Vagueness: A Minimal Theory[1]

Patrick Greenough

Vagueness is given a philosophically neutral definition in terms of an epistemic notion of tolerance. Such a notion is intended to capture the thesis that vague terms draw no known boundary across their range of signification and contrasts sharply with the semantic notion of tolerance given by Wright (1975, 1976). This allows us to distinguish vagueness from superficially similar but distinct phenomena such as semantic incompleteness. Two proofs are given which show that vagueness *qua* epistemic tolerance and vagueness *qua* borderline cases (when properly construed to exclude terms which are stipulated to give rise to borderline cases) are in fact conceptually equivalent dimensions of vagueness, contrary to what might initially be expected. It is also argued that the common confusion of tolerance and epistemic tolerance has skewed the vagueness debate in favour of indeterminist over epistemic conceptions of vagueness. Clearing up that confusion provides an indirect argument in favour of epistemicism. Finally, given the equation of vagueness with epistemic tolerance, it is shown that there must be radical higher-order vagueness, contrary to what many authors have argued.

1. Overview

The broad aim of this paper is to give a rigorous characterization of vagueness from a perspective which is as neutral as possible on matters both logical and philosophical. In so doing, the foundation is laid for what may be called a *minimal theory of vagueness*. One key merit of this theory is that it promises to ensure that the dialectic of the vagueness debate can at least begin at a mutually agreed point—this theory can at least ensure that we are all taking about the same thing from the outset in our inquiry into the nature and source of vagueness. In setting forth this minimal theory, three related dimensions of vagueness are distin-

[1] The basic ideas in this article were conceived in the Summer of 2000, but they did not receive their first outing until a workshop on vagueness held at the Institute of Philosophy, University of Bologna, 22–23 November, 2001. Particular thanks to Andrea Sereni for his stimulating reply. A shorter version of this talk was given at the Fourth European Congress for Analytic Philosophy, Lund University, 14–18 June, 2002. Very useful feedback on earlier drafts of this paper (or of the ideas therein) was kindly given by John Burgess, Timothy Chambers, Annalisa Coliva, Richard Dietz, Katherine Hawley, James Ladyman, Michael Lynch, Fraser MacBride, David McCarthy, Sebastiano Moruzzi, Graham Priest, Mark Sainsbury, and Stewart Shapiro. Particular thanks are due to Dominic Hyde, Sven Rosenkranz, Crispin Wright, and an anonymous referee, who all supplied many suggestions for improvement.

guished: vagueness *qua* sorites-susceptibility, vagueness *qua* borderline cases, and vagueness *qua* tolerance. Hitherto, the relationship between these dimensions has remained somewhat unclear. The minimal theory of vagueness is equipped to remove much, if not all, of that unclarity. Of perhaps greater interest, is that this theory entails that there must be radical higher-order vagueness—a subject about which there has been vigorous dispute. So while the axioms of this minimal theory are uncontroversial in the first instance, some of its theorems turn out to be decidedly controversial.

As a preliminary to such investigations, it is necessary to inquire as to what we might reasonably expect or demand from a minimal theory of vagueness. Can this theory solve the sorites paradox? Can it isolate the source of linguistic vagueness? Can this theory successfully rehabilitate what Sainsbury (1991) has called the 'characteristic sentence approach' to defining vagueness? These are the sorts of questions addressed in section 2. In section 3, it is found that vagueness defined as sorites-susceptibility offers the least controversial characterization of vagueness. However, this characterization proves to be too insubstantial for the promises of the minimal theory to be properly satisfied. On what is perhaps the most prevalent conception, vagueness is the phenomenon of borderline cases (Sorensen 1985; Williamson 1994; Sainsbury 1995; Tye 1995). Whether or not it is plausible to give an uncontroversial definition by reference to such a phenomenon is the key issue of section 4 through to section 6. A number of non-epistemic and epistemic accounts of what it is to be a borderline case are scrutinized. For the purpose of finding a neutral definition of vagueness, none of these proves entirely satisfactory. The particular bug-bear proves to be the possibility of terms which we can stipulate to give rise to borderline cases but which draw sharp and clearly identifiable divisions across their associated dimension of comparison. Prima facie, it is far more plausible to define vagueness minimally by reference to an epistemic notion of tolerance. Such a notion is intended to capture the thesis that vague terms draw no clear or known boundary across their range of signification and contrasts sharply with the semantic notion of tolerance given by Wright (1975, 1976.) In section 7, the identification of vagueness with epistemic tolerance is exploited so as to give a rigorous but nonetheless neutral definition of 'is vague'. This definition allows us to distinguish vagueness from superficially similar but distinct phenomena such as semantic incompleteness. In section 8, two proofs are given which show that vagueness *qua* borderline cases (when properly construed to exclude terms which are stipulated to give rise to borderline cases) and vagueness *qua* epistemic tolerance are in fact conceptually equivalent, contrary to what might initially be expected. In section 9, it is argued that the common confusion of tolerance and epistemic tolerance has skewed the vagueness debate in favour of indeterminist over epistemic conceptions of vagueness. Clearing up that confusion provides an indirect argument in favour of epistemicism. Finally, in section 10, given the equation of vagueness with epistemic tolerance, it is shown that there must be radical higher-order vagueness, contrary to what many authors (for example, Burgess 1990, 1998; Wright 1987, 1992b; Koons 1994) have argued. Radical higher-order vagueness is a fact of life for everyone.

