It is not even a disagreement over the meaning of 'bald' *plus* whether or not Barry is bald. The disagreement is just over whether or not Barry is bald. Period. Granted, the disagreement can be rephrased in metalinguistic terms, as being over whether the predicate 'bald' applies to the individual Barry. But this is not to frame the issue as being one over the meaning of 'bald'. It is simply to recognize it as being an issue which involves disagreement over the extension of 'bald'. Disagreements over borderline cases and cutoff point are disagreements over use and not meaning: the point of disagreement is over how a vague predicate is to be applied and not what it means.

4.4 Linguistic abduction

Where does this leave us? We began by looking at a possible way of saving Williamson's theory from inconsistency with the two claims that we know what our vague terms mean (P2) and that to count as knowledge our meaning beliefs must be reliable (P1). The solution was to reject P3 by showing how maverick deviance and linguistic abduction are never possible in close worlds, given a certain semantic externalism on which what I mean when I use a vague term, as well as what I take myself to mean, is fixed by the linguistic community I am in. A worry was raised that the individuation of meanings by extension gives independent reason to think that maverick deviance and linguistic abduction were possible. The worry was seen to rest on the dubious premise that a difference in belief about extension entails a difference in belief about meaning, which gave bad predictions about what goes on in cases of disagreement over the application of vague predicates to borderline cases. Now, if this were all there was to our original objection, it would be a bad objection indeed. But our acceptance of P3 in no way rests on that premise, which I agree ought to be rejected. We have independent reasons to believe that cases of maverick deviance and linguistic abduction are possible; hence we have independent reason to believe that P3 is true. In this section, I shall try to show why these reasons are not undercut by Williamson's appeal to semantic externalism.

We concluded in the previous section that if knowledge of meaning requires induction into the relevant practice, cases of counterfactual maverick deviance, counterfactual linguistic abduction, and actual maverick deviance all seem to be rendered impossible. Does the possibility of actual linguistic abduction suffer the same fate? I suggest that it does not.

Suppose that in the actual world the 'bald'-using practice I am currently inducted into, call it p_1, is such that 'bald' means *bald*. I actually know that 'bald' means *bald*. Now suppose that tomorrow I become the victim of a case of linguistic abduction: I am taken out of my current linguistic environment and placed in another, one in which the 'bald'-using practice p_2 is such that 'bald' means *bald**. The abduction is carried out swiftly and surreptitiously, without my consent or knowledge. Suppose enough time then passes for me to interact with the members of my new linguistic community; enough, let us say, so that I have become inducted into p_2 by a certain time t. But 'bald' means *bald** in p_2. So I know that 'bald' means *bald** at t.

Question: can I also believe that 'bald' means *bald* at time t? Nothing seems to prevent us from saying so. By (i), p_2 is the only 'bald'-using practice I am inducted into at t; so I am no longer actually a part of p_1 by t. And by (ii), since 'bald' means *bald** in p_2 at t, 'bald' does not mean *bald* in p_2 at t. So I am not at t inducted into any linguistic practice in which 'bald' means *bald*. It does not follow from any of this however that I lack the concept *bald* at t. For all that has been said, it may still be within my possession at t. For our only constraints were put on how concepts may be acquired; nothing was said about how concepts are lost or what is required to retain them. I possess the concept *bald* when I am still a part of and inducted into p_1. Nothing said so far implies that by t I must no longer be in possession of that concept. If so, then nothing keeps me from believing at t that ‘bald’ means *bald*. But then, I am able to have a false belief about the meaning of ‘bald’ at t, for ‘bald’ does not mean *bald* within p_2 at t. Hence, the attempted externalist reply fails to rule out the possibility of actual linguistic abduction. So P3 of our inconsistent triad remains.

This supports the pretheoretic intuition we had that actual linguistic abduction is indeed a real possibility. If I am abducted by those from another linguistic community—linguistic aliens, so to speak—I could be wrong about what I took ‘bald’ to mean in my new surroundings, even after talking with plenty of ‘bald’-users in my new surroundings.

The intuition is strengthened if we draw out the case a bit. Suppose that on Monday I am still within my old linguistic community in which ‘bald’ means *bald*. That day I judge Mary to be bald. My utterance “Mary is bald” expresses my belief *that Mary is bald*. Then on Tuesday I am abducted into a new linguistic community, one in which ‘bald’ means *bald**. Enough time passes for me to enjoy plenty of interactions with other ‘bald’-users in my new surroundings. Suppose that by the end of the week at time t, I have been inducted into the new ‘bald’-using practice. It seems that I can recall truly at t: “I believed on Monday that Mary is bald.”[14] On the externalist reply, any utterance of ‘bald’ made by me at t will express *bald**, so my utterance “I believed on Monday that Mary is bald” at t will express *that I believed on Monday that Mary is bald**. Yet this is false: what I believed on Monday was not *that Mary is bald** but *that Mary is bald*. So I cannot make true reports about past beliefs I held when I was part of a different linguistic

[14] Granted, such “memory” objections to externalism are not uncontroversial. Cf. Tye & Heal 1998.

practice. That is a bizarre and unattractive consequence of the assumption (ii) made by the externalist reply.

4.5 Ignorant induction

There is also independent reason to doubt (Induction) itself. Suppose that actual linguistic practice has it that 'bald' means *bald* and that the cutoff for 'bald' is at n hairs, so that supervenience of bald facts on hair facts entails that *being bald* and *having n hairs or less* are necessarily coextensive. A consequence of (Induction) is that my induction into the relevant 'bald'-using practice guarantees that I know the meaning of 'bald'. The following considerations are meant to challenge this sufficiency claim.

Scenario 1: zero hairs. It is possible that I might have believed that 'bald' means having exactly zero hairs on one's head. I might do so because I have had it explained to me on numerous occasions that "To be bald is to have no hair," and I have always understood this quite literally. Whenever others have described individuals who are not entirely hairless as being 'bald', I have always taken them to be speaking figuratively or in exaggerating jest. Only when entirely hairless individuals are labeled as 'bald' do I take the speaker to be speaking literally. Likewise, I take myself to speak literally only when applying 'bald' in the presence of hairless scalps; when applying it to non-completely hairless individuals, I expect others to understand me as speaking figuratively. My misunderstanding is understandable, but it would nevertheless be the case that I failed to understand what 'bald' meant. In this scenario it seems I would not count as knowing the meaning of ‘bald’.[15] But intuitively, it also seems that I would count as being inducted into the relevant ‘bald’-using practice. After all, I have interacted with plenty ‘bald’-users in the practice. Moreover, the misconception I have that being bald is having exactly no hairs could be something that lasts for years as I continue to interact with those in the practice, until by good fortune I come across someone who disillusions me. But if I do count as being inducted into the relevant practice, then by (Induction), I also count as knowing the meaning of ‘bald’, which I do not.

Scenario 2: caps. It is possible that I might have believed that ‘bald’ meant wearing a cap. I might do so because the only individuals I have had pointed out to me as being ‘bald’ are those who wish to hide (as we would describe it) their baldness. It is not an unlikely mistake to be made, say, by a child when he sees aged hatted men walking around or someone storming out of a barbershop with his baseball cap, and is told that these individuals are ‘bald’. But if I were to make it, I would hardly count as knowing the meaning of ‘bald’. Definite cases of baldness abound who do not wear any caps. But intuitively, it seems I could still count as being inducted into the relevant ‘bald’-using practice. Perhaps it is a rather impoverished form of being a part of the practice. But (Induction) nonetheless would imply that I thereby know the meaning of ‘bald’, which I do not.

Scenario 3: lucky guess. It is possible that I might have believed that the cutoff for the vague term ‘bald’ was at n hairs. It would be a lucky accident that in so believing I hit upon the truth, since it is in fact true (we are supposing) that anyone is bald just in case he has n hairs or less. On Williamson’s account, I do not count as knowing this because the cutoff could have easily been

[15] One may reply that, although it is uncontroversial that I am mistaken in my use of ‘bald’, it is nonetheless unclear whether I thereby fail to know what ‘bald’ means. After all, it may be said, Burge’s Alfred still understands the term ‘arthritis’; he merely misapplies it.

one or two hairs removed. Indeed, on anyone's account, I should not count as knowing where the cutoff of 'bald' is, if there is a cutoff. That so much is clear. But do I count as knowing what the meaning of 'bald' is? That much is not so clear. On the one hand, if I believe that vague predicates have sharp cutoffs (supposing I am an epistemicist), it seems perfectly felicitous for me to try and guess where those cutoff points are. On the other hand, the moment I begin insisting, for instance, that the cutoff for 'bald' is at some specific number of hairs and not anywhere else—as I might if I am a fervent believer that 'bald' means having n or fewer hairs—it seems I have failed to understand something about the vague nature of 'bald', namely, that its cutoff point (if it exists) resists any attempt to be located. In general, it seems essential to any vague predicate that, not only is it impossible to know where its cutoff is, but one also cannot possibly reasonably believe to have located its cutoff point; cutoffs are things that frustrate not just knowledge but even reasonable belief.[16] If knowledge of the meaning of vague predicates requires appreciating this point, then in particular I will not count as knowing the meaning of 'bald'. But intuitively, it seems I could still count as being (at least partially) inducted in the relevant 'bald'-using practice, in which case (Induction) would imply that I know the meaning of 'bald', which I do not.

Scenario 4: disregarded vagueness. It is possible that I might have believed that the cutoff for 'bald' was at n hairs but have had significantly different beliefs about how 'bald' is used. Suppose I believe that 'bald' was introduced into the linguistic community as having a stipulated cutoff point at n hairs and that it continues to be used in this way; and that the only reason why

---

[16] Williamson provides an account for this intuition, rather unsurprisingly, in terms of margin for error: variants of the margin for error constraints that held in the case knowledge are said to hold also in the case of reasonable belief. The key adjustment is to require (for the latter case) that one's belief be true in *most* similar scenarios, instead of in *all* similar scenarios (as in the former case). Cf. Williamson 1994 §8.7.

people hedge their judgments about baldness is that they lack complete perceptual information about the relevant hair facts. If people had the time, patience and means to count the number of hairs on someone's scalp at a given time, I believe, we would be able to partition all individuals as either 'bald' or 'not bald'; it would simply be a matter of identifying every individual at a given time who has n hairs or less on his head. Again, it seems I should not count as knowing the meaning of 'bald', because I have misunderstood the vague nature of the term. In fact, it is arguable that in such a scenario I have ceased to believe that 'bald' is vague altogether; my conception of how it is used has by this point radically departed from the normal conception of how vague expressions are actually used. But intuitively, it seems I could still count as being (at least partially) inducted in the relevant 'bald'-using practice, in which case (Induction) would imply that I know the meaning of 'bald', which I do not.

The four scenarios above purport to show that, contrary to Williamson's account, induction into the relevant practice does not suffice for knowledge of meaning. In each scenario, it seems I can still be inducted into the relevant 'bald'-using practice, given ample interaction with other 'bald'-users in community. Perhaps it is a rather impoverished form of being a part of the practice. But (Induction) nonetheless would imply that I thereby know the meaning of 'bald'. Yet in each scenario I do not count as knowing what 'bald' means because of my deviant beliefs about how 'bald' is used. I may have false beliefs about its extension, locating the cutoff either at a far different number of hairs than it actually is (zero hairs scenario) or on the entirely wrong continuum (caps scenario). Or I may accidentally be right about the extension, but misunderstand the vagueness of 'bald', because my belief that the cutoff is at n hairs is either based on no good reason at all (lucky guess scenario) or based on the entirely wrong sort of reason (disregarded vagueness scenario).

Thus, there is in each case the intuition that linguistic induction is possible without knowledge of meaning. The intuition merely relies on our having a pretheoretic grip on the notion of induction into a practice, which I am willing to grant we do have. Taken seriously, it is further evidence that the sort of externalist reply envisaged by Williamson is unworkable.

4.6 Margins for error

The previous section was meant to show that P3 cannot easily be denied. One who thinks P1 and P2 are in comparable standing may draw the conclusion that the outright denial of any of these three statements would be unwise and ultimately untenable. Another option open to him would be to be revisionist about one of the claims. We start with P1.

The idea behind margin for error principles is that knowledge precludes untrue belief in sufficiently close cases. Consider the claim that I know what 'bald' means. What are the relevant beliefs in question whose falsity is precluded by my knowing this? And what makes a case sufficiently close? Since the latter is obviously the more difficult question, let us first address the former. I have stated the margin for error principle governing knowledge of our vague expressions as P1. When applied to my knowledge of meaning with respect to the vague term 'bald', this yields M1.

M1 If I know what 'bald' means, then in all similar cases, I have a true belief about what 'bald' means.

The most obvious shortcoming with this formulation is that what I know and what I believe are not given by 'that'-clauses. To avoid this, we might reformulate the relevant margin for error principle as M2.

M2 If I know that 'bald' means *bald*, then in all similar cases, I have a true belief of the form: 'bald' means *F*.[17]

Here, the object of my knowledge is specified: I know the meaning of 'bald' to be *bald*. However, the object of my belief must remain unspecified, for the belief content may vary from case to case. Williamson allows for the possibility that in one similar case I believe truly that 'bald' means *bald* and in another similar case I believe truly that 'bald means *bald**. That my beliefs about what 'bald' means would shift from one similar case to another does not automatically violate the relevant margin for error constraint; so long as each belief is true in its own respective world, I may still be said to know that 'bald' in fact means *bald*.

I am assuming that beliefs are individuated by their contents, although nothing essential rests on this. With some minor adjustments, we could have chosen to speak of the same belief as having a different content in different circumstances, but with the same result: truth evaluation would still be relative to the world in question; in this case, it would depend on what content the belief had in that world.

Another way to achieve the same effect would be to formulate things metalinguistically, as in M3.

M3 If I know that 'bald' means *bald*, then in all similar cases, my uttering a sentence of the form "'Bald' means *bald*" would express a true belief.

Two features are noteworthy. First, the metalinguistic formulation allows for variation in belief across similar cases, because the belief that would be expressed by my uttering a sentence of the

---

[17] It deserves pointing out that the relevant belief of mine is to be understood as a true *de re* belief about the concept F to the effect that it is the meaning of 'bald'. Thus, merely asserting the disquotational principle "'Bald' means bald" doesn't suffice to express a belief that would satisfy the margin for error condition. For instance, I might lack the concept *bald*, in which case the assertion might express a de dicto belief of the form *'bald' means F*, but not a de re belief about the concept *bald* to the effect that 'bald' means it.

form "'Bald' means *bald*" might vary from case to case.[18] Second, the subjunctive form of the consequent is merely meant to account for the possibility that, although it may not be true that in every similar case I utter the relevant words, it would nevertheless be true that in any similar case were I to utter the relevant words I would thereby be expressing a true belief. These two features assure that what is required for knowledge, as specified in the consequent, is that I believe truly about what 'bald' means in all similar cases—exactly as predicted by M2. For our purposes then, M3 is essentially no different from M2: it is just a torturously and unnecessarily complicated rephrasing. For ease of exposition, let us stick with M2.

Our problem can now be recast in terms of M2 (or equally M3). It seems jointly inconsistent with two claims entailed respectively by P2 and P3: first, that I do know (say) that 'bald' means *bald*; second, that there exist similar cases in which I would lack any true belief of the form: 'bald' means *F*. The first claim follows from my status as an ordinary English speaker who is competent in using the vague predicate 'bald'. The second claim follows if we accept that either of the following are easily possible: cases of maverick deviance, in which I believe that 'bald' means *bald**, or cases of linguistic abduction, in which I believe that 'bald' means *bald*; in either case, my belief of the form *'bald' means F* would be false.

Is there another way of capturing the thought that knowledge requires a margin for error that does not run into this difficulty?

Williamson has proposed other formulations for margin for error principles. In one place, he considers: "If x is similar enough to y, and x is known to be F, then y is F."[19] Call this the *similar-objects formula*. Notice here that the mention of world- or case-similarity in M2 has been

---

[18] Assuming we understand "the belief …" as an unrigidified definite description.

[19] Cf. Williamson 2002: 53.

dropped in favor of the notion of object-similarity. A natural way of applying this principle to the issue at hand is to allow ‘x’ and ‘y’ to range over concepts such as *bald* and *bald** and then interpret ‘F’ as “being the meaning of ‘bald’”. This yields M4.

M4 If I know that *bald* is the meaning of ‘bald’, and *bald* is similar enough to *bald**, then *bald** is the meaning of ‘bald’.

The problem with M4 is that it is obviously false. On any measure of similarity, it should come out true that *bald* and *bald** are “similar enough”; for they are after all indiscriminable concepts. It follows on M4 that in order for me to know that *bald* is the meaning of ‘bald’, *bald** must also be the meaning of ‘bald’. But surely, it is impossible that *bald* and *bald** are both at once the meaning of ‘bald’. To endorse M4 would be to rule out knowledge of meaning in an obviously incorrect way.

Suppose we instead understood the similar-objects formula differently. We allow ‘x’ and ‘y’ to range over predicates and interpret ‘F’ as “meaning *bald*”. This yields M5.

M5 If I know that ‘bald’ means *bald*, and ‘bald’ is similar enough to some term T, then T means *bald*.

The problem with M5 is that it too looks false. It seems there could be a predicate whose overall patterns of use closely resembles (but not exactly) those of ‘bald’ as it is actually used, whose meaning hence closely resembles (but not exactly) that which ‘bald’ actually carries, and whose extension hence closely resembles (but not exactly) that which ‘bald’ actually has. On any measure of similarity, it should come out true that two terms are “similar enough” if they closely resemble each other with respect to these three things: use, meaning and extension. A predicate that was similar enough in all these respects to ‘bald’ might come close to meaning the same as ‘bald’, but would not in fact mean the same; in particular it would not mean *bald*.

Perhaps we ought to turn to Williamson's preferred, original formulation of margin for error principles: "A is true in all similar cases where 'A' is known."[20] Call this the *counterfactive formula*. This version of a margin for error principle easily delivers the result that knowledge is factive, for among the cases similar to that where A is known is the (actual) case where A is known. The same thought is merely extended to all other similar cases. Applied to our case at hand, the formula yields M6.

M6 If I know that 'bald' means *bald*, then 'bald' means *bald* in all similar cases.

The problem with M6 is that it is falsified by Williamson's own theory. It is essential to his account that a vague term such as 'bald' could have easily meant something other than what it in fact means. Such meaning instability is exactly what underlies his employment of margin for error principles to explain our vagueness-related ignorance in the first place. Whatever his standards of similarity, it should therefore come out true that there are "similar enough" counterfactual scenarios in which 'bald' means something different, in which case 'bald' does not mean *bald* in all similar cases.

Another problem with the counterfactive formulation is that it isn't enough to rule out unstable judgments from counting as knowledge. Suppose I judge borderline Barry to be bald. Suppose moreover that this judgment is true: despite his borderline status Barry has few enough hairs to be bald. Still, I could have easily judged differently. My skill at assessing baldness and my competence in applying the predicate 'bald' surely preclude no such move. Thus, my judgment that borderline Barry is bald is unstable. In general, my use of 'bald' begins to exhibit judgment instability when it encounters borderline cases. Even if the facts about baldness are

---

[20] Cf. Williamson 1994 §8.3.

stable, my judgments about baldness may not be. Yet according to the counterfactive formula, my judgment that borderline Barry is bald comes out as satisfying the relevant margin for error principle. For all the formula requires for me to know that Barry is bald is that Barry is bald in all similar cases—which he is, assuming he has the same number of hairs in all similar cases. In fact it is necessarily true that anyone with that same number of hairs is bald. Hence it is true in all similar cases that anyone with that same number of hairs is bald.[21]

What makes the counterfactive formula inadequate is its failure to place any appropriate constraints on the subject's beliefs in similar cases. Unlike M1 and M2, for instance, M6 does not say what beliefs of mine about the meaning of 'bald' must be true in the relevant cases. In fact, the consequent makes no reference to beliefs at all.

Williamson recognizes that such instability in judgment exists and is endemic in our use of vague predicates. He denies however that this explains our ignorance in borderline cases. It is instability in meaning and not instability in judgment that explains why we can't know whether a vague predicate applies to a borderline case.[22] It may be said in Williamson's defense that

---

[21] Williamson is well aware of this worry. Necessary truths such as that expressed by 'Everyone with physical measurements m is thin' (his example), he says, "seem trivially to satisfy the necessary condition for being known imposed by a margin for error principle. A necessary truth is true in all cases; *a fortiori*, it is true in all cases similar to the case in which it is a candidate for being known. How then can a margin for error principle explain our ignorance of a necessary truth?" He goes on to offer a solution: "Someone who asserts 'Everyone with physical measurements m is thin' is asserting a necessary truth, but he is still lucky to be speaking the truth. He does not know the truth of what he says. Although he could not have asserted the proposition he actually asserted without speaking truly, he could very easily have asserted a different and necessarily false proposition with the same words." Peculiarly, however, Williamson does not go on to explicitly revise his formulation of the relevant margin for error principle. It is left as it is, in its inadequate form, as given by the counterfactive formula. Cf. Williamson 1994: 230.

[22] The rationale can be reconstructed as follows. Perhaps in cases of phenomenal sorites, instability in judgment looks to be a promising way of explaining vagueness-related ignorance. That is because we know where to locate the source of the judgment instability: the inaccuracy of our perceptual faculties is to blame. But where our faculties of perception, memory, etc. are not at issue, as in a non-phenomenal sorites, the appeal to instability in judgment looks less convincing. Our inability to decide on a borderline case of tallness is not due to any defect in our discriminatory capacities, especially if we know all the

because judgment instability, although itself a feature of vagueness-related ignorance, is not the cause of vagueness-related ignorance, a margin for error principle which explains vagueness-related ignorance thereby need not also explain judgment instability. This would not however address the worry that M6 is already falsified by the meaning instability of 'bald'.

I am forced to conclude that although alternative versions of the relevant margin for error principles are available, those alternatives are either defective in their own right or just as much, if not even more, ill-suited for reconciling P2 and P3.

4.7 Supervenience and skepticism

In the rest of the paper I present what I believe to be the best solution for saving Williamson's theory from incoherence, which is to revise P1. On the proposal, Williamson should not commit himself to the claim that we know the meanings of our vague terms exactly; rather he ought to say that our knowledge of vague meanings is a form of inexact knowledge.

First, I wish to point to a puzzling feature about Williamson's account. Williamson accepts the pair of supervenience claims: (a) meaning supervenes on use, in the sense that holding fixed the overall community-wide patterns of use for a given predicate thereby fixes its meaning; and (b) extension supervenes on meaning, in the sense that holding fixed the meaning of a predicate thereby fixes its satisfaction conditions, and where certain facts about the environment are held

---

relevant physical measurements, i.e. the man's exact height. The connection between vagueness and instability in judgment in such cases is not clear. In cases where reliance upon perception is entirely absent, such as Wang's Paradox, it seems even less so. It is more plausible that meaning instability, and not judgment instability, is what explains our ignorance regarding borderline cases and cutoff points in non-phenomenal sorites cases. If good, this point should not be restricted to non-phenomenal sorites, but should extend to phenomenal sorites, on pains of failing to deliver a unified treatment of vagueness in both phenomenal and non-phenomenal sorites cases.

fixed, thereby fixes the extension of term. The supervenience claims hold for vague predicates too, like ‘bald’: the overall use patterns for ‘bald’ as used in a certain practice determine its meaning within that practice, and the meaning of ‘bald’ within a practice in turn determine its cutoff point, and thus together with the relevant hair facts, determines its exact extension. What is our epistemic status with respect to the facts about use which determine meaning and to the facts about extension which are in part determined by meaning? Williamson claims that we lack any exact knowledge of such facts. We have no way of knowing exactly what the overall use patterns for a vague term are in our linguistic community.[23] Nor can ever know exactly what the extension of a vague term is, as shown by the unknowable status of borderline cases. Yet, surely Williamson would concede that we have knowledge of some sort of such facts, however rough or inexact. Competent users of ‘bald’ know a good deal about how the term is used and which individuals it definitely applies to or definitely doesn’t apply to. Perhaps we can be said to have with respect to a vague term, inexact knowledge of its use and of its extension.