2. Minimalism and vagueness

The expression 'minimal theory of vagueness' is ambiguous. On the one hand, it can be used to mean the sort of theory that is, can, or ought to be endorsed by those who sponsor some form of minimalism concerning truth (see for example, Horwich 1997, pp. 929–935; Field 2001, Ch. 8). On the other hand, it can be used to mean the sort of theory which endeavours to set forth some a priori, basic, and platitudinous principles which provide an uncontroversial definition of vagueness, a definition which isolates the constitution of vagueness from a perspective which is as neutral as possible on matters logical and philosophical. It is the latter sort of theory, which we may simply call *the minimal theory of vagueness*, which will be the preoccupation of this paper.[2] What should we reasonably expect or demand from such a theory?

The traditional goal of any theory of vagueness has broadly been three-fold: to solve the sorites paradox (in all its many guises), to identify the source of vagueness, and to distinguish the vague from the non-vague. We should not expect or demand that a minimal theory of vagueness be able to furnish a generally acceptable solution to the sorites paradox. Were such an uncontroversial solution available that would indeed be gratifying; but no such solution seems in prospect. Likewise, while all parties can agree that much of natural language is vague we should not demand that our minimal theory identify the

[2] Wright (1992a) has argued that the content of 'is true' (and cognate expressions) is entirely exhausted by a 'network' of platitudes concerning truth (such as Tarski's T-schema, the thesis that to assert is to present as true, the thesis that warranted assertibility and truth are distinct, the thesis that truth is absolute, and so on). Here, the idea is to apply a similar methodology to the predicate 'is vague' and cognate expressions. The challenge will be to locate a set of platitudes concerning vagueness which at least exhaust our everyday understanding of this term while not pretending to thereby identify the underlying nature of vagueness.

source of this linguistic vagueness—that is the business of some substantive, controversial conception. The theory developed below is nonetheless equipped to distinguish the vague from the non-vague; that is, this theory is able to isolate the necessary and sufficient conditions for when a sentence counts as vague. As such, this theory is still able to address the following range of questions: Is vagueness at bottom the phenomenon of borderline cases? In what way can we say that vague language is *tolerant*? How do the different dimensions of vagueness relate? What is it to say that an expression suffers from higher-order vagueness? Is there higher-order vagueness? Perhaps the minimal theory is rich enough to address many more questions than these; but these are the ones which are addressed below. To that end, two immediate qualifications are in order.

First, given that this minimal theory will not locate the source of linguistic vagueness, nor furnish a solution to the sorites paradox, then clearly this theory will not be the last word on vagueness. Nonetheless, it is important to note that this does not entail that our minimal theory of vagueness is unable to give a rigorous definition of 'is vague'. Arguably, one can grasp the constitution of vagueness without thereby grasping how one might solve the sorites paradox. This is entirely analogous to the predicament faced by the truth-minimalist who holds that one can grasp the essence of truth via grasping the veracity of certain truth-platitudes (such as Tarski's T-schema) without thereby grasping how one should solve the liar paradox.

Secondly, it was mentioned above that the minimal theory of vagueness should be as neutral as possible not only on philosophical matters *but on logical matters also*. It thus seems that the minimal theory of vagueness can only be developed using the some suitably uncontroversial (and thereby very weak) background logic. Since there is little agreement as to the correct logic of vague language, and moreover since there is no general agreement about the correct logic for the language we are entitled to employ in theorizing about vagueness (that is, what we may loosely call the meta-language), then the project of developing an acceptable minimal theory of vagueness looks rather bleak. If we take that worry seriously then it looks as though there is no scope for the vagueness debate to begin at a mutually agreed point. Every candidate minimal characterization of vagueness proposed would presuppose some background logical principles which have been, or at least might be, disputed by partisans to the debate as a whole. There would thus be a very real sense in which there would be no genuine disagreement about the character of vague language at all since each partisan would mean something different by the predicate 'is vague'. This (familiar) worry is deep, but not insurmountable.

One way to combat this concern is to adopt what may be termed the *bold* strategy of adopting classical logic as the background logic in which we theorize about particular features of vagueness until such point as such an adoption proves to generate tangible controversy. Williamson (1997) has recommended a similar bold strategy as the most sensible methodology one can adopt when approaching philosophical problems (particularly the problem of vagueness) which might initially seem to demand a revision of classical logic in the object-language or meta-language. Williamson urges that

> one holds one's logic fixed, to discipline one's philosophical thinking [because] in the long-run the results of the discipline will be more satisfying from a philosophical as well as from a logical point of view. (p. 218)

In contrast to Williamson, the suggestion here is not that we should adopt classical logic (in either the meta-language or object-language) to discipline one's philosophical thinking about *every* aspect of vagueness (including those aspects of vagueness that we must take account of in addressing the sorites or identifying the source of linguistic vagueness). Rather, the suggestion is that we should retain classical logic until such time as this proves to undermine the goal of the minimal theory to furnish an uncontroversial basic characterization of vagueness. It is a further question whether Williamson's methodology is appropriate to the development of a substantial theory of vagueness. That further question need not worry us here. It is enough that we have a rationale with which to begin our minimal investigations.

So how then might we minimally define vagueness? One way in which we might do so is via what Sainsbury has called 'the characteristic sentence' approach. This involves finding

> a sentence schema, containing a schematic predicate position, such that the sentence resulting by replacing the schematic element by a predicate is true if that substitute is a vague predicate. (Sainsbury 1991, p. 170)

Something like the characteristic sentence approach was first offered by Wright (1987, pp. 282–8), while Sainsbury's employment of this approach is merely provisional as he proceeds to argue that one cannot satisfactorily identify vagueness in this way. One aim of this paper to isolate a characteristic sentence which is entailed by all conceptions of vagueness, including, I take it, the conception offered in Sainsbury (1990, 1991).

To put the characteristic sentence approach to work, we need to adjust it in two key respects. Firstly, we must generalize this schema to accommodate the possibility of sentential vagueness which does not depend on predicate vagueness. (In so doing, in much of what follows we can conveniently focus on sentential vagueness rather than predicate-vagueness or subject-vagueness.) Secondly, in order for this approach to offer a rigorous minimal definition of vagueness we must demand that satisfaction of the sentence schema is not merely necessary but also *sufficient* for the vagueness of some sentence. We thus need to find a schema, containing a schematic sentence position or positions, such that the sentence resulting by uniformly replacing the schematic elements by a particular sentence is true if and only if that substitute is a vague sentence.

Can we isolate a sentence schema which would be acceptable to all partisans? One natural place to start is by looking at the property of being *sorites-susceptible*.

3. Vagueness *qua* sorites-susceptibility

Arguably the most general (and least controversial) way to characterize (sentential) vagueness is by reference to the sorites paradox.[3] Say that: a declarative sentence is vague just in case this sentence is *sorites-susceptible*. Can we isolate an uncontroversial characteristic sentence which exploits this basic feature of vague expressions?

Suppose we have some sentence S such that the truth of S in a case α depends only on the value $\nu(\alpha)$ taken by some discretely or continuously varying parameter ν in α, where ν (let us say) takes non-negative (real) numbers as values. For example, in a simple case, if S is the sentence 'the bath is hot', then ν will be the temperature of the water in the bath.[4] Where c is some small positive real number, then it seems ini-

[3] In what follows, we shall take the primary bearers of vagueness to be declarative sentences, rather than statements or propositions. That may be a controversial step in developing a substantial conception of vagueness (see for example, Williamson 1994, p. 187), but nothing particularly turns on this issue when developing the minimal theory of vagueness.