When it comes to knowledge of meaning, however, Williamson insists that our knowledge of this is exact: I know *exactly* what ‘bald’ means. But how can this be? How can I know only inexactly how ‘bald’ is used overall and what ‘bald’ applies to and yet somehow know exactly what ‘bald’ means, when the use facts fix the meaning facts and the meaning facts (together with the hair facts) fix the extension facts? Given that the A-facts supervene on the B-facts and the B-facts supervene on the C-facts, if I know the A-facts and C-facts only inexactly, should I not also know the B-facts only inexactly? It is a strange feature of the view that it gives us an inexplicable privileged epistemic status with respect to the middle link in this chain of supervenience claims. This doesn’t automatically make the view incoherent, but it certainly calls for explanation.

[23] Cf. Williamson 1994 §8.4.

In certain places though, Williamson does appear to suggest that sometimes our knowledge of what a vague term means falls short of full understanding. For instance:

> "On the epistemic view, our understanding of vague terms is not partial … When I have heard a word used only once or twice, my understanding is partial because there is more to the community's use of it than I yet know. I have not got fully inside the practice; I am to some extent still an outsider. Indeed, I probably think of myself as an outsider, knowing that there is more to the practice than I yet know; my use of the term will be correspondingly tentative and deferential … To know what a word means is to be completely inducted into a practice that does in fact determine a meaning." (1994:211)

and later:

> "… one can think of actual meanings as located on a continuum of possible meanings: but it does not follow that to recognize a meaning one must locate it on that continuum. It is enough to know which of the actual meanings it is. To do that, it is enough to use the term within the appropriate practice." (1994:236-7)

All the talk about 'partial understanding' and 'complete induction' suggests the following idea. The typical English speaker knows fully well what the meaning of 'bald' is. Yet on the rare occasion someone may understand it only partially. Knowledge of meaning is had by induction into the relevant practice. Whether that knowledge makes for complete or partial understanding depends on whether the induction is complete or partial. Partial induction is characterized by a number of things such as: insufficient exposure to others' usage of 'bald' ("…heard a word used only once or twice…"), insufficient experience in one's own usage of 'bald' ("…it is enough to use the term…"), and insufficient confidence in one's own usage of 'bald' ("…tentative and deferential…"). In contrast, complete induction is presumably had when one's exposure, experience and confidence meet a certain threshold.

Induction into the relevant practice now looks to be vague notion. For it exhibits many of the standard features belonging to vague expressions. It admits of borderline cases. It is susceptible to Sorites constructions by varying the above parameters. And so on. On the epistemicist theory, there must then be a tripartite division: one is inducted either fully or partially or not at all. A corresponding division holds for understanding: I understand 'bald' either fully or partially or not at all.

What about knowledge? It is tempting to think that a parallel treatment awaits knowledge of meaning: I know what 'bald' means either exactly or inexactly or not at all. However, the thought is mistaken. Where understanding is either full or nonexistent, the subject either definitely knows the meaning of the term or definitely does not know it, respectively. But where understanding is partial, it is vague whether the subject knows the meaning of the term. Cases of partial understanding are just borderline cases of knowledge. On the epistemicist theory, there is always a fact of the matter in such cases of whether one knows the meaning in question or not—it is simply unknowable which it is. For instance, I can perhaps understand 'bald' partially while knowing exactly what it means, or I may understand it partially without knowing what it means at all; just that no one could know if I knew it or not. Hence, inexact knowledge is not the same as partial understanding. (As we shall soon see, it is rather knowledge that something of a certain sort falls within a certain range.) To say that our knowledge of our vague terms mean is made inexact because their meanings are unstable forces one to say that all our knowledge of meaning for vague terms is inexact, because meaning instability pervades all of our vague discourse. That is to deny that we have any exact knowledge of what our vague words mean. Whereas Williamson insists that exact knowledge is the norm. That is because for him, induction into the relevant practice guarantees exact knowledge of meaning, and linguistic induction is the norm. Moreover, the denial of exact knowledge weakens the need for a tripartite division among knowledge being exact or inexact or entirely absent, for it renders it altogether irrelevant to the question of vague meaning. These points, I should add, are not made explicitly in Williamson

1994. But they are nonetheless, as I understand him, what he probably would say to fend off any such misunderstanding of the view.

What is Williamson's motivation for insisting that we know exactly the meanings of our vague terms? He seems to think that it is to be preferred over its alternative, which is to say that we know the meanings of our vague terms only inexactly. He says:

> "One might react to the phenomenon of indiscriminable semantic differences by concluding that speakers only roughly know what their utterances mean; they cannot uniquely identify their meanings. If this reaction is open to anyone, it is open to the epistemic theorist. However, a less sceptical line of thought deserves to be explored." (1994:236)

I fail to see why the fallback position of inexact knowledge should deserve to be called a "sceptical line of thought" in any but the most superficial sense. Since skepticism questions what we purport to know, presumably Williamson means to say that the fallback position contradicts what we purport to know. As he puts it: "There is a sense in which we often know exactly what an utterance means." (1994:236) But do we purport to know this? I do not see how. One might purport to know exactly what he means by 'bald', but I doubt anyone would seriously purport to know *that* he knew exactly what he means by 'bald'. Having knowledge of meaning does not entail having knowledge of one's own knowledge of meaning. Thus, even if, as a matter of fact, we had exact knowledge of what our vague terms mean, it would not follow that to say we had only inexact knowledge of what they mean is to challenge our claim to know this fact. In that case, the accusation that the alternative is somehow skeptical is ungrounded. It is surprising that this point escapes Williamson's attention, given his own rejection of the KK principle.

Hence, I do not share the worry that saying we have only inexact knowledge of what our vague terms mean is a skeptical position to hold. In fact, I believe it is the right thing to say. More to the point, I believe it is the right thing for Williamson to say on his theory. I shall say why later (§11). For now, let us concentrate on the alternative: the claim that we have exact knowledge of what our vague terms mean. I agree with Williamson that this claim "deserves to be explored" (as much as I disagree with his reasons for why he thinks it is worth exploring). Let us then explore it. In particular, let us examine two objections to the claim and Williamson's responses to those objections. I shall argue that the responses do not mitigate the objections.

4.8 Indiscriminable meanings

Does the indiscriminability of *bald** from *bald* not undercut my claim to knowing exactly that 'bald' means *bald*? Williamson's reply to this objection is as follows.[24] Knowledge of what a vague term means is constituted by an ability to recognize its meaning in different linguistic contexts. It is to be likened to knowledge of who someone is based on what she looks like, which is constituted by an ability to, say, recognize that person upon seeing her face. I may possess such recognitional knowledge of who someone is even if, should a lookalike be present, I would be unable to discriminate the two. Similarly, I should still count as knowing the actual meaning of 'bald', even if I would be unable to discriminate it from counterfactual meanings. Exact knowledge of what something is is not undercut by the mere possibility of indiscriminably different things. If this is true for persons, it is true for meanings.

Mark Sainsbury has suggested (in conversation) that the point can be further developed given a distinction between comparative knowledge and non-comparative knowledge. The distinction is most easily seen with respect to phenomenal looks, so it will be best to begin there.

Say that things *look comparatively the same* when they are compared simultaneously and look to have the same appearance, as when two identical twins Shawn and Shane are standing

[24] Cf. Williamson 1994 §8.5.

before me side by side and look the same. Things *look comparatively different* to me when they are compared simultaneously and look to have different appearances, as when two fraternal twins Shania and Bob stand before me side by side and look entirely different. I can be said to have *comparative knowledge* that two things look the same or different if my knowledge is based on such comparative looks. Thus, I know in the comparative sense that Shawn and Shane look the same, and that Shania and Bob look different.

Say that things *look non-comparatively the same* when I experience them individually, as when I have an experience of a rose at one second and another experience (distinct from the first) of a different rose the next second, and the two roses look the same way. They would look *non-comparatively different* if, upon being experienced individually, they looked to have different colors. Or it may be two experiences of one and the same rose, where either it looks the same in both instances (as when seen from the same vantage point) or it does not (as when seen from different vantage points). Or it may be one single rose which looks the same or different to distinct subjects. Whatever the case, knowledge gained in this way of things looking the same or different is said to be *non-comparative knowledge*.

The distinction can arguably be extended from looks to meanings. I may be said to know comparatively that two concepts F and F* are distinct when I come to know this upon direct comparison of F and F*, say, by entertaining them simultaneously before my mind. Non-comparative knowledge of concept difference is that gained in ways other than direct comparison, say, by inference or memory. For instance, when a listener badly misinterprets something you said and goes on to voice his misunderstanding, you might insist, "No, that is not what I meant," thus voicing your non-comparative knowledge that what he took you to be saying was very different from what you actually said.

The point then would be this. Discriminating counterfactual from actual meanings of predicates, such as *bald** and *bald*, involves comparative knowledge. Knowing the actual meaning of 'bald' does not. I know that 'bald' means *bald*. So I have non-comparative knowledge of the actual meaning of 'bald'. Had 'bald' meant *bald**, I would have known that 'bald' meant *bald** instead. So if bald had meant something different, I would have non-comparative knowledge of that counterfactual meaning of 'bald'. It doesn't follow however that I could know that counterfactual meaning *bald** to be the same as or different from the actual meaning *bald*. I might lack comparative knowledge of whether these meanings are the same. Williamson's analogy of facial recognition not being undermined by the mere possibility of lookalikes may be seen as illustrating this general point that non-comparative knowledge needn't imply comparative knowledge of difference.

But why should I have non-comparative knowledge of the meaning of 'bald' in all similar scenarios? The same problem reemerges of how we can know the meanings of our vague terms when knowledge requires reliability and our judgments about meaning are not always reliable enough to constitute knowledge. Only now, it reemerges under a new guise: it is now the problem of how we can *non-comparatively* know the meanings of our vague terms when *non-comparative* knowledge requires reliability and our judgments about meaning are not always reliable enough to constitute *non-comparative* knowledge. But a problem under a new guise is still a problem. And as we saw earlier, the externalist reply Williamson desires does not offer a suitable solution to the problem.

4.9 Knowledge of non-meaning

A related problem is how I can know that a vague expression means one thing, to the exclusion of other things, if this requires knowing also that it doesn't mean those other things. Knowing the latter would seem impossible, given that vague expressions are unstable in meaning: there are many candidate meanings I cannot know to be ruled out, for there are many things my expression could easily have meant.

Against this objection Williamson responds basically as follows.[25] Suppose the vague term T has actual meaning F and counterfactual meaning F*: as things actually are, T means F, although it could have easily meant F*. I know that T means F. But I don't know that T doesn't mean F*. So I know what my vague words mean, just not what they don't mean. How can that be? Meaning one thing excludes meaning something else. So: if T means F, then T doesn't mean F*. But if I know the antecedent of this conditional, should I not also know its consequent? Only if I knew the entire conditional, since knowledge is closed not under entailment, but under known entailment. But I can't be said to know the conditional: if T means F, then T doesn't mean F*. The reason why is because I can't discriminate what T actually means from what it counterfactually means, that is, I can't discriminate F from F*.

The thought is, there are certain counterfactual meanings of any vague term which, because of their close similarity to its actual meaning, we cannot know to be counterfactual. Perhaps our being unable to know that our vague terms don't mean certain things is partly why Williamson insists on our knowing exactly what they do mean. But why does he think that we must be unable to know the non-meanings of our vague terms? Williamson's underlying concern seems to be that the claim to knowledge of non-meaning violates margin for error principles.

To be sure, there are violations on at least one way of formulating margin for error principles. Recall that on the counterfactive formulation, these principles say: if A is known, then A is true in all similar cases. Thus, if I know that T doesn't mean F*, then T doesn't mean F* in all similar

[25] Cf. Williamson 1994: 235.

cases. But T is unstable in meaning. Even though it in fact means F, it could have easily meant F*. So T does mean F* in at least some similar cases. So it is not the case that T doesn't mean F* in all similar cases. So I don't know that T doesn't mean F*.

However, we saw that the counterfactive formulation faced various problems, including its failure to rule out unstable judgments from counting as knowledge, as well as its incompatibility with our vague expressions being unstable in meaning.

There are better ways to formulate principles of margin for error. For instance, the general formula corresponding to M1 and M2: if one knows, then one believes truly in all similar cases. Thus, if I know that T doesn't mean F*, then I believe truly in all similar cases.

But believe what truly? The relevant belief might vary from world to world. The problem case, though, as we just saw, is the counterfactual scenario where T does mean F*. Therefore, let us consider some similar case w* in which T means F*. Let us suppose for the sake of argument that Williamson is right in saying I would know the counterfactual meanings of an expression should they obtain. So in w*, I know that T means F*.

What then is the relevant belief in w*? It cannot be a belief that T does not mean F*. For in w*, I know that T means F*; so I believe that T means F* (since knowledge entails belief); so I do not believe that T does not mean F* (on pain of holding contradictory beliefs). If the relevant belief in w* is a belief that T means F*, then that belief is true. For in w*, I know that T means F*; so I believe truly that T means F* (since knowledge entails true belief). If the relevant belief in w* is something else, for instance, a belief that T does not mean F, then that belief is true, for T does not mean F; and likewise for all other candidate meanings other than F*.

Therefore, whatever the relevant belief is in w*, it is true in w*. But w* was our problem case. Hence, nothing has been shown to violate the relevant margin for error principle for the claim that I know that T doesn't mean F*. Generalizing to other close counterfactual meanings, nothing has shown that a margin for error principle is incompatible with my claim to knowing that T doesn't mean those things.

Given that knowledge of non-meaning does appear to leave an appropriate margin for error, the threat of inconsistency disappears. For that reason, I fail to see why Williamson denies it on his theory. I must admit that Williamson's motivations here entirely elude my understanding.

4.10 Inexact knowledge

Regardless, I wish to draw attention to a peculiar fact. Far from providing an alternative to the view that our knowledge of the meanings of our vague terms is inexact, the account seems to necessitate this very conclusion. To see this we must first consider what is inexact knowledge.

To use Williamson's own example: I may know on the basis of sight roughly how many people there are in a stadium.[26] My knowledge would be inexact because there is no n for which I know that there are exactly n people. Yet I would be able to place lower and upper bounds on my estimate. Whatever the exact number of people in the stadium is, I know that it is not 200, for I know that there are more than 200 people; I also know that it is not 2,000, for I know that there are fewer than 2,000 people. The range within which n falls lies somewhere between 200 and 2,000. Let 'S(m)' stand for "there are exactly m people in the stadium" and 'K' stand for "I know that". The set of natural numbers such that I do not know that there are not exactly that many people, {x: ¬K¬S(x)}, is finite and nonempty. But every finite and nonempty set of natural numbers has both a least member and a greatest member. So the set has a least member L and a

[26] Cf. Williamson 1994 §8.2.

greatest member U. Hence I know that the exact number of people in the stadium lies somewhere in a range [L, U] where 200 < L and U < 2,000 (of course, I should not be said to know exactly what those numbers L and U are).

We have here a prototypical kind of inexact knowledge: knowledge that something of a certain sort falls within a certain range.[27] My inexact knowledge of how many people are in the stadium is knowledge that the exact number of people in the stadium falls within a certain range. That range consists in all and only those numbers n for which I don't know that there aren't exactly n people. It is reasonable to conclude: I know the number of people there are in the stadium only inexactly just in case there exists a range such that for no number n in that range do I know that there aren't exactly n people, even though the exact number of people does in fact fall within that range. More generally: I know what m is only inexactly just in case there exists L and U such that for no n within [L, U] do I know that $m \neq n$, even though m does in fact fall within [L, U].

Knowledge of meaning, if inexact, would share much of the same structure. I am able to place bounds on the range of possible meanings for 'bald'. Whatever the exact meaning of 'bald' is, I know that it is not *having 30 hairs or less*, for I know that the cutoff for 'bald' is more than 30 hairs; I also know that it is not *having 30,000 hairs or less*, for I know that the cutoff for 'bald' is less than 30,000 hairs. The range within which the cutoff falls lies somewhere between 30 and 30,000. Let 'C(m)' stand for "anyone is bald just in case he has exactly m hairs or less", or

[27] An older example perhaps is Aristotle's question of how many good friends one should have: "But as regards good friends, should we have as many as possible, or is there a limit to the number of one's friends, as there is to the size of a city? You cannot make a city of ten men, and if there are a hundred thousand it is a city no longer. But the proper number is presumably not a single number, but anything that falls between certain fixed points. So for friends too there is a fixed number perhaps the largest number with whom one can live together…" (Nichomachean Ethics IX.10)

equivalently, “the exact cutoff for ‘bald’ is at m hairs”. The set of natural numbers such that I don’t know that anyone is bald just in case he has exactly that many hairs or less, $\{x: \neg K \neg C(x)\}$, is finite and nonempty. By a least number principle, it has a least member L. By a greatest number principle, it has a greatest member U. Hence I know that the exact meaning of ‘bald’ lies on a naturally ordered range of possible meanings all of the form *having n hairs or less* such that its exact cutoff falls somewhere in a range [L, U] for unknown $L > 30$ and $U < 30,000$.

It seems that Williamson must accept this as an accurate picture of how we use ‘bald’. For he claims that there are many candidate meanings within the vicinity. Indeed, there is a whole range of meanings surrounding the actual meaning of ‘bald’ on the continuum of hair numbers, each of which I cannot know to not be the meaning of ‘bald’. On the picture being considered, this is to accept that there is a finite nonempty set $\{x: \neg K \neg C(x)\}$ bounded by L and U. So for all n between L and U, I do not know that n is the exact cutoff for ‘bald’. Hence there is no n between L and U for which I know that n is not the exact cutoff for ‘bald’. Hence there is no n between L and U for which I know that the exact meaning of ‘bald’ is not *having n hairs or less*. Hence there exists some L and U such that for no n within [L, U] do I know that the meaning of ‘bald’ is not *having n hairs or less*. Yet the meaning of ‘bald’ does in fact fall within [L, U]. By our definition of inexact knowledge, I must count as knowing what the meaning of ‘bald’ is only inexactly—contrary to Williamson’s claim that we know it exactly.

Chapter 5: Ignorance and Open Texture[1]

5.1 Preliminary: vagueness, promises, obligations

Epistemicism is the view that there is an unknowable fact of the matter for every vague matter.[2] Stewart Shapiro (2006) offers the following case against epistemicism.[3] A father promises his children they will go to the ballgame if it is sunny, or the movies if it is not. Assuming his promises are obligation-generating, he seems to have the following conditional obligations:

(1) If it is sunny, we ought to go to the ballgame

(2) If it is not sunny, we ought to go to the movies

However, it turns out to be borderline sunny that day: not quite sunny, not quite not sunny either. What should he do? His obligations depend on the weather, which is vague. But, according to the epistemicist, there is a fact of the matter about the weather:

(3) Either it is sunny or it is not

despite the fact that

(4) It is vague whether it is sunny

---

[1] An earlier version was previously published as "Epistemicism, Paradox, and Conditional Obligation", *Philosophical Studies* (in press).

[2] Defenders of epistemicism include Williamson (1994), Sorensen (2001), Hawthorne & McGonigal (2008). Deniers include just about everyone else.

[3] Shapiro's main target is supervaluationism, though he targets epistemicism too. In this paper, I shall be concerned only with the potential threat to epistemicism. See Shapiro 2006:82-86 (ch.3 §6); see also §2.

It is simply unknowable which it is, sunny or not, since the matter is vague. It follows, given the conditionality of his promises, that there is a fact of the matter about what is required:

(5) Either we ought to go to the ballgame or we ought to go to the movies

It is simply unknowable which he should do, ballgame or movies, since the only basis or source of evidence for believing one way or the other—the weather—is itself unknown.[4] This is already objectionable, insofar as one finds the idea of unknowable obligations objectionable (at least in such ordinary circumstances as these). Here is Shapiro:

> "According to epistemicism [...] vague expressions have sharp, but unknowable boundaries. So if the weather on Sunday is borderline nice (i.e. near the sharp boundary), then either it is [sunny] or it is not [sunny], but neither the father nor the children have any way of knowing which. Thus, the family will not know what they have to do in order to fulfil the father's promise. For example, if they go to a ball game and the weather is, in fact, not [sunny], then the father has broken his second promise (given that they do not also go to the movie). And if they go to a movie and the weather is, in fact, [sunny], then he has broken his first promise. I conclude that according to epistemicism, no one should make a promise that has a potentially vague antecedent and a sharp consequent, since if the antecedent falls near the border, she will not know what to do (since she cannot know if the antecedent is true)." (Shapiro 2006: 85-86)

The problem is one of permission. Intuitively, it is entirely up to the father's own discretion to choose either the ballgame or the movies; he is not at fault for choosing either, perhaps so long as he chooses one. In other words, neither option is individually mandated, since each is permissible:

(6) It is not the case that we ought to go to the ballgame and it is not the case that we ought to go to the movies

Yet (5) and (6) are inconsistent. Hence, epistemicism denies our permissibility intuitions.

---

[4] Otherwise, if it is knowable that (e.g.) the ballgame option (or movies) is mandated, this can only be because it is knowable that it is sunny (not sunny), as given by the terms of his first (second) promise; so the state of the weather turns out to be knowable after all—contrary to assumption.

Or so one might claim. In this paper I dispute that claim. I contend that the epistemicist need not endorse the troubling conclusion of unknowable obligations. Better yet, fuller reflection on the information-sensitive nature of conditional obligation shows that epistemicism—far from rejecting the permissibilist conclusion of no obligated choice, as Shapiro thinks it must—actually predicts it. In this way, the epistemicist theory is able to accommodate the very permissibility phenomena its critics say it shuns.

I first argue that Shapiro's own presentation of the example is unconvincing, if it is meant to cause trouble for epistemicism. Stronger arguments are available. I reconstruct two such arguments. However, I show why there is independent reason to reject these lines of reasoning, given the failure of analogous paradoxical reasoning in the Miners Paradox. In particular, we have independent reason to think the underlying inference forms fail for deontic conditionals—namely, constructive dilemma and contraposition. I argue that the epistemicist has special reason to think these inferences fail, provided a certain understanding of deontic *ought* on which the deontic facts exhibit sensitivity to one's epistemic situation. This suggests that permissibilist objections to epistemicism of the sort inspired by Shapiro's considerations overlook the information-sensitive nature of conditional obligation. In doing so, they presuppose an erroneous logic for deontic conditionals. Finally, I suggest where the true difference lies between epistemicist and semantic treatments of vagueness regarding permissibility intuitions: instead of yielding rival predictions, they should be thought of as offering rival explanations for the same phenomena.

Alternatively, an epistemicist could give an error theory about our permissibilist intuitions or defend the possibility of vague (and hence unknowable) obligations. My discussion ignores these alternative responses. Against the first, I will simply assume that our permissibility intuitions like

(6) are robust enough to warrant accommodation on any plausible theory of vagueness—including epistemicism.[5] Against the second, I treat the existence of vague obligations as an orthogonal issue, to be addressed elsewhere.[6] All I intend to show here is that epistemicists need not affirm or deny their possibility in order to circumvent Shapiro-style objections.

5.2 Presentation: Shapiro

First, a quick exegetical detour. In Shapiro's own presentation of the case, the father keeps his promises only by acting so as to make both future-tensed conditionals turn out true:

(7) If it is sunny, we will go to the ballgame

(8) If it is not sunny, we will go to the movies

In the case of vague weather, the father won't know what to do, because he won't know which antecedent is true (and hence which conditional to act upon). He is not in a position to rule out either sunny or non-sunny weather though because, assuming epistemicism, exactly one of these must obtain; therefore he runs the risk of breaking one of his promises. By going to the movies, he risks breaking his promise to verify (7), since it might after all be sunny—in which case they should, as promised, go to the ballgame instead. By going to the ballgame, he risks breaking his promise to verify (8), since it might after all *not* be sunny—in which case they should, as promised, go to the movies instead. By doing nothing, he is guaranteed to break at least one

[5] This assumption, and hence the reply on offer, is not available to epistemicists whose—strictly optional—commitment to a knowledge norm of assertion forces them to say that the unknowability of vague matters forbids, rather than permits, judgment (see Williamson 1994, 2000). The reply is meant for the more permissive epistemicist who does not take vagueness-induced ignorance as having this sort of silencing effect on conversationalists.

[6] Deniers of vague obligations include epistemicists like Sorensen (2001). Defenders include Sider (1995).

promise (corresponding to whichever conditional turns out to have the true antecedent, though he won't know which). And (we may assume) he cannot do both. Either action incurs a risk.

This alone poses no threat to the epistemicist theory. For epistemicism correctly prescribes doing one activity rather than none, since doing nothing guarantees that both (7) and (8) will be false, whereas doing *something* at least has a nonzero chance of making both conditionals true. It merely predicts that people can place themselves in ethically risky situations by making promises with vague conditions—a surprising consequence, maybe, but not fatal for the theory.

The real worry, it appears, lies elsewhere. Should making a set of conditional promises indeed generate certain unknowable, hidden obligations, then this means the father is not free to pursue either option after all. What is problematic is this apparent denial of free choice. The issue does not concern the riskiness of violating hidden requirements incurred by promising to make certain conditionals true; the issue is whether there exist any such requirements at all. Epistemicists who wish to accommodate our free choice permissibility intuitions face a puzzle about conditional promises, not promised conditionals. Shapiro's presentation, in attending exclusively to the truth-valuation of the conditionals themselves while avoiding all use of deontic language, appears unable to capture this fact.[7] Reformulating the objection with deontic conditionals, as in (1)-(6), rather than bare indicatives, I submit, does better.

---

[7] The father's freedom cannot, for instance, consist in the fact that his promises remain unbroken whatever he chooses to do, where this is understood as requiring the truth of both conditionals (7) and (8). For in choosing one option, he must forgo the other; so at least one of the two conditionals (corresponding to whichever he fails to do), under some possible circumstance of action (i.e. not picking *that* option), will have a false consequent. This is enough to make the entire conditional untrue, on any standard analysis of the indicative, assuming its antecedent is not guaranteed to be false either—which it can't be, since it either is indeterminate in truth-value (on either a Weak or Strong Kleene logic), has intermediate truth-value (on a many-valued logic), has no truth-value, i.e. no fixed truth-value across all admissible precisifications (on a supervaluationist logic), or is potentially false in truth-value (on a bivalent logic; see above remarks). Thus there is no guarantee that both (7) and (8) will be true *no matter what he does*.

Presumably, this denial of permissibility is what, crucially, sets apart Shapiro's argument from the following variant argument (which Shapiro himself doesn't advance). Assuming we swear to speak truthfully when describing the weather:

(1*) If it is sunny, we ought to call the weather 'sunny'

(2*) If it is not sunny, we ought to call the weather as 'not sunny'

(3*) Either it is sunny or it is not sunny

(4*) Either we ought to call the weather 'sunny' or we ought to call it 'not sunny'

Here, the epistemicist's commitment to (3*) appears to force an objectionable either-or choice in usage for 'sunny' in cases of vague weather. Now, most semantic treatments of vagueness will surely reject (4*), since according to these accounts it is part of the meaning of any vague predicate 'F' that polar judgments—of the form 'F' or 'not-F'—are either permitted but purely optional (and hence not mandated) or simply inappropriate (hence, again, not mandated) for borderline cases. But to object to the epistemicist theory on these grounds is obviously question-begging, if such semantic assumptions are already thought to be incompatible with epistemicism.[8] In contrast, the permissibility intuitions backing (6), as against (5), in Shapiro's argument are theory-neutral: denying the father any freedom of choice in activity—unlike the sanction or licence of certain semantic choices regarding word use—is implausible *independently* of any controversial claims about the semantic nature of vague expressions.[9]

---

[8] This should be unsurprising, since the argument in (1*)–(4*) is essentially an elaboration of the common "incredulous stare" objection against epistemicism—which, however convincing, is itself question-begging. Although see §8 below. Thanks to an anonymous referee for suggesting this alternative argument.

[9] Another difference is: the fact that hedging ("It's sort of *F*") is an appropriate response when presented with a borderline case makes it easier to read "ought…" in (4*) as denoting some deflationary notion (e.g. *speaking in a truth-apt way* or *what an idealized speaker would say*), which would make (4*) true but unproblematic. (The idea would be: hedging allows us to approximate these truth-aiming ideals as best we

Thus, while Shapiro's own presentation of the case is unthreatening, the case itself genuinely poses a non-question-begging objection to epistemicism, and therefore deserves addressing.

5.3 Paradox: The Miners

As a first pass, one might extract the following argument from Shapiro's case. In abbreviated notation:[10]

The Forced Oughts Argument

(1) If SUNNY, Ought(BALLGAME)

(2) If ¬SUNNY, Ought(MOVIES)

(3) SUNNY ∨ ¬SUNNY

(5) Ought(BALLGAME) ∨ Ought(MOVIES)

Forced Oughts reasoning contends that the epistemicist commitment to bivalence about the weather forces one option to be required.[11] Yet comparison with the Miners Paradox quickly casts doubt on this claim.

The Miners case goes as follows.[12] Ten miners are trapped either in shaft A or in shaft B, but we do not know where. Water threatens to flood the shafts. We have enough sandbags to block

---

can, however imperfectly.) In contrast, there is no optional choice in Shapiro's case analogous to hedging (what would it be to only *sort of* go to the ballgame?); hence no easy deflationary response is available.

[10] A general note about terminology. Throughout the paper I shall use the following notions interchangeably: *should* / *ought* / *required / obligated / mandated* and *vague* / *borderline* / *unclear*.

[11] Here, the difference in choice of presentation seems to me unobjectionable, since Shapiro nevertheless appears to think the epistemicist is committed to the soundness of the Forced Oughts argument (see quote above).

[12] The example is from Derek Parfit (unpublished), who credits Donald Regan (1980: 265 n.1).

one shaft but not both. If one shaft is blocked, the other will be completely flooded, killing every miner inside. If neither shaft is blocked, both will be partially flooded, killing one miner.

Action	**If miners in A**	**If miners in B**
Block shaft A	All saved	All drowned
Block shaft B	All drowned	All saved
Block neither shaft	One drowned	One drowned

Lacking knowledge of the miners' exact whereabouts, it seems right to say

(9) We ought to block neither shaft

However, we accept both

(10) If the miners are in shaft A, we ought to block shaft A

(11) If the miners are in shaft B, we ought to block shaft B

And we know

(12) Either the miners are in shaft A or they are in shaft B

But (10)-(12) seem to entail

(13) Either we ought to block shaft A or we ought to block shaft B

which contradicts (9). Paradox ensues.

A detailed discussion of the various solutions to the Miners Paradox is unnecessary here. I limit myself to two observations. First, arguably the most viable strategy for dissolving the paradox is to show why premises (10)-(12) are not jointly sufficient to derive the paradoxical conclusion (13) (see Kolodny & MacFarlane 2010, Willer 2012).[13] The inference is invalid

[13] Niko Kolodny and John MacFarlane (2010) reject modus ponens to block the inference. Malte Willer (2012) blames disjunction elimination (or what I call constructive dilemma).

despite initial appearances. Second, the paradoxical inference in the miners case (reproduced below) shares the same logical form as the derivation of unknowable obligations in Shapiro's weather case. Both arguments are instances of constructive dilemma featuring statements of conditional obligation as premises.

The Miners Argument

(10) If IN A, Ought(BLOCK A)

(11) If IN B, Ought(BLOCK B)

(12) IN A ∨ IN B

(13) Ought(BLOCK A) ∨ Ought(BLOCK B)

Critics like Shapiro accuse epistemicism for licensing the Forced Oughts argument. Yet this is structurally identical to the Miners argument, whose failure we should already anticipate. Thus there is strong reason to believe the Forced Oughts argument fails too. Its conclusion—contrary to criticism—cannot follow simply from the epistemicist's commitment to (1)-(3).

What blocks the inference? Is an epistemicist committed to additional premises that would license it, or something equally bad, after all? I take up these questions in the next two sections.

5.4 Permissibility: epistemic explanations

I propose that failure of Forced Oughts reasoning is to be explained in epistemic terms. Again, comparison with the Miners Paradox is instructive.

The miners case demonstrates how deontic facts exhibit sensitivity to one's epistemic situation. Ignorance can be *obligation-defeating* in the following sense. If we knew which shaft the miners were in, we would be obligated to block that shaft. But as it is, we do not know, so no such requirement is in place; rather, our ignorance mandates blocking neither shaft.

The Miners premises capture this situation. Conditional obligation statements (10)-(11) merely express what would be required of us *if* we came to know the information given by the antecedent. Learning that information, either by being told of the miners' exact whereabouts or perhaps discovering it ourselves, *would* cause one of those obligations to become actual. But what is required of us *now* in our current state of ignorance is to block neither, as given by (9).[14] Our ignorance concerning the antecedent conditions of (10)-(11) serves to explain why neither action is mandated: we should not act as if we knew when we don't. Such ignorance is consistent with knowing (12): although we know that they are in one place or the other, we nevertheless do not know which.

This is why we know that the Miners argument fails well in advance of arriving at any particular solution to the Miners Paradox. We already find its conclusion (13) to be paradoxical prior to any substantive theorizing, because it contradicts (9) in the initial setup of the problem.

A parallel explanation awaits the Forced Oughts argument. Clarity in the weather would dictate which family outing is required, ballgame or movies. But when the skies are unclear, permission prevails: the father is not obligated to choose one over the other. This is because it is only borderline sunny outside, and the consequences of his promises are contingent upon the weather.

Here too, ignorance is obligation-defeating. Knowing the state of the weather would allow the family to know what they should do. But as it is, the weather is vague and, according to the epistemicist account, therefore unknowable. This ignorance serves to explain why neither

[14] Nothing significant turns on the fact that (9) is stated in the first person. Any omniscient third-party observer who does know the whereabouts of the miners (but does not communicate this information), when reporting on the situation, would agree that the agent should endorse (9), so long as he remains in the same epistemic situation.

activity is mandated. On this picture, the deontic facts still exhibit sensitivity to one's epistemic situation—epistemicism merely predicts that one's epistemic situation is worse off due to the vague state of one's environment.

The Forced Oughts premises capture this situation. Conditional obligation statements (1)-(2) merely express what the father should do *if* the antecedent were known to be true. But his ignorance of the facts—a consequence of the vague weather—prevents the relevant requirements from taking effect. What is required of him at present (if anything) is to simply choose one—even though neither activity is individually mandated, as stated by (6). Such ignorance is consistent with (3): there may be a fact of the matter about the weather, it simply remains unknowable. Thus the epistemicist commitment to bivalence—there being a truth or fact of the matter—is innocuous here.

In this way, we are able to recognize that the Forced Oughts argument, although apparently compelling, must ultimately fail. For we recognize that its conclusion (5) contradicts the permissibilist claim (6).[15]

---

[15] Other diagnoses of the Miners Paradox are available for appropriation. An obvious alternative is to deny the joint truth of the conditional premises (10)-(11) in the Miners argument, on grounds that neither shaft should be blocked, wherever the miners are (i.e. given (9) and (12), exactly one of the conditionals must be false—corresponding to whichever shaft the miners are in). One can likewise deny the joint truth of the conditional Forced Oughts premises (1)-(2), on grounds that neither action is mandated, whatever the weather is.

Another strategy is to deny that obligation must always be information-sensitive: perhaps the *oughts* in (10)-(11) express an "objective" notion (e.g. being the overall best course of action), in which case the conclusion of the Miners argument is unparadoxical. One can likewise insist that the Forced Oughts argument employs an information-*in*sensitive notion of *ought* (e.g. acting in accordance with one's promises); failure of epistemic transparency will then be unproblematic.

I leave aside other solutions to the Miners (see Kolodny and MacFarlane 2010 for details). I expect that, whatever one ends up saying about the Miners situation, a similar response may be made about Shapiro's case, given the structural similarity of the two. The present paper details merely one such response.

The epistemicist is entitled to deny (5) too. At least nothing in the epistemicist account considered so far suggests otherwise. Better yet, he can affirm (6). He simply needs to give an epistemic explanation, roughly of the form above, of how ignorance of the relevant facts defeats the relevant obligations. Fortunately for him, his very theory entitles him to that explanation by securing the link between vagueness and ignorance. It seems that epistemicists can make room for permissibilist intuitions after all. The situation is essentially the same as with the Miners: ignorance is obligation-defeating.

Although not entirely the same. One disanalogy is worth flagging. In the miners case, the agent should perform neither action, whereas in the weather case, the father should perform one.[16] Fair enough. But the point of the present analogy is simply to highlight one key similarity between the two cases: ignorance has the deontic consequence of defeating obligations that would otherwise obtain under different epistemic circumstances. Why ignorance should have other deontic consequences beyond that requires further substantive moral theorizing.

5.5 Proof: unknowable obligations

One may suspect the epistemicist of being secretly committed to principles that ultimately commit him to (5).

---

[16] Does the requirement to do something, rather than nothing, not mean that *some* thing—a single option—is required after all? No. Obligation can fail to distribute over disjunction, as when one must decide between choices, neither of which is mandated. Charlie is told to eat his vegetables, and given a choice between peas or carrots. Neither is individually required. Nevertheless he must choose one. At any rate, such worries properly concern the permissibilist claim that "Either is okay, but he must do one", rather than the epistemicist's commitment to bivalence.

To illustrate, here are some principles it seems an epistemicist may happily endorse:[17]

CLOSURE — If Φ entails Ψ, then: if clearly Φ then clearly Ψ

CONTRAPOSITION — 'If φ, ψ' entails 'If $\neg\psi$, $\neg\varphi$'

where what is clear is definite, not vague (in the relevant epistemicist sense). If vagueness and clarity are to be understood in epistemic terms of ignorance and knowledge, as epistemicists would have it, then those who have no qualms about epistemic closure will readily accept CLOSURE. The epistemicist's commitment to classical logic also commits him to CONTRAPOSITION, assuming a material conditional reading of *if*.

Yet these principles, together with some of the original premises in Shapiro's weather case, spell trouble. To see why, let us reconsider the conditional obligation expressed by the father's first conditional promise:

(1) If it is sunny, we ought to go to the ballgame

This seems not just true but *clearly* true. For, although the antecedent admits of borderline cases (and perhaps the consequent as well, as we shall see), the statement as a whole should hardly count as vague. After all, nothing in how the father makes the promise is vague; he speaks quite clearly and unambiguously when asserting it, without hedging or hesitation. At least it does not hurt to assume this. But then we may reason by way of closure and contraposition:

(14) a. Clearly(If SUNNY, Ought(BALLGAME)) — *by the clear truth of (1)*

b. Clearly(If ¬Ought(BALLGAME), ¬SUNNY) — *by CLOSURE, CONTRAPOSITION*

c. If Clearly(¬Ought(BALLGAME)), Clearly(¬SUNNY) — *by CLOSURE*

---

[17] Williamson himself accepts modified versions of these two principles. See Williamson 1994: esp. §§5.3, 8.2, 8.5.

d. If ¬Clearly(¬SUNNY), ¬Clearly(¬Ought(BALLGAME)) *by CONTRAPOSITION*

Now consider the converse of (1):

(1′) If we ought to go to the ballgame, it is sunny

This seems true, since if his obligation really is to bring his kids to the game, this can only be because things are in accordance with his first promise, and it is sunny outside. Nothing vague lurks here, so we may safely assume (1′) is not just true but *clearly* true. Then reasoning by way of closure and contraposition gives us:

(15) a. Clearly(If Ought(BALLGAME), SUNNY) *by the clear truth of (1′)*

b. If Clearly(Ought(BALLGAME)), Clearly(SUNNY) *by CLOSURE*

c. If ¬Clearly(SUNNY), ¬Clearly(Ought(BALLGAME)) *by CONTRAPOSITION*

The antecedents of (15c) and (14d) say that it is neither clearly sunny nor clearly not sunny—or simply, it is borderline sunny. The consequents of (15c) and (14d) say that the ballgame option is neither clearly mandated nor clearly not mandated—or simply, it is borderline mandated.

Consolidating (14)-(15), a second argument for unknowable obligations emerges from Shapiro's case.

(16) **The Vague Oughts Argument**

a. Clearly(If SUNNY, Ought(BALLGAME))
b. Clearly(If Ought(BALLGAME), SUNNY)
c. Vague(SUNNY)
d. Vague(Ought(BALLGAME))

A parallel argument begins with the clear truth of (2) and its converse, and proceeds to derive the result that the movie option is also borderline mandated.

Vague Oughts reasoning contends that vagueness about the weather makes each option vaguely required, such that one cannot know what is required, on the epistemicist assumption that vagueness entails ignorance.[18] The father is in a sort of double bind. He is vaguely obligated to take his kids to the ballgame, because it is vaguely sunny. But he is also vaguely obligated to take his kids to the movies, because it is vaguely not sunny. On the epistemicist account, since bivalence is preserved, there is a fact of the matter what he is obligated to do—it is simply unknowable.[19]

Poor man—to be potentially stuck with contrary obligations, each vague! One cannot do everything at once. Yet it appears that vagueness in the weather has somehow made the deontic situation vague too—all a result of making a couple of careless promises, though through no fault of his own (the weather is rather to blame). Such vagueness, uncontained, threatens to spread to the very keeping or shirking of those obligations (should any exist). This poses unsettling consequences. Suppose the father does not do what he should. Then whichever he fails to do, he will only count either as flouting an obligation which was *vague* to begin with or as *vaguely* breaking his promise—hardly an excuse for neglected duties. Alternatively, suppose he acts as he should. Then whichever he succeeds in doing, he only counts either as fulfilling an obligation which was *vague* to begin with or as *vaguely* keeping his promise—hardly a vindication of dutiful action. Both these possibilities fly in the face of the permissibilist intuition that *neither* option—

[18] By the locution *x is vaguely F*, I mean *it is vague whether x is F*, which is consistent with *x is not F*. Hence it does not follow from something's being vaguely obligated that it is obligated, but only vaguely.

[19] Perhaps neither is mandated, though this fact would be vague and hence unknowable. This possibility differentiates the current argument from the original weather argument, according to which (on one reading of the example) it is mandatory to do at least *one* of the two options, ballgame or movies. See also n.20.

ballgame or movies—is mandated, not even vaguely.[20] What has gone wrong? The next section describes a way out for the epistemicist.

5.6 Proposal: non-classical norms

Taking seriously the analogy with the miners case again offers a solution. Arguably, any viable solution to the Miners Paradox must reject modus tollens. This can be seen by considering a simpler paradoxical argument which Kolodny and MacFarlane (2010) present about the miners:

(17) **The Simple Miners Argument**

a. If the miners are in shaft A, we ought to block shaft A

b. <u>It is not the case that we ought to block shaft A</u>

c. The miners are not in shaft A

where (17a) is simply premise (10) in the original miners argument and (17b) is entailed by our original premise (9). Yet concluding (17c) is unjustified: assessment of our deontic situation should not allow us to deduce the miners' location. Rejecting modus tollens is required to block the simple miners argument. Theorists who accept the premises of the original miners argument thereby incur the responsibility of explaining why modus tollens fails. It should be unsurprising

---

[20] Granted, being vaguely obligated to do one thing is (we may concede) compatible with being allowed to do otherwise (see n.18). For it is arguably fallacious to think vagueness precludes truth or falsity (see Wright 2003). If so, the presence of vague obligations is strictly speaking compatible with the permissibilist assumption (6). Such a combination of claims nonetheless remains unattractive to all. The epistemicist would be forced to conclude, implausibly, that although the father is free to do either, this itself is unknowable. Otherwise, knowing this would entail knowing that neither option is individually mandated (assuming epistemic closure and the duality of permission and obligation)—contrary to its being vague and hence unknowable of each whether it is obligated. Nor can the semantic permissibilist (see §7) accept the Vague Oughts conclusion (16d), since this would allow the relevant requirements to be arbitrarily settled—contrary to our assumption (6) that the matter is already settled since neither option is mandated. I presume our permissibility intuitions are strong enough to warrant strengthening (6) to be not just true but *clearly* true. Thanks to Jeff Snapper for pressing this point.

that failure of modus tollens should translate into failure for the closely related inference rule of contraposition. Indeed, both rules fail on the semantics proposed by Kolodny and MacFarlane (2010) and Willer (2012) in their solutions to the Miners Paradox.[21]

This leaves us with a conditional that does not contrapose or license modus tollens. Such is the nature of statements expressing conditional obligations. This is at least an alleged lesson of the Miners Paradox. We have, I am suggesting, equal reason to believe this about cases like Shapiro's involving conditional obligations with vague antecedents.[22]

---

[21] Willer (2012) and Kolodny and MacFarlane (2010) give competing explanations of why modus tollens fails. And although contraposition goes unmentioned in their discussions, they would agree this fails too, though disagree in the details. Consider the canonical proof for contraposition.

1	$\mid \underline{\text{If } \phi, \psi}$	
2	$\mid \quad \mid \underline{\neg(\text{If } \neg\psi, \neg\phi)}$	assume for reductio
3	$\mid \quad \mid \quad \neg\psi$	$\rightarrow$E, 2
4	$\mid \quad \mid \quad \neg\neg\phi$	$\rightarrow$E, 2
5	$\mid \quad \mid \quad \phi$	$\neg\neg$E, 4
6	$\mid \quad \mid \quad \psi$	MP, 1, 5
7	$\mid \quad \mid \quad \bot$	$\bot$I, 1, 5
8	$\mid \text{ If } \neg\psi, \neg\phi$	reductio, 2-7

Both accounts would object to line 6, but for different reasons. Kolodny and MacFarlane would object to the use of modus ponens—presumably because modus ponens only delivers information about what happens in ϕ-shifted information states; it fails to say what follows given the current unshifted information state. Whereas Willer would object to importing line 1 within the subproof—presumably for reasons of nonmonotonicity: the truth of 'if ϕ, ψ' fails to be preserved under the additional assumption of $\neg$(if $\neg\psi,\neg\phi$).

[22] Whether the epistemicist must surrender classical logic and semantics depends on the details of the analysis for *if*. It is at any rate not obvious why this should be so bad a result. Efforts to relax traditional epistemicist assumptions are not unprecedented, however much these are in the minority. Hybrid epistemicist accounts that combine epistemic and semantic approaches have been proposed by Goguen (1969), Koons (1994), Graff (2000), Barker (2002), Raffman (2005), MacFarlane (2010), Akerman & Greenough (2010).

Rejecting contraposition blocks the derivation of vague obligations.[23] More exactly: the use of CONTRAPOSITION in (14b) is dubious. (I shall set aside the occurrences of contraposition elsewhere, as I think the problem can be isolated here.[24]) That is because (1) does not entail its contrapositive. Although it is true that sunny weather mandates going to the game, it is not true that the absence of that mandate means the weather is not sunny. After all, the father is not required to take his kids to the game, but recognizing that fact does not somehow enable us to deduce facts about the weather.

Rejecting modus tollens invalidates the analogue of the simple miners argument:

(18) **The Simple Weather Argument**

a. If it is sunny, we ought to go to the ballgame

b. It is not the case that we ought to go to the ballgame

c. It is not sunny

where (18a) is simply premise (1) in the original weather argument and (18b) is entailed by the original permissibilist claim (6). The argument is unacceptable. Surely, concluding (18c) is unjustified: acknowledging the father's free choice should not allow us to deduce facts about the

---

[23] An alternative derivation applies modus tollens to (14c) and (15b), together with (3), to deliver the paradoxical result of vague obligations. However, this fails to show that contraposition is inessential to the proof. For it is hard to see how to get (14c)—which, considered alone, is already implausible—without first deriving it from (14b), which is in turn gotten from the objectionable rule of contraposition (or so I complain).

[24] The other occurrences of contraposition in (14d) and (15c) appear unproblematic for the following reasons.