[4] Vague terms are typically associated with some dimension or dimensions of comparison. The predicate 'is tall' is one-dimensional (with respect to some comparison class) as it merely governs the dimension of heights. The concept *tall* characteristically takes a positive, a comparative, and a superlative: a person can be tall, taller, and the tallest. The predicate 'is humid' governs (at least) two dimensions: the temperature and water content of air. Colour predicates govern the three dimensions of hue, saturation, and brightness. The predicate 'is hirsute' is multi-dimensional: the thickness, length, colour, texture, distribution, and number of hairs all affect its application. It is to be noted that type of vagueness I shall be concerned with in this paper is what Alston (1967, p. 219) has called 'degree-vagueness', which he defines as 'lack of precise boundaries'. I shall not be concerned with another kind of vagueness isolated by Alston (p. 220), namely, where there is 'a variety of conditions, all of which have something to do with the application of [a] term, yet [we] are not able to make any sharp discriminations between those combinations of conditions which are, and those combination of conditions which are not, sufficient and/or necessary for application'.

tially plausible to say that:

(**SS1**) $\forall\alpha\ \forall\beta$, if $|\nu(\beta)-\nu(\alpha)|<c$ then S is true in α if and only if S is true in β

(where α and β range over actual and counterfactual cases). This is to say that where the difference between the value taken by ν in α and the value taken by ν in β is suitably small, the sentence S will be true in both cases or neither. If the temperature of the bath-water in the α case and the temperature of the bath-water in the β case differs only slightly, then the bath is either hot in both cases or not-hot in both cases. Another (classically equivalent) way of formulating this claim is to say that there are no cases α and β (across which the parameter ν varies by some small amount) whereby S is true α and not-S is true in β, which we express as follows:

(**SS2**) $\neg\exists\alpha\ \exists\beta$ such that $|\nu(\beta)-\nu(\alpha)|<c$ and S is true in α and not-S is true in β

(again, where α and β range over actual and counterfactual cases). The immediate suggestion, then, is to employ the schemas **SS1** and **SS2** as characteristic sentences. The idea is that if a given substitution S_1 of S, makes the characteristic sentence **SS1** (or **SS2**) true then the sentence S_1 is vague; conversely, if S_1 is vague then **SS1** (and **SS2**) will be true.[5] Any conception of vagueness which can or does define vagueness via the characteristic sentences **SS1** or **SS2** we may call a minimal conception of vagueness *qua* sorites-susceptibility. On this conception vagueness just is sorites-susceptibility. Have we given a satisfactory minimal definition of vagueness?

The trouble with the schemas **SS1** and **SS2** is that they can both be used to generate paradox — at least given further (prima facie) plausible assumptions.[6] Let's take each schema in turn. The most familiar sorites template can be given as follows, where we employ the sentence

[5] Notice that (i) satisfaction of the characteristic sentence leaves it open whether the vagueness of S_1 issues from the predicate or subject terms contained in S_1, or indeed from both types of term, and (ii) any substitution of S must be the sort of sentence whose truth is determined by the degree of variation in one or more graded or continuous parameters $\nu_1,\ldots\nu_n$. (These qualifications will often be left inexplicit in the rest of this paper.)

[6] Such as the assumption that certain rules of inference are valid, but also the further assumption that vague sentences do indeed express propositions. Not every substantial theory of vagueness ratifies this further assumption (and kindred assumptions). For example, Frege (1979, p. 155) seems to have thought that the sorites paradox shows us that certain concept expressions (for example, 'heap') fail to properly circumscribe a concept, such that a sentence such as '*a* is a heap' fails to properly express a thought (or proposition). Whether or not such a view, or variants of such a view, can successfully combat the sorites lies outwith the scope of the present paper.

'A bath of $n°$ is hot', and where premiss **A2** (the so-called induction step) is effectively derived from **SS1:**[7]

(**A1**) A bath of water temperature 100° is hot

(**A2**) For all n, if a bath of $n°$ is hot then a bath of n–1° is hot

(**A3**) A bath of water temperature 0° is hot.