If *clearly* neither option is mandated (as I claim; see n.20), then the ballgame option is clearly not mandated; so the antecedent of (14c) is true and the consequent of (14d) is false. But since the weather is vague (in particular, it is not clear that it is not sunny), the consequent of (14c) is false and the antecedent of (14d) is true. Hence (14c) and (14d) are both classically false. This makes the inference from (14c) to (14d) unproblematic, even if valid.

Similarly, since neither option (including going to the ballgame) is clearly mandated, the antecedent of (15b) is false and the consequent of (15c) is true. And since the weather is vague (in particular, it is not clear that it is sunny), the consequent of (15b) is false and the antecedent of (15c) is true. Hence (15b) and (15c) are both classically true. This makes the inference from (15b) to (15c) unproblematic, even if valid.

weather. It appears that both simple miners-type reasoning and simple weather-type reasoning rely solely upon modus tollens, so the culprit is obvious: it is modus tollens we should get rid of (or otherwise restrict).

5.7 Parity: semantic vs. epistemic permissibilism

We have seen how the epistemicist can make room for permissibility intuitions, provided an appropriate, independently motivated logic for deontic conditionals. Freed of inference rules like constructive dilemma and contraposition, he can avoid commitments to the otherwise troubling results of forced or vague obligations.

This undercuts Shapiro's motivation for moving away from traditional theories of vagueness like epistemicism.[25] The alternative he favors is a kind of *semantic permissibilism* on which, as a constitutive feature of their vagueness, the meanings of vague terms exhibit an element of "open texture": competent speakers are generally free to judge either way in vague matters.[26] Such semantic permissibilism about vagueness is not the only way to secure "no forced choice" predictions like (6). Epistemicism can too. It would seem the two sorts of theories are on par

---

[25] "Traditional" in the sense of *old* (in contrast e.g. with the more recent upsurge of contextualist approaches). One may worry the reply undercuts traditional motivations for epistemicism too—namely, the preservation of classical logic and semantics in their entirety, including the classically valid inferences of constructive dilemma, modus tollens, and contraposition. This is worrisome insofar as one follows Williamson—whose own defense of bivalence employs these very such rules—in taking this to be *the* driving motivation for epistemicism. But Williamson's way of motivating epistemicism is not the only way. Realist intuitions in vague matters ("there is a fact of the matter, whether we know it or not") are enough to offer *direct* support for the view, without detouring through classical inferences. Nor is Williamson's version of epistemicism the only version (see n.22). As pointed out earlier, the reply developed here is meant for the epistemicist who does not already subscribe to Williamson's other views on assertion (see n.5). I wish to explore these points in a future paper.

[26] Defenders include Waisman (1968), Kamp (1981), Sainsbury (1990), Soames (1999), Kyburg & Morreau (2000), Wright (2001), Shapiro (2006), Akerman & Greenough (2009), Gaifman (2010).

with regard to predictive power. Their differences, I suggest, nevertheless lie in competing explanations for this fact.

The epistemicist, I have argued, should say that the father's ignorance allows him to do either without fault. In deciding, however, he does *not* thereby come to fulfill either of his original conditional promises. For which would it be? He would not know which has been satisfied, since the weather remains unknowable. Rather, what gets fulfilled, if anything, is the further, tacitly implied (or perhaps implicated) promise that he and the children will do one or the other.[27]

The semantic permissibilist, I suspect, will offer a competing line of thought.[28] The father is free to choose, but only by arbitrarily settling upon a decision about the weather: sunny or not? Recall that the semantic permissibilist's claim of free choice applies to the weather, not to the promised action—for it is the weather, not his obligations, that are vague. Upon choosing, he thereby decides which of the two original conditional promises gets to be kept. Saying it is sunny means they must go to the ballgame; saying otherwise means the movies. Either way, a decision has the effect of reimposing one of the conditional obligations. Otherwise, the father could legitimately say, for instance, "We'll say it's sunny, but let's go to the movies anyways"—which is to go back on his first promise.

---

[27] Stewart Shapiro (in private communication) has clarified that the intended reading of his example did not involve any such promise. Nothing in the setup, however, is essentially altered by including it. Filling in the details of the example makes this salient. The father might, for instance, explicitly say they'll do one or the other, where the weather is simply a useful heuristic to help them randomly decide. Alternatively, he might promise a ballgame outing as a first option in case of sunny weather, then promise a movie as a backup option in case of non-sunny weather. Regardless, there seems at least to be an implicature that they will do something fun, where this is limited to two options. This would explain why the father is at fault should he fail to do anything fun, and why he remains at fault (perhaps to a lesser degree) should he pursue some fun activity *other than* a ballgame or movie, in the case of vague weather. (Although nothing substantial should rest on this.)

[28] Notably, Shapiro's own line of thought (2006:82-3).

What better explains (6)? Is the father exempt from mandated choice because it is unknown which conditions of his promises have been satisfied, or because these are easily and arbitrarily decided? The epistemicist's explanation seems to win out on the score of plausibility.[29] To claim that the father and his children cannot simply choose an outing and leave the weather undecided, but must rather *first* settle on a decision about the weather—sunny or not—in order to remain faultless in their quest for family fun, is nothing short of incredible. One decides what to do by deciding to do it, not by deciding the weather.

5.8 Pairings: epistemic open texture

My partial defense of epistemicism is non-committal about any "open texture" thesis about the alleged flexibility of vague words. The present point is simply that no such assumption of open texture needs to be made in order to make sense of Shapiro's case. Indeed, the semantic permissibilist's insistence that the father, in order to remain blameless, *must* exploit the open texture of 'sunny' by first declaring the state of the weather before acting is objectionable.

That said, nothing said so far precludes the epistemicist from endorsing limited claims of open texture, while denying that these are constitutively characteristic of vagueness (which for him is fundamentally epistemic, not semantic, in nature). He may allow for the possibility of reimposing one of the conditional obligations, either by settling the facts about the weather (by fiat, as in "Let's say it's sunny") or by repositioning one's own epistemic situation with respect to

---

[29] On other scores too—such as explaining why weather-type cases of conditional obligation are often best solved through compromise (taking the kids to a movie about baseball, like *Angels in the Outfield*) or explaining why vague matters of serious moral importance (not aborting a fetus if it is e.g. a person or past the point of viability) strongly appear *not* to admit of arbitrary decisions in settling the matter. I wish to elaborate on these issues in a future paper.

the facts (perhaps by consulting weather experts who are better at judging whether the skies are sunny—if any such division of epistemic labor exists).[30] In that case, stipulating or coming to know that it is sunny (or not sunny) *would* mandate choosing the ballgame (or movies) after all. Normative considerations of obligation and liability, previously rendered inert by ignorance, are reintroduced once the relevant obstacles to knowledge are removed. Yet the nature of the explanations underlying these normative facts remains thoroughly epistemic throughout: it is ignorance that acquits and knowledge that commits.

One might dispute any such epistemicist endorsement of open texture, on grounds that committing to speak truthfully about vague matters would, according to epistemicism, require one to give a polar verdict of 'F' or 'not F' even in borderline cases—contrary to any open texture permission to judge borderline cases either way. Recall argument (1^*)–(4^*) from §2 (reproduced below). Assuming we promise to speak the truth when describing the weather:

(19) **The Modified Weather Argument**

a. If it is sunny, we ought to call the weather 'sunny'
b. If it is not sunny, we ought to call the weather 'not sunny'
c. <u>Either it is sunny or it is not sunny</u>
d. Either we ought to call the weather 'sunny' or we ought to call it 'not sunny'

---

[30] One might resist the idea of giving any such local resolutions, or "precisifications" as it were, to borderline cases, should he already believe the (however unfairly) strong thesis that any vagueness-related ignorance is impossible to remove, even partially, by making further stipulations or gaining more evidence. This, however—like so many other features unique to Williamson's own peculiar brand of epistemicism—is just another purely optional commitment for the epistemicist. Although it is worth noting that Williamson himself admits a kind of open texture for partially defined predicates: stipulative completions of their meanings, he claims, are able to reverse the old truth-values of borderline cases, and thereby make precise what was formerly imprecise. See Williamson 1997: 226.

The falsity of (19d) in the case of borderline sunny weather, given open texture, would appear to serve as a reductio of the epistemicist premise (19c), demonstrating that this unusual pairing of views—epistemicism with open texture—is incoherent after all.

Such reasoning, however, is invalid. The diagnosis should be obvious, given the general line of argument I have pursued so far. The conditional obligation to call an F 'F' (or to call a non-F 'not F') is defeated by one's ignorance of its Fness (or its non-Fness) when it is a borderline case. As with the weather and the miners arguments, (19) exemplifies invalid reasoning by constructive dilemma.[31] Reasoning by modus tollens in (20) fails for analogous reasons:[32]

(20) **The Simple Modified Weather Argument**

a. If it is sunny, we ought to call the weather 'sunny'

b. <u>It is not the case that we ought to call the weather 'sunny'</u>

c. It is not sunny

In this way, epistemicism leaves room for individual discretion when applying vague predicates to borderline cases and acting accordingly.[33] Epistemicism does predict that there are facts of the matter underlying all vague matters, but these facts, I hope to have shown, are (at least oftentimes) normatively inert. Whether they must remain so is an open question. Just as the presence of vagueness is not without normative consequences, neither should the removal of

---

[31] Disanalogies persist. Unlike the miners case, one needn't (given open texture) refrain from either choice. Unlike Shapiro's case, one also needn't call it one or the other (unless forced to, as in a "forced march" Sorites).

[32] Alternatively, an epistemicist unsympathetic to open texture, like Williamson, may dismiss (19) and (20) as question-beggingly assuming the truth of something (i.e. open texture) already incompatible with the target view.

[33] Soames (2012) in effect disputes this claim within legal contexts, on grounds that epistemicism leaves no room for individual judiciary discretion over interpreting vague statutes of the law when deciding the legal status of borderline cases. I leave it an open question as to whether epistemicism is objectionable on these other grounds.

vagueness be automatically regarded as normatively inconsequential: removing epistemic obstacles deriving from vagueness can alter one's normative situation. That someone might achieve this effect by exploiting the "open texture" nature of vague terms is, for all that has been said, compatible with the epistemicist account.

In the following chapter, I develop a theory of vagueness-related reasoning with deontic conditionals, according to which the Forced Oughts, Vague Oughts, Simple Weather, Modified Weather, and Simple Modified Weather arguments are all deductively invalid but nonetheless defeasibly valid. The resulting account will exploit certain features of normative reasoning about vagueness already explored here, including the epistemic construal of open texture, in order to provide a proper diagnosis of the sorites paradox. This will demonstrate just how much a permissibilist version of epistemicism can achieve.

Chapter 6: Defeasible Tolerance and the Sorites

6.1 Worries for contextualism

Consider a series of 1,000 men, each with just slightly more hair on his head than the man before him, the first of which is Patrick Stewart who is completely bald, the last of which is Howard Stern who is completely hairy. Standard sorites reasoning attempts to take us from the reasonable judgment that Patrick Stewart is bald to the unreasonable conclusion that Howard Stern is also bald, by way of the inductive principle:

(1) For each n < 1,000: if #n is bald, then #n+1 is bald

or, alternatively, by way of a "step-by-step" sorites argument that does not invoke a single inductive premise but rather lists out its individual instances:

(2) #1 is bald
If #1 is bald, #2 is bald
If #2 is bald, #3 is bald
Etc.
<u>If #999 is bald, #1,000 is bald</u>
#1,000 is bald

Both (1) and its instances in (2) are independently motivated by the alleged tolerance of vague predicates like 'bald': comparing any two adjacent #n and #n+1 in our series, if I judge that #n is bald, I cannot go on to judge that #n+1 is not bald, given that the two differ only slightly in hair number.

Contextualists maintain that sorites contradiction can be avoided without jettisoning our tolerance intuitions. Instances of the inductive premise can be true when relativized to appropriate contexts. Sorites reasoning then relies crucially upon deploying each conditional in a context in which it is true—be this the shared conversational scoreboard (Shapiro 2006), a relevant set of speaker-relative interests (Fara 2000), mutual agreements or stipulations (Soames 1999), an individual's internal psychological state (Raffman 1994), or salient background information (Kamp 1981). There is no way, however, of fixing or otherwise coordinating the contexts of evaluation to make this possible; vagueness removes any guarantee for stability of context. Contextual variation thus prevents us from carrying such reasoning all the way through to absurdity. Sorites reasoning must ultimately fail, however compelling each step of it may be.[1]

Contextualist solutions face the following worry. They all seem to predict failure of sorites reasoning relatively late in the sorites series, toward the last few borderline cases of baldness where people start to definitely *not* look bald—presumably, because this is where the subject's judgments are clearly mistaken, should he continue to extend his use of 'bald' as required by tolerance and begin applying it to those who count as not being bald in any context. But intuitively, sorites reasoning starts to look suspicious *well before* the end of the borderline range. Yet no existing contextualist account ventures to explain why.

I shall lay out and develop an alternative view—really, an instance of an alternative *kind* of view—which I shall argue better accounts for our unease when reasoning through a sorites series. Doing so involves formulating sorites conditionals in a way that properly expresses the deontic content of tolerance principles—something left out of formulations like (1) and (2), which, as we shall see, fail to adequately capture the sort of reasoning actually deployed when deliberating how far to extend a vague predicate through a sorites series. This makes room for more than one type

---

[1] Even in a "forced march" scenario (see §6), where one is forced to make a judgment about each #n—bald or not—sorites reasoning must give out: one will "jump" and start calling things not-bald before the end of the series.

of tolerance. Properly accounting for the failure of sorites reasoning, in turn, requires differentiating among these options.

This undercuts the chief motivation behind contextualist theories, whose preservation of tolerance is an alleged major advantage over competitor theories, including standard supervaluationist treatments. Even if contextualists succeed in this respect, this achievement is not unique to contextualism: viable alternatives exist. Indeed, existing contextualist theories are to be found wanting in two critical respects concerning the very phenomena they purport to explain: one, they leave certain facts about the phenomenology of sorites reasoning unaccounted for, and two, they fail to distinguish between different types of tolerance. My account, I propose, does better on both scores, by fitting sorites reasoning within the framework of a nonmonotonic deontic logic, without appeal to any contextual relativization of evaluation. The general lesson that will emerge, if I am correct, is that contextualist accounts are generally right in maintaining that sorites reasoning should be rejected without faulting tolerance, but wrong in how they implement this idea.

6.2 Tolerance principles

Various constraints guide our use of vague predicates. Among the relatively uncontested are: recognizing clear cases for what they are (Patrick Stewart counts as 'bald' by anyone's lights), and not what they aren't (by nobody's lights does Howard Stern count as 'bald' or Patrick Stewart as 'hairy'); maintaining consistency in one's judgments, either about a single case (not calling someone both 'bald' and 'not bald') or across cases (whatever standard is in place for 'bald', anyone in the same hair situation as or balder than someone who meets the standard himself meets the standard); coordinating one's standards of evaluation when using related predicates (whatever the standards are for ‘bald’ and ‘hairy’, nothing can meet both, although quite plausibly some things also meet neither). Applied to our series for ‘bald’, the relevant constraints entail, for $0 < n < 1,000$:[2]

CLEAR CASES	For some k, #k is clearly bald (not-bald, hairy)
CLEAR CALL	If #n is clearly bald (not-bald), one should call #n bald (not-bald)
CLEAR SANCTION	If #n is clearly bald (not-bald), one should not call #n not-bald (bald)[3]
CONTRAST	If one calls #n bald (hairy), he should call #n not-hairy (not-bald)
MONOTONICITY	For $n \geq k$, if one calls #n bald, he should call #k bald

More controversial are so-called constraints of tolerance. The alleged tolerance of vague predicates is generally understood to be a feature of our judgments: a constraint on the way those predicates may be used. Take Crispin Wright’s classic definition:

> “*F* is tolerant with respect to ϕ if there is also some positive degree of change in respect of ϕ insufficient ever to affect the *justice* with which *F* is applied to a particular case.” (Wright 1975:334, my emphasis)

On one widespread understanding, tolerance is the constraint that marginal differences between two cases in the relevant parameter of application never allow for differential verdicts. Suppose a subject is led through our sorites series. For any pair of adjacent men in the series he is presented

---

[2] These are left stated in their ordinary English formulations, open to further refinement on each individual theory (e.g. by adding contextual parameters). Given that these principles, properly understood, are supposed to be exceptionless, I see no reason to analyze them as making generic conditional claims, but shall save the notion of exception-allowing generic implication for more controversial constraints, such as tolerance, below. Here, I shall simply assume that the standard material conditional reading of *if* in these principles is unproblematic.

[3] This is redundant and follows from CLEAR CALL, on the assumption that Op and $O\neg p$ are contradictory (see §3).

with, if he judges that one is bald, he is thereby committed to judging that the next one is bald too.

Many have followed Wright in tying vagueness to the initial plausibility of tolerance.[4] This includes nearly all contextualists about vagueness. They maintain that tolerance principles govern our use of vague predicates, but contend that these principles are free of sorites-related contradictions, so long as the relevant judgments are relativized to appropriate contexts: contextual variation then explains how a subject who is led through a sorites series can retain consistency in his judgments without ever violating the relevant tolerance principles (Kamp 1981, Raffman 1994, Soames 1999, Fara 2000, Shapiro 2006).[5] There may be circumstances in which

---

[4] "*Initial* plausibility", since Wright in fact disavows their tolerance, given the threat of contradiction—and with it, classical logic and a rule-based conception of meaning—though he thinks any adequate account should explain the allure of these things (see Wright 1987:§4). In contrast, the current proposal (as well as most contextualist theories) promises to salvage all three features of vague predicates—their tolerance, coherence, and rule-based meaning—through less radical revision of classical logic and semantics.

[5] "[…] we must recognize an object as falling within the positive or negative extension of such a predicate if it is indistinguishable (in the relevant respect) from an object which has already been accepted as belonging there." (Kamp 1981:243)

"There is […] no sharp division between objects that are clearly red and objects that aren't (clearly red), people who are clearly rich and people who aren't. The vagueness of these predicates goes hand in hand with what Crispin Wright has aptly called their "tolerance": they tolerate marginal changes in the parameters decisive of their application." (Raffman 1994:41)

"For any two patches of color *x* and *y* that are perceptually indistinguishable to competent speakers under normal conditions, if someone who is presented with *x* characterizes the predicate *looks green* as applying to it, then that person is thereby committed to a standard that counts the predicate as applying to *y* as well." (Soames 1999:215)

"[If] two men are pretty much the same height—one is just noticeably shorter than the other—then the option is not available to me to say that one is tall but the other is not. Because the similarity of their heights is so perceptually salient—and now that you've asked me whether they're tall, also conversationally somewhat salient—I may not choose a standard that one meets but the other doesn't." (Fara 2000:59)

"Suppose that two objects *a, a′* in the field of *P* differ only marginally in the relevant respect (on which *P* is tolerant). Then if one competently judges *a* to have *P*, then she cannot competently judge *a′* in any other manner." (Shapiro 2006:14) tolerance is no longer in force—as with Sainsbury's paint shop owner who, in order to divide the reds from the oranges along a chromatically ordered line of paint jars, must choose an arbitrary cutoff point and label at least one pair of adjacent paints "red" and "orange", though these are indistinguishable in color (Sainsbury 1990). However, such examples notwithstanding, contextualists agree that tolerance remains in force in *most* circumstances of use for a good many paradigm cases of vague predicates, like "red", "tall", "bald" (see Shapiro 2006:9).

Tolerance is widely understood to be a semantic feature of vague predicates. Tolerance principles embody semantic norms. Compliance with these principles is part and parcel of semantic competence. To flout tolerance is somehow to sin against language, betray semantic incompetence, violate linguistic convention, or use words in ways that offend against their meaning. Alternatively, tolerance principles may be seen as encoding a type of epistemic rather than semantic norm, and their deontic content properly understood in epistemic or broadly psychological terms (Fara 2000; Raffman 1994, 1996).[6] Regardless, tolerance principles are upheld by many, even by those who affirm its contradictory consequences.[7]

---

[6] Alternative views prescribe treating tolerance as arising from both distinctively epistemic and distinctively semantic phenomena (Koons 1994); as involving both epistemic and semantic components, where these are inseparable features of the general psychology of vague language (Sorensen 2001); or as not reducing to any combination of epistemic or semantic factors, since vagueness is to be understood as *sui generis* (Barnett 2010).

[7] *Incoherentists* about tolerance claim that they commit us to sorites-type inconsistencies: full compliance with tolerance *would* issue in contradictory judgments (Sorensen 2001, Eklund 2002). *Nihilists* about vagueness take this as grounds for denying the very coherence of vague predicates themselves, and not just their rules for use: either vague predicates are entirely empty in their application (Dummett 1975) or vagueness as traditionally conceived is simply impossible (Horgan 1994). Both camps claim, nonetheless, that semantic competence requires being disposed to judge things in accordance with tolerance principles. Of course, plenty of perfectly competent speakers when led through a sorites series, in order to avoid contradiction, do *not* fully comply with tolerance, yet their semantic competence remains (on the face of it) fully intact. If anything, a willingness to obey tolerance and make the contradictory judgments would show them to be semantic *in*competent. Thus the incoherentist must conclude that semantic competence requires *both* that certain judgment dispositions be in place *and* that those same dispositions never be completely

Indeed, much of the language employed in articulating these principles is blatantly *deontic*. It is said, for example, that when comparing two things that are indistinguishable or marginally different in the relevant respect, one who judges one to be F is constrained in how he may judge the other insofar as he:

- "*must* recognize [it] as falling within the positive or negative extension" (Kamp 1981:243)
- "is thereby *committed to a standard* that counts the predicate as applying to [it] as well" (Soames 1999:215)
- "*may not choose a standard* that one meets but the other doesn't" (Fara 2000:59)
- "*cannot competently judge* [it] in any other manner" (Shapiro 2006:14)

Such uses of deontic concepts as these (italicized for emphasis) proliferate in glosses of tolerance.

Whatever their nature, the requirements of tolerance are *conditional* in form: they prescribe how one should judge the next thing *if* he has already judged the first of a pair of things that are only marginally different.

All this strongly suggests that the prescriptions given by tolerance, when properly articulated, take the general form of deontic conditionals. An appropriate tolerance principle for, say, 'bald' should entail:[8]

(3) For each n < 1,000: If one calls #n bald, then one should call #n+1 bald

---

given into. Linguistic competence is on this view a sort of self-control: one must have, but never follow, his inclinations. My view avoids this odd consequence by freeing tolerance of contradiction.

[8] Properly formulated, 'bald' should be encompassed in quotes, to indicate that tolerance directly concerns e.g. our use of the predicate 'bald', rather than facts about baldness. For simplicity of exposition, however, I shall ignore conventions of mention vs. use and drop the quotes in what follows.

Thus an adequate semantics for *ifs* and *oughts* is desired if we are to express properly the requirements of tolerance in the object language. Granted, the deontic content of tolerance conditionals rarely if ever gets explicitly expressed. Standard formulations of tolerance principles tend to suppress any mention of "obligation", "norm", "competence" and the like—a regrettable if not odd fact, given that such deontic talk is liberally used to motivate these principles in the first place.

Against this trend, the view on offer makes explicit the deontic content of tolerance principles, in a way that properly captures their conditional form while rendering them free of sorites-type contradiction. Against contextualist approaches, I claim this can be done without relying upon contextual relativization of judgments. The presentation shall remain neutral on whether the obligations in question are epistemic or semantic in nature, though I shall later offer some reasons for favoring the epistemic interpretation.

6.3 Generic ifs and oughts

I contend that a theory of conditional obligation, in which the underlying logic is nonmonotonic, can explain both the appeal and the paradox-free nature of tolerance principles governing our use of vague predicates that have motivated contextualist views. The view I shall develop renders our tolerance intuitions intact and unproblematic, provided a notion of defeasible entailment that allows for conclusions to be overridden by further information. Tolerance principles license inferences in the absence of contradictory information. The failure of sorites arguments is then easily attributed to the failure of reasoning by tolerance, which, although defeasibly valid, cannot be sustained throughout the entire series. Resolving the sorites in this way removes a core motivation behind contextualist theories—that of preserving tolerance by the relativization of judgments to contexts. The resulting account, I hope to show, does not face the problems suffered by competing contextualist accounts.