Call this general form of the sorites the *A*-sorites. Premiss **A2** appears to be highly plausible; premiss **A1** (the induction base) appears unimpugnable; and the absurd **A3** is derived either by mathematical induction or via one hundred applications of ∀-elimination and *modus ponens*. All sides can agree that the *A*-sorites represents a logical paradox in that by apparently valid reasoning from apparently sound premisses one can derive a patently absurd conclusion. The key premiss **A2** codifies the (initially) highly plausible thought that a drop of temperature of one degree cannot make the difference between a hot bath and a bath which is not hot. (Were one to think that one degree could mark the difference then we need only consider a smaller c-value such as 0.001°.) Despite the initial plausibility of **A2**, one might nonetheless feel logically obligated to treat this paradox as a *reductio* of **A2**. To do so is to be committed to

(**A4**) There is an n, such that a bath of $n°$ is hot and a bath of $n-1°$ is not hot.

which on the face of it is just to say that 'is hot' is not after all a vague predicate. Our naïve intuitions seem to tell us that principles like **A2** are true of vague sentences and false of non-vague sentences. Our naïve intuitions thus generate paradox.[8]

[7] Strictly speaking, we need the Tarskian schema *S* is true if and only if *p* (where *p* is a translation of *S*) to make the derivation, but one could equally frame the *A*-sorites in the formal mode of speech.

[8] The principle **SS2** generates a different form of the sorites paradox (given further uncontroversial assumptions), as follows: Our naïve intuitions also tell us that

(**B1**) There is no n such that a bath $n°$ is hot and a bath of $n-1°$ is not hot

(where **B1** is derived from **SS2**). It is uncontroversial that a bath of 0° is not hot; but let us also suppose for *reductio* that a bath of 1° is hot:

(**B2**) A bath of 0° is not hot

(**B3**) A bath of 1° is hot

which entails, given ∧-I and ∃-I:

Since **SS1** and **SS2** lead to inconsistency (in many formal systems—including both intuitionistic and classical logic), these schemas cannot ground an uncontroversial definition of vagueness. If they did, they would forbid the view that vague language is both consistent and subject to classical logic or intuitionistic logic (cf. Sainsbury 1991, p. 172). In general, very few substantive conceptions of vagueness do indeed sanction either **SS1** or **SS2**—plausible as these principles might initially seem.[9] In response to this worry, one might offer a more anodyne characteristic sentence of the same general form. Perhaps something like the following can be employed to capture the minimal constitution of vagueness:

(**SS3**) Pre-theoretically (or, according to our naïve intuitions) **SS1** appears to be true

(**SS4**) Pre-theoretically (or, according to our naïve intuitions) **SS2** appears to be true.

A characteristic sentence like **SS3** would, I take it, reflect the fact that the major premiss **A2** of the *A*-sorites is *apparently* or *seemingly* true when one is first exposed to this paradox. (Likewise for the key premiss **B1** of the *B*-sorites.) To be sorites-susceptible is not to be committed to absurdity *per se* but is rather to be *seemingly* subject to soundness of the sorites paradox. On this basis, say that a sentence is vague just in case it is sorites-susceptible in the sense just given, just in case when substi-

---

(**B4**) There is an n such that a bath $n°$ is hot and a bath of $n-1°$ is not hot

which contradicts **B1**; and so by negation-introduction we infer:

(**B5**) A bath of 1° is not hot.

If we further suppose that

(**B6**) A bath of 2° is hot

then by parallel reasoning we can infer that there is an n such that a bath $n°$ is hot and a bath of $n-1°$ is not hot. Contradiction. Reject **B6** to infer that a bath of 2° is not hot. One hundred applications of this inference pattern allow us to infer the absurd result that a bath of 100° is not hot. Paradox. Call this general form of the sorites the *B*-sorites. (Wright (1987, p. 261) dubs this form of the paradox the 'No Sharp Boundaries Paradox'.) Were we to feel logically obligated to treat this as a *reductio* of **B1** then (given classical logic) we would be committed to **A4**, which again just seems on the face of it to rule out the obvious vagueness of the predicate 'is hot'. The *A*-sorites and *B*-sorites are the two main sorites templates which any substantial theory of vagueness must seek to defuse in some appropriate fashion. How one might do this need not concern us here.