Classical deductive logic is monotonic in the sense that conclusions, once established, stay established. Adding premises to a valid argument always produces another valid argument. In nonmonotonic logic, however, adding a premise may make a valid argument an invalid one. Conclusions drawn in a nonmonotonic system are defeasible—licensed in ordinary circumstances but possibly surrendered in light of further information. To make good on this idea, one needs a variable, decentered conditional. There are a variety of ways to make this precise. Here, I shall adopt Asher and Morreau's (1991) nonmonotonic system of commonsense entailment (see also Morreau 1997). They introduce a generic conditional, symbolized as $>$, with truth conditions

$p > q$ is true at a world w iff q holds in all p-normal worlds relative to w

and constraints

FACTICITY	p is true in all p-normal worlds
DISJUNCTION	$(p \vee q)$-normal worlds are either p-normal or q-normal

FACTICITY guarantees the truth of $p > p$. DISJUNCTION secures $((p > r) \& (q > r) \rightarrow ((p \vee q) > r)$. The absence of any centering constraint allows for the possibility that w is not a p-normal world relative to itself. This makes the generic weaker than the Lewis counterfactual, since the p-normal worlds relative to w need not be the closest p-worlds to w.

How can we make precise the idea that some conclusions are reasonable even though their truth is not guaranteed? The key idea (from Morreau 1997) is that some worlds are more regular than others, in the sense that they involve fewer exceptions to principles, or conflicts between principles, than others.[9] Here is Morreau's definition. Say that a world w is *irregular with respect to p* iff for some q, (p > q) ⊃ (p ⊃ q) is false at w (i.e. p > q and p are true, but q is false, so there is a modus ponens failure with respect to p). A world is *regular with respect to p* if it is not irregular with respect to p. Notions of absolute, equal and comparative (ir)regularity are then defined in terms of relative irregularity.

- world w is *irregular simpliciter* iff w is irregular with respect to some p
- world w is *as regular as* w′ *with respect to p* iff whenever w is irregular with respect to p, so is w′
- world w is *more regular* (or *less irregular*) *than* w′ *with respect to p* iff w is as regular as w′ but not vice versa

A *counterexample* to an argument is a world where the premises are true and the conclusion is false. Say that a counterexample to an argument is *unnecessarily irregular* or *gratuitous* if it contains irregularities not required by the truth of the argument's premises—that is, if there are worlds more regular than it in which the premises and conclusion are all true. We shall call an argument *deductively valid* or *valid* if it has no counterexamples, and *allowed* or *defeasibly valid* if all counterexamples to it are gratuitous. Defeasible validity is the weaker notion: all valid arguments are (vacuously) allowed, but not all allowed arguments are valid. For any defeasibly valid argument, we say that its premises *defeasibly imply* or *defeasibly entail* its conclusion (or that the latter is a *defeasible consequence* of the former).

---

[9] The account of defeasible entailment outlined below is informally presented in Morreau (1997) and developed by Bonevac (forthcoming). An alternative (though more complicated) treatment of defeasible validity may be found in Asher & Morreau's (1991) original presentation of their theory of commensense entailment. This is summarized in Bonevac (1998:44); see also Asher & Bonevac (1996:§2).

With these notions in place, let us now comment on several key features of the generic conditional.

I. Default Detachment. The generic > does not support modus ponens: p, (p > q) |/= q. This is because p and p > q can fail to entail q, so long as the world of evaluation w is not p-normal relative to itself. Nevertheless, modus ponens on > is defeasibly valid: p, (p > q) |≈ q, where '|≈' represents defeasible entailment. The reason is that any counterexample to the argument p, (p > q) ⇒ q would have to be unnecessarily irregular. That is because for any counterexample demonstrating a failure of modus ponens on p > q, there will be a world where modus ponens survives, such that {p, p > q, q} are all true and moreover every generic consequence of p is true (i.e. any r is true if p > r is true). The model world will thus be more regular than any counterexample world, thereby making all counterexamples gratuitous.

II. Defeated Detachment. Detachment fails in particular when information incompatible with the conclusion is added to one's starting premises: p, (p > q), ¬q |/= q. The argument is not even defeasibly valid. This is because no counterexample to the argument is gratuitous, since there can be no worlds that are more regular (indeed, no worlds at all) that make both its premises and conclusion true.

III. Nixon Diamond. Given the defeasibility of detachment, conclusions that would otherwise follow are suspended in Nixon Diamond situations, where conditional principles conflict but neither is stronger than the other. It is generally true that Quakers are pacifists and that Republicans are not pacifists. Yet this is insufficient information to conclude anything about Nixon, who is both Quaker and Republican. Symbolically: {p > r, q > ¬r, p, q} entails neither r nor ¬r, not even defeasibly. Consider the two arguments p > r, q > ¬r, p, q ⇒ r and p > r, q > ¬r, p, q ⇒ ¬r. Any model for one (making all its premises and conclusion true) will be a counterexample for the other, where neither is more regular than the other; so neither r nor ¬r is a defeasible consequence.

IV. Generic Specificity. We also want specific conditionals to take default precedence over less specific ones, so that {p > r, (p & q) > ¬r, p, q} |≈ ¬r, in order to validate Penguin Principle inferences:

(4)	Tweety is a bird	b
	Birds fly	b > f
	Penguins are birds	p ≺ b
	Penguins don't fly	p > ¬f
	Tweety is a penguin	p
	Tweety doesn't fly	¬f

First observe that FACTICITY and DISJUNCTION allow more specific defaults to take precedence over less specific ones, which effectively allows us to derive the conclusion that birds aren't normally penguins (b > ¬p).[10] Now consider a counterexample to the Penguin Principle where {b, b > f, p ≺ b, p > ¬f, p, f} are true. This will exhibit a modus ponens failure on p. By contrast, any model making both premises and conclusion true (i.e. all of {b, b > f, p ≺ b, p > ¬f, p, ¬f}) will exhibit a modus ponens failure on b. Yet this does not share the symmetry of the Nixon Diamond setup, because the counterexample involves an *additional* modus ponens failure on b—

[10] *Proof.* Let p: Tweety is a penguin, b: Tweety is a bird, f: Tweety flies. Assume FACTICITY and DISJUNCTION. Further assume that p ⊆ b and *(w, p) ∩ *(w, b) = ∅, where *(w, φ) is the set of φ-normal worlds relative to w. Since p ⊆ b, p ∩ b = p. So b = p ∪ (b – p). By DISJUNCTION, *(w, b) ⊆ *(w, p) ∪ *(w, b – p). Because we have *(w, p) ∩ *(w, b) = ∅, we get *(w, b) ⊆ *(w, b – p). But by FACTICITY, *(w, b – p) ⊆ b – p ⊆ ¬p. It follows that *(w, b) ∩ p = ∅ (i.e. b > ¬p is true in w). Hence, no penguins are normal birds; so it can't be inferred that Tweety flies, since Tweety, being a penguin, counts as an abnormal bird. It can be defeasibly inferred, however, that Tweety *doesn't* fly, given *(w, p) ⊆ ¬f, on the assumption that Tweety is a normal penguin (see below). (Note: the symbolization, which should properly contain quantifiers, is simplified here. I postpone the introduction of quantifiers for later; see note 25.)

since the premises already entail $b > \neg p$, provided FACTICITY and DISJUNCTION. Both model and counterexample are therefore irregular with respect to b, but only the counterexample is irregular with respect to p as well. The counterexample is thus gratuitous, and the argument is defeasibly valid.

For the purposes of formulating tolerance principles in a way that makes explicit their deontic content, I shall adopt the system of deontic logic developed in Asher and Bonevac (1996) and Bonevac (1998). Building off Asher and Morreau's system of commonsense entailment, their account represents obligation and weak permission with dual unary deontic operators O, P where[11]

> $O\phi$ is true at world w iff ϕ is true in all w's ideal worlds
>
> $P\phi$ is true at world w iff ϕ is true in some of w's ideal worlds

The logic of O is classical in that true conflicts of obligation cannot arise: Op and $O\neg p$ are contradictory. Asher and Bonevac have proposed analyzing statements of conditional obligation (Bonevac 1998) as well as certain species of *prima facie* obligation (Asher and Bonevac 1996) in terms of deontic generics. These have the form $p > Oq$, with truth conditions

> $p > Oq$ is true at world w iff q holds in all ideals of p-normal worlds relative to w

So analyzed, conditional obligation and prima facie obligation inherit all the nonmonotonic features of >.

V. Duty Detachment. Modus ponens on $p > Oq$ is a special case of default detachment, and is therefore defeasible: if p is true, then q is *normally* obligatory; from p we can infer Oq in the

[11] Weak permission, or the absence of obligation, licenses &E inferences (i.e. $P(\phi\&\psi) \therefore P(\phi)$)—but not ∨E inferences (i.e. $P(\phi\vee\psi) \therefore P(\phi)$). Strong permission licenses ∨E but not &E. See Asher & Bonevac (2005).

absence of intervening moral rules or considerations. Additional contradictory information may require withdrawing that conclusion, as in the following situations.

VI. Moral Exceptions. Cases of moral exception block inferences made on otherwise reasonable moral principles: $p, p > Oq, \neg q \not\models q$.

VII. Deontic Conflict. It follows from the defeasibility of deontic detachment that, in cases of *moral conflict* (or *deontic conflict* more generally) where p, $p > Oq$, r and $r > O\neg q$ all hold, we can infer neither Oq nor $O\neg q$. Equally strong but conflicting moral requirements are susceptible to the same effect of suspended conclusions as in the Nixon Diamond setup. If keeping one promise means breaking another, figuring out which to keep requires substantive moral reasoning; logic itself will not resolve one's moral dilemmas.

VIII. Deontic Specificity. More specific requirements outweigh less specific ones. Like Penguin Principles, rules of exception take precedence over general moral rules. The logic defaults to the more specific of any two deontic requirements: $\{p > Or, (p \& q) > O\neg r, p, q\}$ defeasibly entails $O\neg r$. You should generally keep your promises, but if keeping a certain promise means you will die, you should not keep it.

6.4 Obligation-based sorites

I propose extending the Asher-Bonevac analysis beyond moral obligation, to other kinds of *oughts*—in particular, as a way of capturing the deontic content of tolerance principles for vague predicates. That analysis was designed to provide solutions to various deontic paradoxes. I claim it can do more—namely, explain away the paradoxical nature of tolerance. Viewing the sorites paradox as a species of deontic paradox in this way, I shall argue, serves to advance our understanding of vagueness.

Introducing a generic conditional and deontic operator lets us formulate new expressions of sorites arguments. I shall first consider two. These give different scope readings for the deontic conditional.

Suppose a subject is led through our series of decreasingly bald men, and is asked in sequential order whether each man is bald. For each man #n, given that he and the next man differ only marginally, tolerance requires that calling #n bald commits one to calling #n+1 bald too. This is to reason using inference by *detachment*:

(5)	If one calls #n bald, one should call #n+1 bald	$Fa_n > OFa_{n+1}$
	One calls #n bald	Fa_n
	One should call #n+1 bald	OFa_{n+1}

The deontic conclusion in (5) defeasibly follows from the argument premises, since this is just a special instance of modus ponens, which on our semantics is defeasibly (though not fully) valid. Aggregating enough of these conditional commitments threatens sorites contradiction in the familiar way.

(6) One should call #1 bald
If one calls #1 bald, one should call #2 bald
If one calls #2 bald, one should call #3 bald
Etc.
If one calls #999 bald, one should call #1000 bald
One should call #1000 bald

Although, stated as such, (6) is not valid, not even defeasibly, since what detaches from each conditional is a deontic claim, whereas each antecedent is a non-deontic claim. What is required to chain together these multiple detachment inferences is a way of getting back Fa_{n+1} from each detached OFa_{n+1}, so as to license the next instance of detachment. But in general it is not true that Op implies p. To fix this, we'll add that the subject generally makes each judgment he ought to make, i.e. $OFa_n > Fa_n$ for all n. The argument then relies upon the validity of inference by *deontic chaining*:

(7)	One calls #n bald	Fa_n
	If one calls #n bald, one should call #n+1 bald	$Fa_n > OFa_{n+1}$
	<u>If one should call #n+1 bald, one calls #n+1 bald</u>	<u>$OFa_{n+1} > Fa_{n+1}$</u>
	One calls #n+1 bald	Fa_{n+1}

This is defeasibly valid, since (7) exemplifies the defeasibly valid argument pattern of *generic chaining*: $p, p > q, q > r \Rightarrow r$. Consider any counterexample that makes true $\{p, p > q, q > r, \neg r\}$. There will be a model that makes true $\{p, p > q, q > r, r\}$, which will be more regular than it, because the model contains no irregularities, while the counterexample must be irregular with respect to either p or q.

In this way, multiple deontic conditionals may be chained together in order to defeasibly license detaching the very last consequent. Therefore, supplementing the premises in (6) with $OFa_n > Fa_n$ for each n renders the argument defeasibly valid.

Crucially, arguing in this manner does not appeal to any inference of *deontic transitivity*:

(8)	If one calls #n bald, one should call #n+1 bald	$Fa_n > OFa_{n+1}$
	<u>If one should call #n+1 bald, one calls #n+1 bald</u>	<u>$OFa_{n+1} > Fa_{n+1}$</u>
	If one calls #n bald, one calls #n+1 bald	$Fa_n > Fa_{n+1}$

If (8) were defeasibly valid, then together with deontic chaining, that would license defeasibly concluding $Fa_1 > OFa_{1000}$ on the basis of Fa_1, $Fa_1 > OFa_2$, $OFa_2 > Fa_2, \ldots, Fa_{999} > OFa_{1000}$. Fortunately, inference by transitivity is in general not allowed: $p > q, q > r \mid\!\not\approx p > r$. Any counterexample w must make true $p > q$, $q > r$, and $\neg(p > r)$—that is, all p-normal worlds relative to w make q true and all q-normal worlds relative to w make r true, but not all p-normal worlds make r true. So there is a w′ such that w′ is a p-normal world relative to w, and thus makes true both p (by FACTICITY) and q, but not r, and so is not a q-normal world. This means that at worst there is an irregularity with respect to q in w′ (provided q > r is true in w′). But this is not an irregularity with respect to q in w. So there is no way to show that a world where {p > q, q > r, p > r} are all true is more regular than w. Therefore, transitivity is not defeasibly valid.

We have seen how sorites reasoning can exploit modus ponens on deontic conditionals of the form $Fa_n > OFa_{n+1}$. Alternatively, a wide scope reading of the deontic conditional instead makes *deontic detachment* the key inference:

(9)	It ought to be that: if one calls #n bald, one calls #n+1 bald	$O(Fa_n > Fa_{n+1})$
	One should call #n bald	OFa_n
	One should call #n+1 bald	OFa_{n+1}

Here, there is no need to assume that obligatory judgments are made, i.e. $OFa_n > Fa_n$, since (9) alone would license concluding OFa_{1000} on the basis of OFa_1, $O(Fa_1 > Fa_2)$,…, $O(Fa_{999} > Fa_{1000})$.[12]

Unrestricted license of modus ponens, as in classical logic, would validate both forms of detachment, regular and deontic, in which case the sorites argument (6) is valid on either reading of the deontic conditional. A nonmonotonic system, however, avoids this result. For (5) is merely an instance of *defeasible* modus ponens; detached conclusions must be retracted in light of additional contradictory information. Similarly for (9): oughts distribute across generics only defeasibly. This makes (5) and (9) both defeasibly valid; they must fail somewhere, but where? One might suppose, along with the contextualist, that this occurs relatively late in the series once the subject encounters cases who are clearly *not* bald, and he can no longer competently call them

[12] Following Thomason (1981), we may say that (7), because it makes assumptions about what the subject says, represents reasoning by tolerance within *the context of deliberation*, in which one takes the facts as given and must decide what to do in light of them (i.e. figure out what to say next in light of previous judgments); whereas an argument like (9), given its more "hypothetical" flavor, is made within *the context of judgment*, where nothing is assumed about what is actually said, only about what should be said.

‘bald’. This would indeed yield contradictory information, OFa_k and $O\neg Fa_k$ for some k, given our assumption that O is classical and allows for no moral dilemmas. We saw that this strategy, however, leaves it unexplained why the subject feels uneasy about extending the concept well before that “jumping” point. In contrast, the present account is able to explain why failure of detachment occurs much earlier.

I propose that we understand tolerance principles as providing general but defeasible rules for using vague predicates, thus imposing the following constraint for $0 < n < 1{,}000$:

O-TOLERANCE If one calls #n bald, then one normally should call #n+1 bald

Two different interpretations are made available when uncovering the logical form of O-TOLERANCE: a narrow reading as $Fa_n > OFa_{n+1}$ and a wide reading as $O(Fa_n > Fa_{n+1})$. Either way, O-TOLERANCE delivers a series of conditional obligations, by which one obligation leads to the next. How far these obligations accrue across the entire series depends on how far reasoning by detachment can be sustained; unrestricted, it threatens to carry these obligations to the last man #1,000. Before proceeding though, one other constraint needs mentioning.

A certain tradition, dating back to Waisman (1951), takes the notion of permissibility as central to vagueness and claims that borderline cases exhibit *open-texture*: in at least some situations concerning a borderline F, a speaker is free to call it F and free to call it not-F. In step with such permissibility-based treatments, I shall assume that free permission over the status of borderline cases *generally* holds.

PERMISSION If #n is borderline (not-)bald, then one normally may call #n (not-)bald

Although controversial, this assumption is nonetheless, for our purposes, dialectically neutral, since stronger claims of permissibility are endorsed by both contextualists (Soames 1999:ch.7, Shapiro 2006: §1.2) and their critics (Wright 1987:244, Sainsbury 1990:§9) alike.[13]

The key observation is that O-TOLERANCE and PERMISSION clash over whether to count as 'not-bald' the first #k that is borderline not-bald, given that one has just called #k-1 bald. PERMISSION says it is okay; O-TOLERANCE says to extend one's application of 'bald' from #k-1 to #k. Hence a conflict of requirement: it cannot be that one who must call #k bald is also free to call it otherwise. Because neither principle is stronger, neither wins out, absent further information. This is the same situation as the Nixon Diamond setup, except now with {Fa_{k-1} >

[13] Attempts to define borderlineness have notoriously eluded consensus (see Greenough 2003:§§4-6). The issue is highly disputed. For that reason, I wish not to commit myself to any particular definition. It is worth pointing out, nonetheless, that PERMISSION is compatible with a number of definitions given by contextualists. I mention three.

1. Shapiro (2006 ch.1), following McGee and McLaughlin (1997), defines a borderline case of F as an object x such that the thoughts and linguistic practices of speakers of the language, together with "external" contextual factors (comparison class, paradigm cases, contrast cases and categories), and relevant non-linguistic facts (for 'bald', one's total hair situation—including hair count, head shape, hair length, hair follicle distribution, hair thickness, etc.), neither determine that x is F nor determine that x is not F. The status of individual borderline cases is then left up to particular conversationalists to decide or leave unsettled; the extension and anti-extension of F vary over the course of the conversation as items pertaining to those cases—assumptions, presuppositions, or other propositions implicitly or explicitly agreed to—get added to or taken off the conversational scoreboard. Borderline cases can be settled one way or the other so long as this does not contradict judgments that are kept on record or otherwise violate constraints governing the use of F.

2. Soames (1999:210), in a similar vein, describes the borderline F-range as a realm of discretion, containing individuals about whom the semantic rules governing F issue no verdict, but whom the conversational participants are (provided they agree) free to characterize in either way. See also Dorr 2003:§3.

3. Raffman (1994:53, 1996:178-80)—when she was still a contextualist—claimed that borderline cases of color predicates undergo category shifts under different psychological states. A borderline red-orange patch can look either red or orange—it undergoes a Gestalt-like shift between looking red at one time and looking orange at another—and can be truly judged accordingly each time; these judgments are merely relativized to different internal psychological contexts.

That said, PERMISSION is *not* compatible with many other conceptions of borderlineness—including those on traditional supervaluationism (Fine 1975, Keefe 2007), epistemicism (Williamson 1994), degree theory (Edgington 1992), and incompatibilism (Raffman 2005a). These accounts all deny PERMISSION by treating borderlineness as a status that precludes freedom of verdict. See Shapiro 2006:11 for discussion.

OFa_k, $Indet(a_k) > P\neg Fa_k$, Fa_{k-1}, $Indet(a_k)\}$, where this defeasibly entails neither OFa_k nor $P\neg Fa_k$. So the conclusion in (5), that #n+1 must count as bald, gets withdrawn when #n+1 is considered to be borderline non-bald, even if #n has just been called bald. Similarly for (9). This reflects how strengthening of the antecedent is not generally allowed for deontic conditionals, or generic conditional claims more generally (as shown by deontic and generic specificity).[14]

The proposal carries psychological merit. Confronted with vagueness, one finds himself in a state of not knowing what to say. Wright (2001, 2003) calls this a state of "quandary". Borderlineness brings an onset of uncertainty.[15] This is predicted by the account, on which anyone faced with the conflicting requirements of O-TOLERANCE and PERMISSION is, barring further information, left at an impasse as to what to call #k.

Such is the phenomenology of borderlineness. Exactly where this uncertainty begins for any given subject is hard to say, given higher-order vagueness. Locating the first borderline not-bald #k may prove an unreasonable request, given that where the borderline cases begin is, arguably, itself a vague matter.[16] This too is allowed for on the account: a subject may find it hard to say exactly where he ceases to be obligated in extending his use of 'bald' in accord with tolerance. The reason for this, according to PERMISSION, will be because he finds it hard to say exactly where the borderline non-balds begin.

---

[14] We have $Indet(a_k) > PFa_k$ from PERMISSION; so $Fa_{k-1} > OFa_k$ does not entail $(Fa_{k-1} \& Indet(a_k)) > OFa_k$, not even defeasibly. Nonetheless, defeasible strengthening of the antecedent still holds: $\psi > \chi \mid\approx \phi > \chi$ where $\vdash \phi \rightarrow \psi$.

[15] This is often taken to roughly characterize or indicate something's having borderline status (Sainsbury 1995:64). Epistemicists insist more strongly that this *just is* what borderlineness consists in (Williamson 1994:2, 202).

[16] Defenders of higher-order vagueness include Greenough (2003), Williamson (1994). Deniers include Wright (1987, 1992). Doubters include Tye (1994), Koons (1994).

Note that these are claims about how borderline cases are experienced, not how they are conceptualized. To experience #k as borderline non-bald, the subject need not conceptualize #k as having its borderline status. The proposition *that #k is borderline not-bald* need not enter the contents of one's thoughts. Nor must this be articulated in words. One's experience of something as a borderline case requires only thinking of it *in* a certain way that is responsive to its borderline status, not thinking of it *as* being a borderline case.

Identifying the exact point of one's first encounter with a borderline non-F may be a difficult, if not impossible, even pointless, task. What matters is that this happens before one encounters any clear non-Fs. For the borderline non-bald individuals precede those who are clearly not-bald.[17] The account predicts that reasoning by O-TOLERANCE fails on account of the former rather than the latter.

In this way, the compulsory force of sorites reasoning is seen to fail well before reaching the end of the borderline range. This distinguishes the present account from contextualist treatments that attempt to situate the breakdown of sorites reasoning instead at the end of the borderline region.[18] In contrast, the present interpretation of tolerance predicts that its requirements run out

---

[17] Just how far before is debatable. The first borderline not-bald case may in fact be the first borderline bald case, on the standard assumption that borderlineness is (*strongly*) *symmetric* in the sense that anything borderline F is borderline not-F (and vice versa). That assumption is available to the present account, but inessential. One may drop this assumption in favor of an asymmetric conception of borderlineness, on which the borderline Fs and borderline non-Fs may overlap slightly, but never entirely (or never at all; see Raffman 2005a). On such a non-standard account, the borderline non-Fs will come *slightly before* the first clear non-Fs. Otherwise, on standard accounts, the borderline non-Fs will come *well before* the first clear non-Fs, toward the beginning of the borderline range (the range encompassing all borderline Fs and borderline non-Fs, which, given symmetry, overlap entirely). Either way, sorites reasoning is seen to fail before reaching the end of the borderline range.