[9] The conceptions of vagueness offered by Dummett (1975), Unger (1979), Wheeler (1979), and Hyde (1997) all sanction **SS1**.

tuted into **SS3** and **SS4** it renders these schemas true. Surely this ought to be agreeable to all?[10]

While the characteristic sentences **SS3** and **SS4** are, presumably, uncontroversial, they are also unspecific. How are we to give a rigorous (and uncontroversial) account as to what is meant by the qualifiers 'pre-theoretically', 'according to our naïve intuitions' or 'appears to be'? Moreover, we also need an account of just why it is that our naïve intuitions incline us to accept **SS1** and **SS2**. Perhaps some deeper feature of vague expressions explains that inclination in which case vagueness defined in terms of sorites-susceptibility is not a particularly informative characterization. So while a definition employing **SS3** and **SS4** may record a genuine conceptual (if rather unspecific) insight, and while it goes some way to ensuring that partisans to the vagueness debate are not talking past each other, it does not seem to record an *explanatory* insight. There is a strong sense in which a sentence is sorites-susceptible *because* it is vague, and not vice versa. Sorites-susceptibility is secondary in the explanatory order. We should look elsewhere for our minimal definition of vagueness.

4. Minimal vagueness *qua* borderline cases: the minimal indeterminist conception

On what is perhaps the most prevalent conception, vagueness is the phenomenon of borderline cases. This conception is so widespread that Sorensen (1985, pp. 135–6) can confidently say that:

> Although there is considerable disagreement over the nature of the defectiveness and exact nature of vagueness, there is general agreement that predicates which possess borderline cases are vague predicates.[11]

How then might we define the relevant notion of borderline case from within a conception which is as neutral as possible on matters logical and philosophical? Wright has offered the thought that:

> when dealing with vague expressions, it is essential to have the expressive resources afforded by an operator expressing definiteness or determinacy. (1987, p. 262)

[10] Sorensen (1985), for instance, takes sorites-susceptibility to be one hallmark of the vague, as do Keefe and Smith (1996, p. 3).

[11] The tradition of defining vagueness primarily in terms of borderline cases dates back to Peirce (1902, p. 748), was continued by Black (1937, p. 30), and receives its fullest expression in Fine (1975).

If this insight is correct then it will hold just as much for the minimal theory of vagueness we are trying to articulate here as it will for any further substantive conception of vagueness. So in this section, taking our lead from Wright, let us pursue the provisional strategy of defining a notion of vagueness *qua* borderline cases via some (as yet unspecified) notion of definiteness or determinacy. But which notion is to be employed—definiteness or determinacy? It will do no harm to assume, following Williamson (1996, p. 44), that definiteness and determinacy, in the context of vagueness at least, are fully interchangeable notions.

A first promising candidate characteristic sentence for the indeterminist minimal theory might be given as follows:

(**DT1**) $\exists\alpha$ In α, it is not determinately the case that *S* is true and it is not determinately the case that not-*S* is true (where 'α' ranges over both actual and counterfactual cases).

Again, the idea is that if a given substitution S_1 of the schematic sentence *S*, makes the characteristic sentence **DT1** true then the sentence S_1 is vague; conversely, if S_1 is vague then **DT1** will be true.[12] Any conception of vagueness which can or does define vagueness via the characteristic sentence **DT1** we may call an *indeterminist* minimal conception of vagueness *qua* borderline cases. Such a conception is intended to capture the thought that a sentence is vague just in case it takes a status intermediate between determinate truth and determinate non-truth (falsity). Is this characterization defensible?

Satisfaction of **DT1** cannot be sufficient for the presence of vagueness. A familiar complaint in this regard is that it is a mistake to take borderline cases *per se* to be constitutive of vagueness.[13] To illustrate, suppose we stipulate that the open sentence '*x* is an oldster' is determinately true of every person sixty-eight years of age and over, determinately false of those persons sixty-five years of age and under, and neither determinately true nor determinately false of the remainder. If a speaker applies this term to persons who are between sixty-five and sixty-eight then we are entitled to say that they have done something not quite right and done something not quite wrong according to the

[12] The fact that 'α' ranges over both actual and counterfactual situations allows us to capture Fine's distinction between intensional and extensional vagueness (Fine 1975, p. 266). A sentence is extensionally vague just in case it does give rise to borderline cases (given the way the actual world is) and is intensionally vague just in case it *could* give rise to borderline cases. The sentence 'Timothy Williamson is thin' is extensionally vague (as he concedes) and remains intensionally vague in situations where all people are either determinately thin or determinately not thin.

[13] See Dummett (1973, pp. 646–7); Wright (1975, p. 329); Sainsbury (1989, p. 34–5, 1991, pp. 173); Hyde (1994, pp. 35–6).

dictates of the stipulation. But since '*x* is an oldster' is neither determinately true nor determinately false of the intermediate cases then, given **DT1,** it counts as vague, even though intuitively we should be inclined to say that the term is *not* vague but rather, in some sense, semantically incomplete. This species of indeterminacy *per se* is not vagueness, since the term 'oldster' draws a perfectly sharp and clearly identifiable threefold division across its associated dimension of comparison.