[18] More precisely: Call a conditional premise *If #n is bald, #n+1 is bald* in our original sorites argument (2) *confirmable* if a speaker can competently judge both antecedent and consequent to hold. On pain of sorites contradiction, not every premise of (2) is confirmable; there must be a last confirmable premise (or a narrow range of candidates, if *competently judge* is vague). Insofar as each step of sorites reasoning consists in a speaker's accepting *#n+1 is bald* after accepting *#n is bald*, the point of *failure* or *breakdown*

in the borderline region. The norms of O-TOLERANCE and PERMISSION cancel each other out, since both are merely defeasible principles.[19] Contrary to many contextualizers of tolerance, I claim that tolerance does not remain in force throughout the entirety of one's reasoning through a sorites series: once in the borderline range, there really is no requirement to call the next thing F.

---

of sorites reasoning occurs where a speaker can no longer competently do this. Idealizing away from higher-order vagueness, this will be the pair <#k-1, #k> corresponding to the last confirmable premise. Admitting higher-order vagueness, this will be somewhere in the range of pairs corresponding to the range of candidates for being the last confirmable premise—exactly where is not important, since the range of candidates will presumably be relatively narrow. Every contextualist theory in print—ever since Kamp's (1981) original contextualist proposal—seems to predict that this pair (or range of pairs) occurs somewhere toward the end of the borderline range.

1. For Shapiro (2006), #k is the last item whose status is underdetermined by linguistic practice, external contextual factors, and relevant non-linguistic facts. Thus #k+1 is the first definite non-F, i.e. the first item that is determined by these things to be not-F. Conversationalists accordingly cannot judge #k+1 to be anything other than not-F. One is forced to jump at #k+1, calling it not-F, and thereby retract his previous judgment that #k is F (on pain of violating tolerance). This forced change in conversational score marks the limits of sorites reasoning: extended application of F via tolerance cannot go beyond #k.

2. For Soames (1999), #k is the last individual whom conversationalists are free to agree, in line with tolerance, to characterize as F, after which the semantic rules dictate that #k+1 is not F. Discretion ends where rules begin.

3. For Raffman (1994, 1996), #k can appear to be either F or not-F when pairwise compared with #k-1, depending on one's internal psychological state. Given their indiscriminability, #k-1 and #k will both appear to be F while viewed in one psychological state, or both appear to be not-F while viewed in another psychological state. However, the first state is no longer available when comparing #k pairwise with #k+1, if the latter looks not-F, in which case pairwise comparison compels one instead to judge both #k and #k+1 as not-F. No retraction of previous judgments is necessary to retain consistency (contra Shapiro), given that the ensuing judgment that #k is not-F and one's previous judgment that #k is F get relativized to different contexts, now understood as different internal psychological states. This forced shift in psychological state marks a limit of sorites reasoning: one cannot indefinitely remain in the sort of psychological state required for tolerant application of F.

Raffman (1994, 1996) is an outlier among her contextualist peers insofar as she hypothesizes that such a shift occurs for ordinary subjects somewhere in the *middle* of the borderline range, rather than toward the end; her later empirical studies (2005a:ch.5) appear to confirm this point. Yet the general criticism still stands: relativization to a single psychological context appears unable to capture the dual nature of tolerance (see §5, n.23).

4. For Kamp (1981), #k is the last individual one can judge to be F on the basis of tolerance, while remaining in a coherent context of judgment. Attempting to incorporate the information that #k is F into the background information of the context, in order to conclude that #k+1 is F via tolerance, renders the overall context incoherent—because by this point people start to look clearly *not* F. The status of #k+1 as not-F is, presumably, also part of one's contextually salient background information. This blocks one from updating the context by confirming (and hence inferring the consequent of) the conditional *If #k is F, #k+1 is F*, on pain of adding contradictory information to the overall context.

[19] Contrary to incoherentists about tolerance and nihilists about vagueness (see n.7), such conflicts among constraints need not reflect any inherent incoherence in vague predicates.

Contrary to deniers of tolerance, however, the view does not jettison our tolerance intuitions: requirements of tolerance are never violated, only suspended.

The cancelation of O-TOLERANCE with PERMISSION only reflects the suspension of epistemic norms. The felt force of such norms may very well persist, despite their suspension. The subject may continue to experience the pull of tolerance and feel compelled to extend the predicate to #k, despite recognizing its borderline status; hence, the onset of uncertainty in how to proceed. In this way, the felt effects of tolerance do not just suddenly disappear. There is still a *felt* conflict, just no conflict in what, all things considered, what the subject *ought* to say.

This explains away the compulsory appearance of sorites reasoning, but what of its allowance? Granting that the first #k who is borderline not-bald is also borderline bald, PERMISSION says it is defeasibly okay to call #k bald.[20] Importantly, this does *not* conflict with O-TOLERANCE.[21] Although it does not reinstate the requirement to comply with tolerance, it makes doing so permissible: one need not extend his application of 'bald' from #k-1 to #k, though he may if he wishes to. At the first sign of borderlineness, judgments of tolerance are neither mandated nor sanctioned, only permitted. But surely only up to a point: the permissive nature of sorites reasoning cannot continue indefinitely. This threatens yet another variety of sorites argument that must be addressed.

---

[20] For presumably, the borderline Fs *are* (or *just are*, barring intuitionism) the borderline non-Fs (see n.17).

[21] Nor does it, for that matter, conflict with calling #k borderline. One may call #k bald despite acknowledging its borderline status ("He's bald, but only borderline bald"). Such judgments seem perfectly compatible. Nor is it a feature of PERMISSION that we *never* have to call anything borderline—that is, unless borderline status already precludes truth and falsity (although such a suggestion is independently implausible; see Wright 2003).

6.5 Permission-based sorites

A weaker interpretation of tolerance constrains what judgments are *allowed* rather than what is *required*:

P-TOLERANCE If one may call #n bald, then one normally may call #n+1 bald

The deontic conditionals delivered by P-TOLERANCE come in the form of conditional permissions. These threaten sorites contradiction in their own right:

(10) One may call #1 bald
If one may call #1 bald, one may call #2 bald
If one may call #2 bald, one may call #3 bald
Etc.
If one may call #999 bald, one may call #1000 bald
One may call #1000 bald

If one allowance led to another without end, one could count Howard Stern (#1000) as 'bald' without fault—*reductio ad absurdum*. So (10) is not valid. The inference to blame here is *conditional allowance*:

(11)	If one may call #n bald, one may call #n+1 bald	$PFa_n > PFa_{n+1}$
	One may call #n bald	PFa_n
	One may call #n+1 bald	PFa_{n+1}

This is yet another instance of detachment, and hence defeasibly valid. Detaching PFa_{n+1} fails in light of the additional information that $\neg PFa_{n+1}$, as when #n+1 is *definitely not-bald*—that is, too hairy to acceptably count as 'bald'. Beyond the borderline range, any permission to call things bald gets overridden by the stricter constraint CLEAR SANCTION. This prohibits extending 'bald' to any case that clearly does not satisfy it, even if one has just applied it to a previous borderline case without fault. Thus CLEAR SANCTION defeats inferences like (11) otherwise licensed by P-TOLERANCE.

Allowance failure occurs at the end of the borderline region, sometime after encountering the *last* borderline non-bald man. Higher-order vagueness in the transition from the borderline non-balds back to the definite non-balds may prevent us from being able to say exactly where this is for any given subject. Nevertheless, we may reasonably suppose that the first definitely not-bald man #h comes well *after* the first borderline non-bald man #k.[22]

Unlike any borderline case, no uncertainty or feeling of quandary shrouds the status of man #h when considered alone. We know what to call him: he is definitely *not* bald. In contrast, we experience unease when reasoning through the borderline range. This difference in phenomenology is captured by the present account. Within the borderline region, O-TOLERANCE and PERMISSION offer competing, mutually cancelling directives. The resulting lack of directive makes one unsure whether to continue applying the concept; one feels uneasy doing so, even if allowed. Whereas past the borderline region, a clear winner *does* emerge between competing constraints, for CLEAR SANCTION takes precedence over P-TOLERANCE. This explains why one feels compelled to "jump" at <#h-1, #h> and switch from calling #h-1 bald to calling #h not-bald. To be sure, one feels uneasy about extending the concept any further—

---

[22] Given PERMISSION, #h cannot be borderline bald (otherwise it would be okay to call him bald, contrary to stipulation). Given STRONG SYMMETRY (x is borderline F iff x is borderline not-F), #h therefore cannot be borderline not-bald. Given MONOTONICITY, h > r for any borderline not-bald #r (otherwise, since #r is borderline bald given STRONG SYMMETRY and so able to count as bald given PERMISSION, #h would have to count as bald too if h ≤ r). In particular, #h comes after the first borderline not-bald man #k—or well after, assuming there are plenty of borderline non-balds. This will be true no matter exactly where the borderline (non-bald) cases begin or end; so higher-order vagueness is not at issue here. Even if the boundaries of the borderline region are admittedly vague, the point stands: #h appears near the end, not the beginning.

Given a suitable constraint like ADJACENT (x is not bald iff x is hairy), these results project onto a contrast predicate such as 'hairy' for a bald-hairy sorites series: #h cannot be borderline hairy and, moreover, must come well after the first borderline hairy man, notwithstanding higher-order vagueness.

that is, after all, what drives him to switch judgments. One may even feel uneasy about making the jump, given that he just called #h-1 bald. But jump, he must—about *this*, there is no feeling of uncertainty. Nor is there any doubt about #h's non-bald status after making the jump. This contrasts with the lingering unease that persists even after deciding to settle a borderline-bald case as 'bald'. No such lingering unease plagues one's decision to jump, which brings relief in its wake rather than more doubt.

These observations mark several key general differences between forms of permissibility-based sorites reasoning (or *may*-type sorites), as guided by P-TOLERANCE, and forms of obligation-based sorites reasoning (or *must*-type sorites), as guided by O-TOLERANCE.

I. Location of failure. Whereas *must*-sorites reasoning fails near the beginning or middle of the borderline region, *may*-sorites reasoning fails toward the end of the borderline region.

II. Source of failure. Whereas *must*-sorites reasoning breaks down at the first signs of borderline non-Fness, *may*-sorites reasoning breaks down at the first signs of definite non-Fness.

III. Feel of failure. Failures of different types of sorites reasoning are experienced differently—felt uncertainty over continued application of F when *must*-type sorites reasoning fails, as opposed to felt compulsion to discontinue application of F when *may*-type sorites reasoning fails.

IV. Effects of failure. Different psychological accompaniments follow the failures—lingering unease after continued application of F when *must*-type sorites reasoning fails, as opposed to relief after discontinued application of F when *may*-type sorites reasoning fails.

No surprises lurk here. How far one *must* extend a concept and how far one *may* extend it are distinct issues. Normal presentations of the sorites paradox neglect this difference and conflate the two by suppressing the deontic content of tolerance principles. Expressing tolerance solely in terms of conditionals, defeasible or otherwise, leaves this deontic content ambiguous. Formulating tolerance in terms of *deontic* conditionals forces one to resolve any ambiguity over which deontic notion is at issue.

In this way, the contextualist's prediction is partly right—tolerance in continuing to call things F really does stop beyond the borderlines. But what stops that far down is the *allowance* of continued usage, not its *requirement*. Contextualists wrongly identify this as a failure in *obligation-based* tolerance when it is rather *permission-based* tolerance that is at issue. They may be right about where things fail, just wrong about what fails. Indeed, it is hard to see how to recover the distinctions between *must*-type and *may*-type sorites reasoning, given that on every existing contextualist proposal, the operative notion of a context seems to have the *must*-interpretation of tolerance simply built into it.[23]

6.6 Defeasible sorites reasoning

So far, the defeasibility of sorites reasoning has been attributed to the local failure of tolerance inferences: consequences of deontic detachment and distribution are surrendered when around borderline non-Fs; and those of conditional allowance, when around clear non-Fs. This disarms sorites arguments (6) and (10) of the "listed conditionals" sort, whose form mirrors how one is actually led through sorites-style reasoning step by step. Notoriously, not all presentations of the

---

[23] Standard contextualism has it that: one who judges a_n is F in one context must (if asked whether the marginally different a_{n+1} is F) judge a_{n+1} is F if he is to remain in that same context, or switch contexts if he judges otherwise. Introducing deontic operators and relativizing the deontic facts to individual contexts, as a fallback strategy, holds little promise of recovering the *may*-interpretation, given that there is only *one* notion of context to relativize to.

sorites are so upfront. Variants of (6) and (10) compact their listed conditionals instead into a single inductive premise.[24]

(12)	One calls #1 bald	Fa_1
	For n < 1000: if one calls #n bald, one should call #n+1 bald	$\forall n_{<1000}\,(Fa_n > OFa_{n+1})$
	For n < 1000: if one should call #n bald, one calls #n+1 bald	$\forall n_{<1000}\,(OFa_n > Fa_n)$
	One should call #1000 bald	OFa_{1000}
(13)	One should call #1 bald	OFa_1
	For n < 1000: if one calls #n bald, one should call #n+1 bald	$\forall n_{<1000}\,O(Fa_n > Fa_{n+1})$
	One should call #1000 bald	OFa_{1000}
(14)	One may call #1 bald	PFa_1
	For n < 1000: if one may call #n bald, one may call #n+1 bald	$\forall n_{<1000}\,(PFa_n > PFa_{n+1})$
	One may call #1000 bald	PFa_{1000}

Each inductive premise in (12)-(14) is a universally quantified generic conditional. To evaluate such statements, we define $\forall$ in the usual way by adding assignments to points of evaluation for truth:[25]

> $\forall x\Phi$ is true at world w under assignment α iff Φ is true at world w under any assignment differing from α at most in its assignment to x.

Consider the inductive premise in (12): $\forall n_{<1000}(Fa_n > OFa_{n+1})$ is true at $\langle w, \alpha\rangle$ iff $Fa_n > OFa_{n+1}$ is true at $\langle w, \alpha'\rangle$ for any α' differing from α in what it assigns to n—i.e. OFa_{n+1} holds in all Fa_n-

[24] A further variant combines *may* with the wide-scope reading of the deontic conditional:

(12*)	One may call #1 bald	PFa_1
	For n < 1000: if one calls #n bald, one should call #n+1 bald	$\forall n_{<1000}\,O(Fa_n > Fa_{n+1})$
	One may call #1000 bald	PFa_{1000}

where permissions accrue over a series of wide-scoped conditional obligations via the defeasible inference of *forced allowance*: $O(Fa_n > Fa_{n+1}), PFa_n \mid\approx PFa_{n+1}$. (12*) is handled in much the same way as the others.

[25] For more details on how quantification iteracts with deontic conditionals, see Asher & Bonevac (1996:§1.2).

normal worlds, for any value of n < 1000. But this is satisfied by assumption. So it makes no truth-conditional difference if we use a single inductive premise or list out each individual tolerance conditional. In that case, argument (12), like its original “step-by-step” formulation, is invalid but defeasibly valid. Likewise for (13)-(14).

A unified picture emerges on which logic, argument and reasoning all neatly converge in the sorites. The inductive premises of (12)-(14) are justified by their respective tolerance principle, O-TOLERANCE (narrow or wide) or P-TOLERANCE. The rules of inference underlying these arguments are chaining, deontic detachment, and conditional obligation, respectively. As before, these are all defeasible.

On this construal, sorites reasoning is not to be seen as inherently defective or never worth utilizing, but as unproblematic in many instances, so long as it is not carried out too far down a sorites series.[26] Such reasoning no doubt incurs the risk of overextending a vague predicate beyond its proper or permitted bounds, though its built-in defeasibility is meant to countervail this effect. Tolerance crucially admits of exceptions.[27] To count (12)-(14) as invalid marks not so much an outright rejection of sorites reasoning as an acknowledgement of its limits. Their

---

[26] Sorites reasoning may hold little utility in considering clear cases, whose status is settled by independent means other than comparison with similar or adjacent cases. However, it may prove potentially useful particularly when sorting through and deciding the status of borderline cases.

[27] To think this claim of defeasibility amounts to stating a sort of triviality, like “sorites reasoning fails when it fails” or “tolerance holds, except when it doesn’t”, is clearly mistaken. The account offers an explanation for the failure of sorites reasoning in terms of a conflict of norms—a diagnosis that is far from trivial or inconsequential. For I claim that the normative consequences of tolerance and borderlineness are *not* compatible. That runs contrary to the standard line of argument (e.g. in Greenough 2003) that these are perfectly compatible features of vague predicates.

defeasible validity reflects failure of sorites reasoning on a global level; locally, it remains unproblematic.[28]

Consider two definitely bald individuals, one slightly balder than the other by just a single hair. Tolerance of either sort, *must* and *may*, would justify calling the man with the extra hair 'bald' *on the basis of* his being just slightly less bald than the other, even though he himself is already definitely bald. Such conclusions are harmless.[29] Reasoning by tolerance merely gives another—albeit superfluous—reason to call him 'bald', on top of looks alone (without comparison to anything else), since CLEAR CALL already requires calling $Baldy_3$ bald; hence, the overdetermination of obligation.[30]

Alternatively, consider extending 'bald' to borderline cases on the basis of comparison to people who are definitely bald. This too is harmless. Although *must*-type tolerance gets suspended in light of PERMISSION, *may*-type tolerance remains operative, offering another extra

[28] The defeasibility of *must/may* sorites reasoning may, in turn, derive from more general generic epistemic principles, such as "You normally shouldn't (can) treat or call things differentially unless (if) they have some relevant detectable difference." Should these general principles be defeasible for reasons independent of contextual variability, this would reinforce the thought that vagueness is not essentially tied to contextual variation.

[29] One may balk. But the burden of proof is on my opponent to show: why such a case of sorites reasoning should somehow be bad, why its conclusion should somehow be unacceptable or infelicitous, and why the reasoner should somehow be unreasonable or at fault. To say that a reasoner using such methods could arrive at contradictory conclusions is a non-starter. The risk of bad consequences is no guarantee of actual badness. At any rate, this is all question-begging from a nonmonotonic perspective, the whole point of which is to dispense with the idea that things, once true (or valid), are guaranteed to remain true (or valid) as more information gets added. What is to keep good reasoning from being, like truth and validity, nonmonotonic in nature? We should expect no guarantee that methods of reasoning, though good when applied to a few cases, would remain so, as more cases are considered.

[30] Note: this is a claim about overdetermination *of* obligation (a single obligation arising from multiple conditions), and not overdetermination *in* obligation (multiple obligations arising from a single condition). More precisely expressed in the logic: $p > Oq, r > Oq \mid\!- (p \vee r) > Oq$. This is guaranteed by DISJUNCTION.

(albeit defeasible) reason to count borderline individuals as 'bald'. Only now, it is the allowance of continued use, rather than its requirement, that gets overdetermined, given PERMISSION.[31]

The defeasibility perspective sheds light on two problem cases that have puzzled philosophers of vagueness. One is Horgan's (1994) "forced march" sorites, in which a subject led through a sorites series is forced to make a judgment about each #n—F or not-F. Contextualists often employ this as a litmus test for a given theory's ability to uphold tolerance principles. Contextualism, they argue, rightly predicts that subjects of forced march trials are required to continue calling things F via tolerance until one can no longer do so, at which point tolerance applies in the reverse direction and they are required to start calling things not-F; whatever direction, tolerance remains in force throughout.

The other is Sainsbury's (1990) paint shop owner, who must pick an arbitrary division point along an arrangement of red and orange paint jars in order to separate the "red" and "orange" shelves. Critics of contextualism argue that the permissibility of drawing such a cutoff, where the last jar on the "red" shelf is only marginally different in color from the first jar on the "orange" shelf, shows that the requirements of tolerance need not always be observed. Alternatively, consider a perhaps more familiar scenario: the task of assigning end-of-term letter grades to students typically requires, and therefore permits, drawing arbitrary cutoffs along a (possibly curved) numerical grade scale.

The defeasibility approach easily accommodates both tolerance-enforcing and cutoff-permitting intuitions simultaneously. How strict the requirements of tolerance are may vary according to circumstance. In forced march cases, the demands of tolerance are presumably reinforced over other defeasible principles, such that one should not stop calling things F unless

[31] In other words: $p > Pq, r > Pq \mid\!- (p \vee r) > Pq$—again, guaranteed by DISJUNCTION.

absolutely necessary. In that case, O-TOLERANCE would supersede, rather than mutually cancel out with, PERMISSION—until past the borderlines, when it is itself overridden by the stronger rule of CLEAR SANCTION. Whereas in paint shop cases, it is the requirement to obey tolerance at every point that gets overridden by the need to violate it at some point, for purposes of organizing one's shelves by color or sorting one's students by grade. Such fairly commonplace situations, in which lifting regular or otherwise enforced norms of tolerance allows arbitrary boundaries to be drawn without fault, demonstrate the defeasibility of tolerance.

6.7 Metaphysics vs. epistemology

What remains of the classical sorites? Traditional presentations of the paradox make use of neither ordinary language generics nor deontic modals.

(15) #1 is bald
#1 is bald ⊃ #2 is bald
#2 is bald ⊃ #3 is bald
Etc.
<u>#999 is bald ⊃ #1000 is bald</u>
#1000 is bald

A story about the defeasible nature of the ordinary English generic, however persuasive, does not erase the fact that the material conditional of classical logic *is* modus ponens-supporting, and unequivocally so. One may complain this leaves the paradoxical nature of (15) unaddressed. Indeed, other formulations of the sorites forgo any use of conditionals whatsoever, classical or otherwise.

(16) #1 is bald

Not: #1 is bald and #2 is not bald

Not: #2 is bald and #3 is not bald

Etc.

Not: #999 is bald and #1000 is not bald

#1000 is bald

(17) #1 is bald

Either #1 is not bald or #2 is bald

Either #2 is not bald or #3 is bald

Etc.

Either #999 is not bald or #1000 is bald

#1000 is bald

(Quantified versions of these use a single inductive premise rather than a string of premises.[32] Similar remarks will apply below.) Concerning these, the account, admittedly, remains silent. It fails to speak to the proper evaluation of the original sorites argument, expressed either in the material mode, as in (15), or some classical equivalent, as in (16) or (17). For classical logic is stripped of all epistemic modals and other normative notions that abound in ordinary idioms. Yet this is a problem (if at all) facing *every* account that produces a revisionist or otherwise non-classical reading for tolerance conditionals. Contextualists are no exception: they too must answer to the original sorites stated in classical terms, which, it would appear, contains no contextually relativized premises and admits of no contextual variation in the evaluation of its premises.[33]

---

[32] These employ a principle of induction. On why *not* to question this principle, see Kamp (1981:§§1–2).

[33] The contextualist may insist that the very act of evaluating the argument inevitably produces a change in context. But this is to sidestep the question of what truth-values sorites premises possess *independently* of any attempt to evaluate them—an odd question, perhaps, of genuine concern only to classical logicians, but

Thus it is no *special* problem for the defeasibility view that it fails to say anything new about these classical sorites arguments. Nor does it need to—I might add. For it is obvious what should be said concerning these arguments, independently of any analysis for the ordinary English *if*. They are valid but unsound: each contains a false premise.[34] Thus, for some n, #n is bald and #n+1 is not bald. The existence of cutoffs for vague predicates is a well-known consequence of classical logic and semantics. (An infamous verdict—but what else *can* be said that isn't revisionary?) I take this to be of substantive metaphysical consequence, if anything: there is a *fact* of the matter about who in the series is the least bald bald man. Whether to take this verdict seriously is an open question that faces every theory.[35]

Tolerance, by contrast, is an epistemic phenomenon. Tolerance principles encode norms of judgment: how 'bald' may or must apply to an individual under consideration in light of other information. They do not determine or otherwise settle metaphysical issues concerning the facts about baldness.[36] Nor do they verify all the premises of classical sorites arguments like (15)–(17). To otherwise deny the existence of cutoffs on account of tolerance is to conflate metaphysics with epistemology. Whether a single hair can make the difference between being

---

reasonable nonetheless. If the contextualist can make no sense of it, neither should this be expected of the defeasibilist.