In reply to this problem, it might be said that the problem of the sentence '*x* is an oldster' stems in essence from failing to accommodate the possibility of higher-order vagueness in setting forth our characteristic sentence. What is meant by higher-order vagueness? Very, very roughly, say that a sentence is higher-order vague just in case it not only gives rise to borderline cases (cases where it is neither determinately true nor determinately false), but borderline cases to those borderline cases (cases where it is neither determinately true nor determinately false that the sentence is neither determinately true nor determinately false), and borderline cases to these borderline cases, and so on. The sentence '*x* is an oldster' would count as genuinely vague if it were also to give rise to borderline cases to the borderline cases, and in turn borderline cases to those borderline cases, and so on. Since it does not, it is not genuinely vague. But can we rehabilitate a constitutive indeterminist minimal account *qua* borderline cases by appealing to some form of higher-order vagueness without at this stage giving an explicit (and perhaps controversial) model of higher-order vagueness?

One way to do this is to borrow the strategy of Hyde (1994) who has offered the thought that the expression 'borderline case' is ambiguous between the type of borderline cases that stem from such terms as 'oldster', and genuine borderline cases of vagueness where higher-order vagueness 'is built in from the very start' (ibid., p. 40). How then might we adjust the characteristic sentence **DT1** to accommodate Hyde's requirement that vagueness *qua* borderline cases automatically ensures that radical higher-order vagueness is built in from the outset? Hyde (1994, p. 39) in effect suggests that one need not make an *explicit* reference to the existence of higher-order borderline cases in our characterization of vagueness at least in so far as we ensure that we have distinguished the type of indeterminacy that is constitutive of vagueness (call it **indeterminacy**) and the type of indeterminacy (just call it indeterminacy) that is characteristic of such terms as 'oldster'. For Hyde, there is no real problem of higher-order vagueness; but rather the

> problem surrounding higher orders of vagueness arises when one tries explicitly to state something about the nature of vagueness that manifests itself in the characterisation anyway—the phenomenon of higher-order vagueness ... There are border border cases for vague predicates, but this need not be stated as part of the analysis of the concept of predicate-vagueness ... One is simply repeating oneself and adding nothing new. (ibid., p. 40)

So while Hyde accepts that higher-order vagueness is genuine feature of vague language (see p. 40), he rejects any suggestion that one can *explain* what vagueness *qua* borderline cases amounts to by reference to the thesis that a vague expression gives rise to borderline cases, and borderline cases of those borderline cases, and so on (or indeed by reference to any more rigorous statement of higher-order vagueness). It is not that Hyde disallows us from expressing what is meant by higher-order vagueness in this way, it is rather that in doing so, one adds nothing to our understanding of vagueness: in the order of explanatory priorities, our grasp of this thesis (or our grasp of a more rigorous formulation of higher-order vagueness) is secondary to our grasp of the basic notion of **indeterminacy.**

If this is right then one can simply side-step the problem of the sentence '*x* is an oldster' by offering the following characteristic sentence:

(**DT2**) $\exists\alpha$ In α, it is not **determinately** the case that *S* is true and it is not **determinately** the case that not-*S* is true.

Again, a sentential substitution is vague just in case **DT2** is true for that substitution.

Should all partisans to the dispute accept **DT2**? An immediate worry with **DT2** is that it is not yet a settled question whether any respectable theory of vagueness should indeed entail that vague terms are higher-order vague—as the requisite notion of **indeterminacy** demands. Since not all theories of vagueness entail the existence of higher-order borderline cases (not even implicitly) then it looks as if we have gained a sufficient condition for the indeterminist minimal definition at the expense of losing a necessary one—and with it we appear to have lost the promise of ensuring that the dialectic of the vagueness debate can begin at a mutually agreed point.[14] The tempting reply is to say that a theory which doesn't recognize higher-order vagueness is just obviously misconceived. We are not yet in a position to settle this matter until section 10, where it is shown that there must be (radical) higher-order vagueness. But even if we ought to recognize the existence of

[14] Wright (1987, 1992b), Burgess (1990, 1998), and Koons (1994), amongst others, have doubted the existence of (non-terminating) higher-order vagueness. See section 10 below.

higher-order vagueness (*qua* borderline cases), and even if we are willing to accede to Hyde's thesis of the ambiguity of 'determinately', one might nonetheless worry that the notion of determinacy is as yet too unspecific to feature in our minimal theory of vagueness. Is this a *sui generis* notion or can we explicate its nature in some substantial way? Is it an entirely non-epistemic notion? Indeed, do we really need a notion of determinacy at all in giving a (minimal) characterization of vagueness (as Hyde, following Wright, thinks)? If we are to put **DT2** to any use in our minimal theory we must first address these questions.