[34] The first premise "#1 is bald" is indisputably true (disregarding nihilism). Hence the false premise will be some instance of the relevant inductive premise: $\forall n(Fa_n \supset Fa_{n+1})$, $\forall n\neg(Fa_n \& \neg Fa_{n+1})$, or $\forall n(\neg Fa_n \vee Fa_{n+1})$. In the universally quantified versions, the inductive premise itself is therefore falsified.

[35] It should be noted that contextualists themselves are divided over the issue of bivalence. Shapiro (2006) denies bivalence, claiming that indeterminacy in truth-value obtains even in individual conversational contexts. Graff-Fara (2000) takes each context of judgment to come pre-equipped with a fully bivalent truth-value distribution. Raffman (2005b) claims her earlier contextualist work (1994, 1996) remained neutral on the question.

[36] The truth-functional premises of (15)–(17) say nothing about our judgments, internal psychological states, shared conversational score, or other facts about our epistemic situation, broadly speaking; as such, any distribution of truth-values is strictly determined independently of tolerance.

bald and not being bald is one issue (of metaphysics). Whether we can so judge or ascertain in certain circumstances is another issue (of epistemology). Teasing the two apart—fact from judgment—is essential in keeping metaphysical matters distinct from epistemological concerns. Contextualists disregard this distinction when they insist on reading the truth-values of sorites premises directly off our tolerance-conforming judgments. Our inability to locate a fixed cutoff is no demonstration of its nonexistence.[37] Nor is the act of refusing to extend a vague predicate past a certain point, as allowed for by the defeasibility of tolerance, any evidence that an actual cutoff exists *there*.[38]

I have argued that we should construe tolerance principles as issuing defeasible norms of judgment. Insofar as sorites reasoning rests entirely upon tolerance, we should therefore construe sorites reasoning as a form of defeasible reasoning. Insofar as sorites arguments are supposed to capture the way one would *actually* reason in a sorites series, we should therefore formulate these in a way that lays bare the defeasible, exception-admitting nature of such reasoning, as in (12)-(14) or similar variants.

As such, classical sorites arguments like (15)-(17) do not constitute a form of sorites argument proper. They are (I think) perfectly coherent arguments, no doubt, and available for anyone to discuss or advance. What they serve to show, by way of reductio, is that the classical inductive premise is false and therefore that cutoffs exist. But whether there *is* a cutoff in a sorites series (an issue of existence) is an entirely different matter from whether one may *draw* a

---

[37] To echo (in spirit) a familiar lesson from the post-Kripkean literature on modality: just as necessary truths may elude a priori knowability, determinate truths about cutoffs may elude the cognitive faculties of competent speakers.

[38] Nor has any contextualist ever cared to suggest otherwise—providing yet further confirmation that tolerance has no bearing on the metaphysics of vagueness.

cutoff when proceeding along a sorites series (an issue of judgment). Insofar as sorites arguments owe their plausibility to tolerance, any sorites inductive premise that ordinary subjects assent to must be a generic—something, as I have argued, that is properly expressed in terms of ordinary *ifs* and *oughts* (or *shoulds*, *musts*, *cans*, *mays*, or other cognates). Once rightly recognized to be a generic, this premise gives *zero* support for the corresponding universal generalizations of the classical material premises in (15)-(17).[39]

One can be revisionist about the classical conditional, as really having generic truth conditions. Or one can keep the material conditional as it is, and maintain that, sorites arguments, properly formulated, do not appear in material mode at all. The difference lies in what one wishes to adopt a revisionist attitude toward: whether the classical connectives themselves or the formulation of arguments initially thought to employ those connectives.[40] What then to revise, formulation or formulae? This is a difficult issue, one I don't intend to take up here. The present account leaves both options open.

The metaphysics vs. epistemology distinction will apply at any higher order of vagueness (if such there be). Whether a single hair can mark the boundary between being *definitely* bald and *borderline* bald is a metaphysical question.[41] Whether one should or can continue to call the next individual definitely bald is a separate, epistemological question. The same goes for other higher-order distinctions—between being borderline bald and definitely not-bald, or definitely definitely bald and borderline definitely bald, or definitely borderline bald and borderline borderline bald, and so on, on up. At each order, some norm of tolerance will mandate or license our extended

[39] Although, as Sorensen (2012) suggests, ordinary speakers may easily conflate the two.

[40] Or one may question what logic has to do with (sorites) reasoning *at all*. For such skepticism, see Harman (1984).

[41] One presupposing there is such a thing as *being definitely bald* and *being borderline bald* to begin with.

use of the vague predicate **F** in question ('definitely bald', 'borderline definitely bald', 'borderline borderline bald', etc.), but only defeasibly. The point of defeat will depend on whether O-TOLERANCE or P-TOLERANCE is in play: one need not continue applying the predicate when running up against cases that are borderline **F** (respectively: borderline definitely bald, definitely borderline bald, borderline borderline borderline bald, etc.), but must discontinue its use when finally encountering cases that are definitely not-**F** (respectively: definitely borderline bald, definitely definitely not-definitely-bald, definitely definitely definitely not-bald, etc.). The phenomenological differences between these failures will be as before at the first-order level.

6.8 Defeasibility vs. incoherence

This paper has argued for an alternative theory that rivals contextualism in its ability to explain the phenomenology of vagueness. A salient example for comparison is Hans Kamp's (1981) pioneering contextualist treatment of vagueness in "The Paradox of the Heap".

Accommodating our tolerance intuitions is what drove Kamp to abandon his earlier supervaluationist theory (Kamp 1975) in favor of a contextualist treatment of vagueness (Kamp 1981). On the later view, every instance of the inductive principle

(1) For each n < 1,000: if #n is bald, then #n+1 is bald

is true in any context. This is achieved by adopting a dynamic semantics for the conditional that is not truth-functional: a conditional is true in a context c just in case either its antecedent is false in c, or its consequent would be true in the updated context that results from incorporating the information carried by the antecedent of the conditional into the background information of c.

Given a suitable tolerance principle—wherever a and a′ are similar and Fa is part of the background information of c, Fa′ is true in c—this verifies each instance of the inductive premise. However, the inductive principle is itself false. Kamp offers a revised semantics for the universal quantifier that involves the notion of an *incoherent context*: a universal generalization is true in a coherent context c just in case each of its instances is true and acceptance of the universally quantified sentence preserves coherence of context. It follows that the inductive premise (1) is false, since attempting to incorporate all the information carried by the universal generalization into the background of one's current context invariably creates an *in*coherent context.

Kamp's view has a number of puzzling features. One, it allows a false universal generalization like (1) to have all true instances.

Two, it has no way of defusing the sort of "step-by-step" sorites argument that does not invoke a single inductive premise but rather lists out its individual instances. Once again, consider:

(2) #1 is bald
If #1 is bald, #2 is bald
If #2 is bald, #3 is bald
Etc.
<u>If #999 is bald, #1,000 is bald</u>
#1,000 is bald

Although each premise is true when evaluated individually, there is no coherent context in which all the premises are true; so the argument turns out trivially valid on any suitable dynamic conception of validity, now perhaps most naturally understood as preservation of truth across *coherent* contexts. Yet one cannot reason through the argument and on that basis legitimately claim its conclusion: we all recognize that sorites reasoning is *bad* reasoning, even when carried out step by step, however compelling each step may be. This leaves us with a valid argument whose premises are each individually true, which nonetheless embodies bad reasoning. Good reasoning requires preserving coherence with each update of context, but by assumption the information carried by the premises cannot be incorporated into a single context all at once; all tolerance guarantees is their individual truth, not their collective truth. This is an odd consequence of the proposed semantics: it divorces good reasoning from truth and validity—contrary to the spirit of dynamic approaches that fashion the logic after how we actually reason.

Three, the view seems to predict failure of sorites reasoning relatively late in the series toward the end of the borderline range—presumably, because this is where the subject begins to make false judgments, should he extend his use of 'bald' via tolerance and begin applying it to those who are clearly not bald; attempting to add such contradictory information to the overall context of judgment renders it incoherent. And yet sorites reasoning appears suspect well before that point. Kamp's contextualist account, like its numerous successors, fails to adequately explain why.

The defeasibility view, in contrast, avoids these pitfalls. In allowing both a sorites inductive premise and all its premises to be true, it respects our tolerance intuitions without having to posit, in radically revisionist fashion, false universal generalizations with all true instances. In diagnosing all sorites arguments as invalidly drawing conclusions from true premises, it recognizes both the appeal and the limits of sorites reasoning without having to divorce reasoning from logic—in particular, we need not accept there are arguments which, despite being valid and having all true premises, nevertheless embody bad reasoning. In distinguishing different types of sorites reasoning, it gets the phenomenology right without having to introduce all the extra technical machinery related to an *incoherent context*—now, properly understood as introducing defeating information incompatible with one's conclusion.[42]

Nevertheless, it preserves the original insight from Kamp's work, largely ignored by later contextualists, that continued reasoning via tolerance really does *fail* beyond a certain point—and subsequently that failure of sorites reasoning is *not*, as later contextualists have maintained, simply a matter of switching one's judgments and reasoning in accordance with tolerance in the reverse direction. The view is thus closer in spirit to Kamp's contextualism than that of Raffman, Soames, Fara, or Shapiro.

6.9 Conditionals and consequence

I have argued that we can exploit the defeasibility of deontic conditionals to defuse the sorites. To achieve this requires some combination of a nonmonotonic deontic logic: nonmonotonic, for modelling the defeasibility of sorites reasoning; deontic for capturing the implicit deontic content of tolerance principles. Both elements are crucial: without defeasiblity, tolerance runs amok;

[42] One subtler difference is worth noting. On Kamp's account, incoherence of context results from attempting to update the context with information carried by the *antecedent* of a sorites conditional. Suppose the "offending" conditional is $Fa_k \rightarrow Fa_{k+1}$. This predicts that one who has updated the context with $\{Fa_1, Fa_2, \ldots, Fa_{k-1}\}$ is still allowed to conclude Fa_k on the basis of $Fa_{k-1} \rightarrow Fa_k$ *without incoherence*, since the truth of $Fa_{k-1} \rightarrow Fa_k$ guarantees the truth of Fa_k in the updated context. The context becomes incoherent only when one attempts to process the offending conditional $Fa_k \rightarrow Fa_{k+1}$. It follows that Fa_k, though unproblematically inferred from other information, cannot unproblematically serve as grounds for further conditional reasoning—which is odd. More generally, the overall view seems to fault over-accumulation of information as the reason why sorites reasoning fails.

In contrast, on my account, incoherence of context—now understood as conflict of requirement—arises from introducing information that is incompatible with the *consequent* of the "offending" sorites conditional. This forces one to withdraw the conclusion Fa_k, which he would otherwise infer on the basis of that conditional. This removes the problem of having information that is *both at once* unproblematic when inferred on the basis of previous conditionals *and* problematic when used to infer further information. Failure of sorites reasoning is also no longer due to over-accumulation of information carried by previous conditionals, but rather a matter of recognizing certain features of the case at hand #k that suspend or overrule conclusions otherwise prescribed by tolerance.

without deontic expressions, there is no way to distinguish *must*-type and *may*-type sorites reasoning.[43]

The choice of analysis for deontic conditionals matters. To see this, we might compare the current Asher-Bonevac-Morreau (ABM) analysis of deontic conditionals with the alternative Kratzer-style accounts defended in recent work by Kolodny & MacFarlane (2010) and Willer (2012). These bear certain similarities to the ABM theory: Kolodny and MacFarlane's account, like ABM's, surrenders modus ponens, while Willer's version shares the feature of nonmonoticity. Although their analyses were not designed with vagueness in mind—rather, their concern is solving the Miners paradox—it may still be worth illustrating the differences, especially if the sorites is to be thought of as simply another species of deontic paradox.

On their semantics, deontic modals function like epistemic modals and so are specifications of informational modal operators: they quantify over a set of possible worlds that is determined by a separate informational parameter. Sentences with deontic modals get evaluated relative to two parameters, a possible world *w* and an information state *i* (modeled as a set of worlds), with truth conditions:

---

[43] Any attempt to reproduce the differences between *must*-type and *may*-type sorites within the conditional formulation itself is unlikely to succeed. One could, for instance, try distinguishing two *subjunctive* versions of tolerance conditionals: "If #n were to count as F, so would #n+1" (*would*-tolerance) vs. "If #n were to count as F, so could #n+1" (*could*-tolerance), where these have their standard Lewis-Stalnaker truth-conditions: $Fa_n \Box\!\!\rightarrow Fa_{n+1}$ (or $Fa_n \Diamond\!\!\rightarrow Fa_{n+1}$) is true at w iff all (or some) of the closest Fa_n-worlds relative to w are Fa_{n+1}-worlds. Presumably, *would*-type sorites reasoning will fail at the *beginning* of the borderline range, where a_{n+1} is now borderline F and thus no longer F in all the closest worlds where its predecessor a_n is F, whereas *could*-type sorites reasoning will fail at the *end* of the borderline range, where a_{n+1} is now definitely not-F and thus no longer F in even some of the closest worlds where the previous a_n is F. This distinction is of course only available on the assumption of weak (rather than strong) centering; otherwise $\Diamond\!\!\rightarrow$ collapses back into $\Box\!\!\rightarrow$. The obvious problem with such subjunctive tolerance conditionals of either form is that any centered conditional like $\Box\!\!\rightarrow$ or $\Diamond\!\!\rightarrow$ will support modus ponens without restriction. On pain of validating sorites arguments, one of these (and hence their universal generalization) must be false—which is hardly a vindication of tolerance. The decentered generic > escapes this snare by only licensing defeasible modus ponens.

$[\![\Box_d \phi]\!]$ is true at <w,i> iff for all $w' \in d(i)$, $[\![\phi]\!]$ is true at <w′,i>

$[\![\Diamond_d \phi]\!]$ is true at <w,i> iff for some $w' \in d(i)$, $[\![\phi]\!]$ is true at <w′,i>

Deontic *ought* is a necessity operator whose domain (i.e. the ideal worlds) is selected by a deontic selection function *d*. Deontic *may* is its dual possibility operator. Conditionals take their familiar Kratzer interpretation on which they function as restrictors on informational modals, or in the case of deontic conditionals, deontic modals. The simplest suggestion is that a conditional antecedent $[\![\textit{if}\ \phi]\!]$ contracts the information state by ruling out worlds in which the antecedent is false. Thus,

$[\![\textit{if}\ \phi, \psi]\!]$ is true at <w,i> iff $[\![\psi]\!]$ is true at <w,i+ϕ>

where the result of strengthening i with ϕ, i+ϕ, is defined by $i+\phi = i \cap \{w: [\![\phi]\!]$ is true at $\langle w,i\rangle\}$.

Where the consequent is a deontic modal such as *ought*, we get the truth conditions for a deontic conditional.

$[\![\textit{if}\ \phi, \Box_d\psi]\!]$ is true at <w,i> iff for all $w' \in d(i+\phi)$, $[\![\psi]\!]$ is true at <w′,i+ϕ>

Whether modus ponens is preserved depends on the notion of logical consequence at play. For Kolodny and MacFarlane, this is the classical notion of truth-preservation at any point of evaluation (now <w,i>):

CLASSICAL VALIDITY $\phi_1,\ldots,\phi_n \models \psi$ iff for all w and i such that $w \in i$: `if $\langle w,i\rangle \in [\![\phi_1]\!]$ and … and $\langle w,i\rangle \in [\![\phi_n]\!]$, then $\langle w,i\rangle \in [\![\psi]\!]$

This spells failure of modus ponens. For both $\ulcorner \textit{if}\ \phi, \psi \urcorner$ and ϕ might be true at some index <w,i>, in which case we know that ψ is true at any ϕ-shifted point of evaluation <w,i+ϕ>, but this is compatible with ψ being false at the original index <w,i>.

As far as the sorites is concerned, this would guarantee the defeasibility of reasoning with tolerance conditionals on their deontic conditional reading. Modus ponens on tolerance conditionals will fail for some n where both Fa_n and ⌜if Fa_n, $\Box_d Fa_{n+1}$⌝ are true at <w,i> but $\Box_d Fa_{n+1}$ is false at <w,i>. Supposing i represents one's knowledge of what to call items in a sorites series in light of what has been said so far, this corresponds to a situation where one has just called #n bald and it is true that if one calls #n bald then one ought to call #n+1 bald—such that in all the updated deontic ideal worlds (relative not to i but to the Fa_n-shifted information state of knowing that one has just called #n bald), upon knowing that one has just called #n bald, one *does* go on to call #n+1 bald—yet $\Box_d Fa_{n+1}$ remains false in one's present unshifted information state: in some ideal world (relative to one's current information state of not knowing that one has just called #n bald) one fails to call the next #n+1 bald. This will be, let us assume, due to familiar reasons: the ideal worlds selected by d(i) will presumably include some where #n+1 is recognized as being a borderline case and therefore (permissibly) counted as not-bald. Why does this not also falsify the conditional ⌜if Fa_n, $\Box_d Fa_{n+1}$⌝ at <w,i>? Here is where tolerance comes into the picture. On an information-theoretic semantics, the tolerance principle is best conceived of as a constraint on our update procedures. The obvious suggestion is:

O-TOLERANCE* For all i and for all a_n, a_{n+1} that are sufficiently similar with respect to F, $\langle w, i+Fa_n \rangle \in [\![\Box_d Fa_{n+1}]\!]$

The idea is that updating with Fa_n in effect commits one to calling the next thing F, by throwing out all the $\neg\Box_d Fa_{n+1}$-worlds from one's information state (thereby eliminating the possibility that #n+1 could be a borderline case and hence permissibly counted to be not-F, i.e. $\Diamond_d \neg Fa_{n+1}$). Thus there can be no world in $d(i+Fa_n)$ where $\neg Fa_{n+1}$, which guarantees $\langle w,i \rangle \in [\![\text{if } Fa_n, \Box_d Fa_{n+1}]\!]$.

The problem with this proposal is that it appears unable to give a corresponding account of permission-based tolerance. Consider the analogue principle for p-tolerance:

P-TOLERANCE* For all i and for all a_n, a_{n+1} that are sufficiently similar with respect to F, $\langle w, i+\Diamond_d Fa_n \rangle \in [\![\Diamond_d Fa_{n+1}]\!]$

This would predict that at the end of the borderline range, supposing one calls the last borderline case #h bald and updates accordingly with $\Diamond_d Fa_h$, one's updated information state will include worlds in which it is permissible to call #h+1 bald, i.e. where $\Diamond_d Fa_{h+1}$ is true (relative to $i+\Diamond_d Fa_h$)—contrary to the fact that, since #h+1 is no longer a borderline case but is definitely not-bald, calling that individual bald is *not* allowed, regardless of one's previous judgments.[44] In this way, Kolodny and MacFarlane's analysis appears only able to accommodate one, not both forms of tolerance.

Willer's version of the analysis proposes an alternative, dynamic conception of logical consequence, one that keeps track of changes each added argument premise may induce on the information parameter:

DYNAMIC VALIDITY $\phi_1, \ldots, \phi_n \models \psi$ iff for all w and i such that $w \in i$:
if $\langle w, i \rangle \in [\![\phi_1]\!]$ and … and $\langle w, i+\ldots+\phi_{n-1} \rangle \in [\![\phi_n]\!]$,
then $\langle w, (i+\ldots+\phi_{n-1})+\phi_n \rangle \in [\![\psi]\!]$

On this dynamic conception of validity, monotonicity is no longer guaranteed. Since adding premises now affects the information parameter in light of which the argument conclusion is evaluated, this leaves open the possibility that the truth-value of the conclusion at a possible

[44] Doing so would violate the constraint of CLEAR SANCTION, now expressed as: $d(i+Det(\phi)) \subseteq [\![\phi]\!]$ for any i. Thus $\langle w, i+\Diamond_d Fa_h \rangle \in [\![\Diamond_d Fa_{h+1}]\!]$, assuming one knows #h+1 is definitely not-F in the first place (i.e. that $i \subseteq [\![Det(a_{h+1})]\!]$).

world may change from true to false.[45] The situation of the sorites may be taken to illustrate this possibility. Suppose one has just called #n bald. Before ruling out the possibility that #n+1 is a borderline case, one needn't call #n+1 bald. In that case, $\langle w,i \rangle \in [\![\neg \Box_d Fa_{n+1}]\!]$. But since calling #n+1 bald is required after eliminating that possibility from one's information state upon updating with Fa_n, we have $\langle w,i+Fa_n \rangle \notin [\![\neg \Box_d Fa_{n+1}]\!]$. So we see that $\neg \Box_d Fa_{n+1} \models \neg \Box_d Fa_{n+1}$ but $\neg \Box_d Fa_{n+1}, Fa_n \not\models \neg \Box_d Fa_{n+1}$. Hence the failure of monotonicity.

This demonstrates how on Willer's account, updating can make previously available information no longer available. This allows arguments to be defeasible insofar as old conclusions drawn from certain premises need not stay established as new premises are added. Unfortunately, this feature is not enough to block certain undesirable *new* conclusions from being drawn. Indeed, it is easily verified that the dynamic construal of logical consequence, together with the semantics for conditionals, validates modus ponens—and with it, any sorites argument like (6) applying modus ponens on tolerance conditionals.[46] By contrast, ABM's theory licenses such inferences only defeasibly. This type of defeasibility is crucial for defusing the sorites, since it is newly derived information (e.g. concluding that borderline #k must be called bald or that definitely not-bald #h can be called bald) that produces soritical results, not the recycling of old information (e.g. calling #k-1 or #h-1 bald), which is harmless. This suggests that it is defeasibility in the stronger sense of *inferential defeat*, not in the weaker sense of *information*

---

[45] That is, unless we impose a constraint of Persistence (for any $i' \subseteq i$: if $\langle w,i \rangle \in [\![\phi]\!]$ then $\langle w,i' \rangle \in [\![\phi]\!]$), which would ensure preservation of truth at a world under information strengthening.

[46] See Willer (2012:458) for proof. I don't, of course, dispute the merits of salvaging modus ponens, as far as the miners is concerned. All this shows is that Willer's style of analysis is not quite suited for addressing the sorites.

loss, that is required for delivering the sort of defeasibility in conditional reasoning needed to block the sorites.

6.10 Permissive consequence

The strategy of logical revisionism as a contextualist response to the sorites was pioneered by Kamp (1981). An information-theoretic semantics is just one alternative way to implement this general strategy, by modifying either the semantics for the logical connectives or the very notion of logical consequence. Contextualists who wish to keep modus ponens may instead weaken the notion of logical consequence in other ways, such as denying the transitivity of validity. This is the strategy taken up by Cobreros et al. (2010, 2012).

They define a weakened notion of *permissive consequence* off two notions of non-classical truth: strict and tolerant. Atomic Fa is *strictly true* (or *s*-true) iff every x similar_F to a is F (i.e. Fx is classically true) and *tolerantly true* (or *t*-true) iff some x similar_F to a is F; and ¬Fa is *s*-true (*t*-true) iff Fa is not *t*-true (*s*-true). Their preferred *st*-notion of entailment (where B is an *st-valid* consequence of $A_1,\ldots,A_n$ iff B is *t*-true wherever $A_1,\ldots,A_n$ are *s*-true) validates all tolerance conditionals $Fa_n \;\&\; a_n I_F a_{n+1} \rightarrow Fa_{n+1}$, where '$I_F$' means 'is similar to (relative to F)', as well as their universal generalization, since these are all *t*-valid. Any inference from Fa_n and $a_n I_F a_{n+1}$ to Fa_{n+1} will be *st*-valid, though these cannot be chained together to reach soritical conclusions, given the non-transitivity of *st*-entailment.

The critical problem with this proposal is its failure to explain how tolerance relations apply to borderline cases. Suppose Fa_n and Fa_{n+1} are both *merely* *t*-true (i.e. *t*-true but not *s*-true) for some borderline cases a_n and a_{n+1} that are similar_F. Assuming we have judged Fa_n, we *cannot* then apply tolerance and go on to judge Fa_{n+1}. That is because Fa_{n+1} cannot be inferred on the basis of Fa_n and $a_n I_F a_{n+1}$, since Fa_n is only *t*-true—even though (oddly enough) the tolerance inference is perfectly valid (i.e. valid in the *st*-sense), as is its corresponding tolerance conditional. Cobreros et al. (2010) presuppose a classical bivalent logic. This guarantees that reasoning by tolerance will fail for borderline cases right where some Fa_n is classically true yet Fa_{n+1} is classically false—which is no better than the traditional epistemicist story of tolerance failing at the unknowable point of a sharp cutoff (Williamson 1994).

Even if we allow for truth-value gaps in the underlying logic (a possibility not considered by Cobreros et al.), the existence of such problem pairs of similar_F borderline cases both merely tolerantly F is still guaranteed. *Proof:* By definition, any similarity_F relation ($\sim_F$) effects a tripartition: the strict Fs, the merely tolerant Fs/non-Fs, and the strict non-Fs. We make no assumptions about how this aligns or deviates from the tripartition of definite Fs, borderline Fs/non-Fs, and definite non-Fs. Consider the last classically true F item, a_T. Fa_T will be merely *t*-true. That is because a_T cannot be strictly F (otherwise, since anything similar_F to a strict F is itself truly F and $a_T \sim_F a_{T+1}$, a_{T+1} would also be truly F—contrary to assumption), nor can it be strictly not-F (otherwise, it would, being self-similar_F, be truly not-F—contrary to assumption). Now, a_{T+1} will be similar_F (since adjacent) to a_T and hence tolerantly F (given $a_T \sim_F a_{T+1}$ and Fa_T). So a_{T+1} must count as "borderline F", whatever the truth-value structure of the underlying logic. If bivalence holds, a_{T+1} will be the first classically false F item (since a_T is the last classically true F) and so must be borderline F (assuming the true-false cutoff lies within the borderline range). Otherwise, if bivalence fails and there are truth-value gaps, Fa_{T+1} is truth-valueless, so again a_{T+1} must be "borderline F" in either sense: as *neither truly F nor falsely F* (since Fa_{T+1} has no classical truth-value) or as *neither definitely F nor definitely not-F* (assuming the true-truthless cutoff lies within the borderline range—which it must, since lacking truth/falsity means lacking definite truth/falsity, i.e. being borderline). Thus, for some borderline a_{T+1}, (even the *t*-truth of) Fa_{T+1} cannot be inferred from Fa_T and $a_T I_F a_{T+1}$, despite the *st*-validity of $a_T I_F a_{T+1}$, $Fa_T \Rightarrow Fa_{T+1}$.

The account therefore fails to explain how tolerance permits extending F across the borderline range. Given that both premises and conclusion for such problematic tolerance inferences are merely *t*-true, the only other suitable version of logical consequence would be *t*-validity (i.e. preservation of *t*-truth). But this (as Cobreros et al. recognize) ceases to validate modus ponens, in which case one can no longer conclude Fa_{T+1} by applying modus ponens to Fa_T & $a_T I_F a_{T+1} \rightarrow Fa_{T+1}$. Therefore, either the axiom or rule form of tolerance fails to license extending F across borderline cases—hardly a vindication of tolerance.

The defeasibility perspective therefore appears to capture what adopters of "contextualist logic" cannot explain: dual deontic distinctions within tolerance-related sorites reasoning and their defeasible nature. In closing, I should stress that the Asher-Bonevac-Morreau theory remains just one among many possible ways to analyze tolerance in terms of deontic conditionals. No doubt, other contenders exist, with potentially more to offer. The current proposal simply demonstrates what a theory of defeasible reasoning, more generally, can achieve in addressing longstanding issues in the philosophy of vagueness.

Outlook

What lies in store for the epistemicist conception of vagueness?

Our discussion has closely examined both pros and cons of traditional versions of the epistemicist theory, as exemplified by the account in Williamson's *Vagueness* (1994). Much of the contemporary criticism was seen to be unfounded. Traditional epistemicism provides a compelling account of both the epistemology and the phenomenology of vagueness. The epistemicist principle of UNKNOWN—that vagueness-related indeterminacy entails ignorance—was seen to withstand the criticisms of Wright (2001), Dorr (2003), Barnett (2009), and Bobzien (2012). In comparison with their alternative proposals for explicating the epistemic consequences of vagueness, it remains the most plausible way to articulate the intuitive connection between vagueness and ignorance (see Chapter 2). It also provides a natural explanation for certain aspects of the phenomenology of vagueness, such as why our actual intuitions about vague matters fail to be response-dependent in the sense of enabling knowledge. Nonetheless, traditional epistemicist theories are perfectly equipped to accommodate the possibility of response-dependence, as in Barnett's (2010) hypothetical scenario of Zengland (see Chapter 1).

Yet traditional epistemicist accounts like Williamson's suffer a number of shortcomings. First, they fail to vindicate the seemingly undeniable principles of tolerance guiding our use of vague predicates. Not even Greenough's (2003) defense of "epistemic tolerance" meets this challenge, since such proposals merely give an epistemic interpretation of the so-called "gap principles" governing higher-order vagueness; as such, they do not constitute a defense of tolerance proper. Second, traditional epistemicism leaves no room for the open texture of vague predicates. This is one of the main reasons why contextualists like Shapiro (2006) and Soames (1999) prefer semantic treatments of vagueness that appear more conducive to explaining the permissibility phenomena related to vagueness. Yet Williamson's independent commitment to a knowledge norm of assertion rules out any freedom of verdict within the borderline range, provided everything there is unknowable (see Williamson 2000). Such an unavoidable posture of quietism strikes one as an unsatisfactory response in the face of vagueness.

Other reasons for discontent with traditional epistemicism are specific to Williamson's own version of the theory. One, Williamson's epistemicist account relies problematically on tenuous grounds for logical conservatism. Williamson offers partly pragmatist considerations in defense of classical logic and semantics: these should not be given up without due cause; fortunately, he argues, vagueness provides no compelling reason to relinquish, for example, the classical principles of bivalence and excluded middle. Williamson famously endeavored a global defense of bivalence: *any* meaningful sentence, whatever the nature of the discourse, must have exactly one classical truth-value, true or false (see his 1994:§§7.1,7.2). Yet the feasibility of the subsequent epistemic account of vagueness he proceeded to give rested entirely upon the success of his case for universal logical conservatism: rather high stakes for those solely concerned about vagueness. What right does a theory of vagueness have to any totalitarian imposition of classical logic and semantics *everywhere*, in all domains of inquiry, even those historically thought to motivate some form of logical revisionism? Witness the vast literature on future contingents, semantic paradoxes, presupposition failure, quantum indeterminacy, etc. No doubt Williamson's appeal to a global logical conservatism alienated many who were attracted to logical revisionism elsewhere but would otherwise have been sympathetic to the epistemicist program. Surely the theory can win wider appeal.

Two, of perhaps greater concern is the very nature of Williamson's arguments for logic conservatism. His pragmatist reasons for retaining classical logic and semantics, on grounds of theoretical simplicity and cross-domain unity—given how they are "vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains" (1994:186)—betray an underlying operative antirealist attitude, as if identifying the correct logic and semantics for treating ordinary vague languages were a mere issue of theoretical preference, constrained by considerations of theoretical choice. Shapiro better conveys the heart of the realist perspective when he quips: "Classical logic was developed with mathematical languages in mind. The logic of vague expressions will be what it will." (2006:72) Indeed, Williamson's epistemic treatment of vagueness, on one reading of his (1994), can be understood as serving to bolster a more general sustained effort to defend classical logic in its entirety, as applied to natural language semantics and beyond, with the upshot that challenges posed by vagueness in ordinary speech present no decisive reason to restrict the scope of applicability for classical logic and semantics.

Three, Williamson often appears to take the preservation of classical logic and semantics, bivalence and all, to be central to the epistemicist theory. This runs together two distinct sets of issues: the metaphysical existence of cutoffs and the epistemic status, unknowable or otherwise, of such facts. In general, Williamson's philosophy of vagueness exhibits a tendency to prioritize the metaphysical issues over the epistemology of vagueness—for instance, in his "Definiteness and Knowability" (1995) his arguments to support the epistemic principle that knowability entails determinacy rely upon appeals to classical rules of reasoning and excluded middle. Issuing controversial metaphysical claims for the purposes of establishing such basic epistemic results seems strategically gratuitous, and runs contrary to the spirit of the K-first methodology developed in his later *Knowledge and Its Limits* (2000). In contrast, Greenough's (2003) discussion served to separate the epistemology from the metaphysics of vagueness, by showing how—bivalence aside—the ignorance-entailing aspect of vagueness could be taken to be a relatively uncontroversial claim.

Four, Williamson's account does *not* strictly speaking constitute a form of epistemicism proper. For it does not take vagueness to be either a type of ignorance or any sort of phenomenon that is epistemic in nature. Instead, it identifies the meaning instability—a *semantic* feature—of our vague terms as the direct source of all vagueness-related ignorance: vagueness, although a distinctive source of ignorance, turns out to be a semantic phenomenon on this view (see Chapter 2). At best, the account presents a hybrid form of epistemicism, if such semantic instability proves to involve a characteristically epistemic element.

Five, Williamson appears to deny the reality of determinacy. There is no such thing as determinacy, according to his account, since such a notion finds no ultimate expression in his theory: any determinacy operator, he claims, lacks any intelligible non-epistemic interpretation and so must be reinterpreted as a kind of epistemic operator (see Chapter 2). Yet without a non-epistemic determinacy operator, it is difficult to see how one may make claims about the *source* of any vagueness-related ignorance (and not just the ignorance itself) or articulate any connections between vagueness and its epistemic consequences.

Six, there are theory-internal reasons to be dissatisfied with Williamson's account. His insistence on *exact* semantic knowledge of the meanings of our vague terms does not sit well with the claim that these terms constantly undergo undetectable shifts in meaning, given his margin for error constraints on knowledge. Nor does Williamson's underdeveloped account of semantic knowledge as induction into a linguistic practice offer much help in precluding the possibility of semantic ignorance or resisting the idea that any semantic knowledge is only inexact (see Chapter 4).

Hence my departure from Williamson. I propose a new direction for the epistemic conception of vagueness. Traditional epistemicism disregards our intuitions about the tolerance and open texture of vague predicates. I argued that these are features of vagueness that, once understood in epistemic terms, can be properly accounted for by the epistemic theory of vagueness. Chapter 5 outlined the general strategy for how to recover open texture on the epistemicist view. Chapter 6 argued for a normative interpretation of tolerance that, when combined with open texture, offers a diagnosis of the sorites paradox. A nonmonotonic logic is needed, however, to represent the defeasibility of such norms. To this extent the account diverges from Williamson's insistence on classical logic and semantics, at least for the purposes of reasoning about vagueness-related norms. As far as tolerance is concerned, sorites reasoning is not classically valid, only defeasibly valid. A fuller treatment would situate this proposal within a broader framework offering a more general explanation of the defeasible nature of normative reasoning. I leave this idea for future work.

The account remains officially neutral on the question of cutoffs. Some may balk: Is it not simply part of the epistemicist thesis that sharp cutoffs exist for vague predicates? I see no reason to believe so. The epistemicist conception of vagueness maintains that vagueness is to be conceived as an inherently or fundamentally *epistemic* phenomenon. Characterized as such, epistemicism makes no obvious or immediate claims about the existence of cutoffs or sharp boundaries. Why tether the view to a controversial metaphysics? Better, instead, to free it of any such metaphysical commitments. To insist otherwise is to fail to appreciate the distinction between the metaphysics and the epistemology of vagueness: these are two domains of inquiry concerning questions that are radically different in nature.

Of course, the normative interpretation of tolerance is entirely compatible with the existence of cutoffs. In that regard, my account of sorites reasoning can be viewed as a modified extension of Williamson's theory (modulo the caveats mentioned above). Accounting for the existence of cutoffs and their unknowability might then follow the general outlines of Williamson's margin for error model. This is a merit not a fault. I myself have never found the prediction of cutoffs to constitute a decisive refutation of the epistemicist theory. Why fault the theory for attracting incredulous stares on account of odd metaphysical claims? Queerness is the norm within metaphysics, which need not answer to intuition.

My own reluctance to endorse bivalence stems from the varying plausibility of claims to cutoffs: insisting there must be a fact of the matter is far more reasonable when vague matters really *do* matter (e.g. determining whether it is morally prohibited to abort a borderline viable fetus) than when matters detached from any serious practical consequence are seemingly open to arbitrary settlement (e.g. sorting paint chips or picking out bald men). Perhaps the verdict on cutoffs must wait until the normative consequences of vagueness have been more fully investigated. In the meantime, the epistemicist theory of vagueness has much to say regarding the epistemology, even if not the metaphysics, of vagueness.

References

Åkerman, J. and Greenough, P. 2009. Vagueness and non-indexical contextualism. In *New Waves in Philosophy of Language*, ed. S. Sawyer, 8–23. Stockholm: Palgrave Macmillan.

Åkerman, J. and Greenough, P. 2010. Hold the context fixed—vagueness still remains. In *Cuts and Clouds*, ed. R. Dietz and S. Moruzzi, 275–288. Oxford: Oxford University Press.

Asher, N. and Morreau, M. 1991. Commonsense entailment: a modal theory of nonmonotonic reasoning. In *Proceedings of the 12th IJCAI*, 387–92. Palo Alto: Morgan Kaufmann.

Asher, N. and Bonevac, D. 1996. *Prima facie* obligation. *Studia Logica* 57: 19–45.

Asher, N. and Bonevac, D. 2005. Free choice permission is strong permission. *Synthese* 145(3): 303–23.

Bacon, A. (manuscript) Vagueness at every order.

Barker, C. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25: 1–36.

Barnett, D. 2009. Is vagueness sui generis? *Australasian Journal of Philosophy* 87: 5–34.

Barnett, D. 2010. Does vagueness exclude knowledge? *Philosophy and Phenomenological Research* 82: 22–45.

Barnett, D. manuscript. Vagueness and rationality.

Bobzien, S. 2010. Higher-order vagueness, radical unclarity, and absolute agnosticism. *Philosophers' Imprint* 10(10): 1–30.

Bobzien, S. 2012. If it's clear, then it's clear that it's clear, or is it? Higher-order vagueness and the S4 axiom. In *Episteme, etc*. eds. B. Morison, K. Ierodiakonou. Oxford: Oxford University Press.

Bonevac, D. 1998. Against conditional obligation. *Noûs* 32(1): 37–53.

Cobreros, P., Egré, P., Ripley, D. and van Rooij, R. 2010. Tolerant, classical, strict. *Journal of Philosophical Logic* 41(2): 347–85.

Cobreros, P., Egré, P., Ripley, D. and van Rooij, R. 2012. Tolerance and mixed consequence in the s'valuationist setting. *Studia Logica* 100(4): 855–77.

Deas, R. 1989. Sorensen's Sorites. *Analysis* 49: 26-31.

Dorr, C. 2003. Vagueness without ignorance. *Philosophical Perspectives* 17: 83–113.

Dorr, C. 2009. Iterating vagueness. In Sebastiano Moruzzi & Richard Dietz (eds.), *Cuts and Clouds. Vagueness: Its Nature and Its Logic*. Oxford University Press, 2009.

Dummett, M. 1975. Wang's paradox. *Synthese* 30: 301–24.

Edgington, D. 1992. Validity, uncertainty, and vagueness. *Analysis* 52(4): 193–204.

Eklund, M. 2002. Inconsistent languages. *Philosophy and Phenomenological Research* 64(2): 251–75.

Fine, K. 1975. Vagueness, truth, and logic. *Synthese* 30: 265–300.

Gaifman, H. 2010. Vagueness, tolerance and contextual logic. *Synthese* 174: 5–46.

Goguen, J.A. 1969. The logic of inexact concepts. *Synthese* 19: 325–373.

Graff Fara, D. 2000. Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics* 28: 45–81. Originally published under the name "Delia Graff".

Graff Fara, D. 2008. Profiling interest relativity. *Analysis* 68(4): 326–35.

Greenough, P. 2003. Vagueness: a minimal theory. *Mind* 112: 235–81.

Harman, G. 1984. Logic and reasoning. *Synthese* 60(1): 107–27.

Hawthorne, J. and McGonigal, A. 2008. The many minds account of vagueness. *Philosophical Studies* 138: 435–440.

Horgan, T. 1994. Robust vagueness and the forced-march Sorites paradox. *Philosophical Perspectives* 8: 159–88.

Hull, G. 2005. Vagueness and ‘vague’: a reply to Varzi. *Mind* 114 (455): 689–93.

Hyde, D. 1994. Why higher-order vagueness is a pseudo-problem. *Mind* 103(409): 35–41.

Hyde, D. 2003. Higher orders of vagueness reinstated. *Mind* 112 (446): 301–5.

Johnston, M. 1989. Dispositional theories of value. *Proceedings of the Aristotelian Society* 63:139–174.

Kamp, H. 1981. The paradox of the heap. In *Aspects of Philosophical Logic*, ed. U. Mönnich, 225–277. Dordrecht: Reidel.

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

Keefe, R. 2007. Vagueness without context change. *Mind* 116: 275–92.

Kolodny, N. and MacFarlane, J. 2010. Ifs and oughts. *Journal of Philosophy* 107: 115–143.

Koons, R. 1994. A new solution to the Sorites problem. *Mind* 103: 439–449.

Kyburg, A. and Morreau, M. 2000. Fitting words: vague language in context. *Linguistics and Philosophy* 23 (6): 577–597.

MacFarlane, J. 2010. Fuzzy epistemicism. In *Cuts and Clouds*, ed. R. Dietz and S. Moruzzi, 438–463. Oxford: Oxford University Press.

McGee, V. and McLaughlin, B. 1994. Distinctions without a difference. *Southern Journal of Philosophy* 33: 203–51.

Morreau, M. 1997. Fainthearted conditionals. *Journal of Philosophy* 94(4): 187–211.

Parfit, D. (unpublished). What we together do.

Pettit, P. 1991. Realism and Response-Dependence. *Mind* 100(4): 587–626.

Raffman, D. 1994. Vagueness without paradox. *The Philosophical Review* 103(1): 41–74.

Raffman, D. 1996. Vagueness and context relativity. *Philosophical Studies* 81: 175–92.

Raffman, D. 2005. How to understand contextualism about vagueness: reply to Stanley. *Analysis* 65(3): 244–8.

Raffman, D. 2005. Borderline cases and bivalence. *The Philosophical Review* 114: 1–31.

Regan, D. 1980. *Utilitarianism and Cooperation*. Oxford: Clarendon Press.

Robertson, T. 2000. On Soames's solution to the Sorites Paradox. *Analysis* 60(4): 328–334.

Sainsbury, R.M. 1990. Concepts without boundaries. Inaugural lecture delivered 6 November 1990. Published by the King's College London Department of Philosophy, London. Reprinted in *Vagueness: A Reader*, eds. R. Keefe and P. Smith, 251–64. Cambridge: MIT Press, 1997.

Sainsbury, R.M. 1997. Easy Possibilities. *Philosophy and Phenomenological Research* 57(4): 907–919.

Shapiro, S. 2006. *Vagueness in Context*. Oxford: Oxford University Press.

Sider, T. 1995. Sorensen on unknowable obligations. *Utilitas* 7: 273–9.

Soames, S. 1999. *Understanding Truth*. Oxford: Oxford University Press.

Soames, S. 2002. Replies. *Philosophy and Phenomenological Research* 62: 429–52.

Soames, S. 2012. Vagueness in the law. In *The Routledge Companion to Philosophy of Law*, ed. M. Andrei, 95–108. New York: Routledge.

Sorensen, R. 1985. An Argument for the vagueness of 'vague'. *Analysis* 45:134–7.

Sorensen, R. 2001. *Vagueness and Contradiction*. Oxford: Oxford University Press.

Sorensen, R. 2008. Semivaluationism: Putting *Vagueness in Context* in Context. *Philosophy and Phenomenological Research* 76(2): 471–83.

Sorensen, R. 2012. The Sorites and the generic overgeneralization effect. *Analysis* 72(3): 444–9. Thomason, R. 1981. Deontic Logic and the role of freedom in moral deliberation. In *New Studies in Deontic Logic*, ed. Hilpinen, R. Dordrecht: D. Reidel.

Tye, M. 1990. Vague Objects. *Mind* 99: 535–57.

Tye, M. 1994. Why the vague need *not* be higher-order vague. *Mind* 103: 42–4.

Varzi, A.C. 2003. Higher-order vagueness and the vagueness of ‘vague’. *Mind* 112 (446): 295–8.

Varzi, A.C. 2005. The vagueness of ‘vague’: rejoinder to Hull. *Mind* 114(455): 695–702.

Waisman, F. 1945. Verifiability. *Proceedings of the Aristotelian Society* 19: 119–50. Reprinted in *Logic and Language*, ed. A. Flew, 122–51 (pp. 117–44 in earlier 1951 edition). Oxford: Blackwell, 1968.

Willer, M. 2012. A remark on iffy oughts. *Journal of Philosophy* 109: 449–461.

Williamson, T. 1992. Vagueness and ignorance. *Proceedings of the Aristotelian Society* 66: 145–77.

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

Williamson, T. 1995. Definiteness and knowability. *The Southern Journal of Philosophy* 33: 171–92.

Williamson, T. 1997. Imagination, stipulation, and vagueness. *Philosophical Issues* 8: 215–28.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Williamson, T. 2000. Margins for error: a reply. *The Philosophical Quarterly*, 50(198): 77–81.

Williamson, T. 2002. Reply to Machina and Deutsch on vagueness, ignorance, and margins for Error. *Acta Analytica* 17(29): 47–61.

Williamson, T. 2005. Vagueness in reality. In *The Oxford Handbook of Metaphysics*, eds. Loux, M. and Zimmerman, D. Oxford: Oxford University Press, 690–715.

Williamson, T. 2007. Knowledge within the margin for error. *Mind* 116(463): 723–6.

Wright, C. 1975. On the coherence of vague predicates. *Synthese* 30(3/4): 325–65.

Wright, C. 1987. Further reflections on the Sorites paradox. *Philosophical Topics* 15(1): 227–90.

Wright, C. 1992. Is higher-order vagueness coherent? *Analysis* 52: 129–39.

Wright, C. 2001. On being in a quandary. *Mind* 110: 45–98.

Wright, C. 2003. Vagueness: a fifth column approach. In *Liars and Heaps: New Essays on the Semantics of Paradox*, ed. Beall, JC. Oxford: Oxford University Press, 2004.

Wright, C. 2010. The illusion of higher-order vagueness. In *Cuts and Clouds: Vagueness, Its Nature, and Its Logic*, eds. Dietz, R. and Morruzzi, S. Oxford: Oxford University Press, 523–49.

Vita

Ivan J Hu was born in Concord, Massachusetts in 1986. The only son of Taiwanese immigrant parents, Michael and Tracy, he was raised in the suburbs of New England, primarily Nashua, New Hampshire and Westford, Massachusetts. He attended Phillips Exeter Academy between 2002 and 2004 and earned a Bachelor of Arts in philosophy at the University of California, Berkeley in 2008. Since 2008, he has resided in Austin, Texas while pursuing a Doctor of Philosophy in philosophy at the University of Texas at Austin.

The author may be reached by email at: ihu@utexas.edu

This manuscript was typed by Ivan J Hu.