5. Determinacy and definiteness: a brief survey

One worry one might have with **DT2** (and indeed **DT1**) is that natural language does not in fact contain the requisite predicates 'is determinately true/false' or their material-mode counterparts 'It is determinately true/false that'. Moreover, if natural language were appropriately extended to included these operators, then the intuitions of natural language speakers could not be reliably employed to determine whether the appropriate instances of some characteristic sentence were true or not (see Sainsbury 1991, p. 174). This worry is well-taken; but in reply it might be said that the intuitions of native speakers are not an issue when one is attempting an *explication* (that is, a logical reconstruction) rather than some mere (descriptive) analysis of the phenomenon of vagueness. So, the reply runs, there ought to be no principled obstacle to specifying an extended language in which a notion of definiteness or determinacy is suitably introduced (cf. Williamson 1999, p. 129, fn.2). (Perhaps indeed such a regimentation would appropriately disambiguate 'determinately' along lines suggested by Hyde.) For the purposes of isolating a minimal theory of vagueness, how might we explicate the relevant notion of determinacy (and its cognate notion of definiteness)?

Dummett (1978, p. 256) has said that 'in connection with vague statements, the only possible meaning we could give to the word "true" is that of "definitely true"'. Likewise, for 'false' and 'definitely false'. This suggests that—contra Wright (1987)—we can do without talk of definiteness and determinacy in our characterization of vagueness, for if Dummett is correct, to say that a statement is (extensionally) vague is really to say no more than that it is neither true nor false. Dummett nonetheless maintains that we should not dispense with the 'determinately' operator for the following reasons: the notion of truth which is relevant to vague statements is a non-distributive notion (see Dummett

1991, pp. 75–6). In particular, Dummett argues that a disjunction '*x* is either orange or red' can be true even though *x* is on the red–orange borderline such that neither disjunct is true (which is not to say that one or more disjuncts is false). But how can we mark the difference between non-distributive true disjunctions and distributive true disjunctions? Enter the adverb 'determinately' (or 'definitely'). This adverb has, for Dummett, a special *force*—it can be used to record the fact that a disjunction is not only true, but that it is true in virtue of the fact that at least one of its disjuncts is true. It is for this reason that Dummett urges we should always formulate the principle of bivalence as saying that every statement is *determinately* either true nor false, so as to rule out the possibility that a class of statements are all either true or false, but that it is not (determinately) true which. Whatever the merits of this proposal, it is clear that this analysis is not compatible with any conception of vagueness in which truth does distribute over disjunctions. In general, we should demand that the minimal theory of vagueness must not exclude from the outset that the logic and semantics of vagueness is classical. To do that is to rule out the possibility that vagueness is an entirely epistemic phenomenon which demands no restriction of classical semantics or classical logic. There is no hope that a Dummettian conception can illuminate the sense of 'definitely' we require for **DT2**.

The point generalizes. From the minimal perspective, it is illegitimate to confer a non-epistemic interpretation to the adverbs 'determinately'/'definitely'—tempting as that construal may be. Determinate/definite truth, on the minimal conception of vagueness, cannot be taken to mean *truth to degree 1* (as on a many-valued conception of vagueness), or *truth under all admissible sharpenings* (as on a supervaluational conception), for again we must allow that vagueness might be, after all, a special species of ignorance. As Wright has said, it cannot be a basic datum that indeterminacy is a non-epistemic phenomenon, for this is just to saddle ourselves from the outset with a 'proto-theory' of vagueness (Wright 1995, pp. 133–4; see also Horwich 1997, pp. 929–30).

The point also applies to those who have offered what we might term *quasi-semantic* interpretations of determinacy or definiteness. McGee and McLaughlin (1995, p. 209) suggest that 'to say that an object *a* is definitely an *F* means that the thoughts and practices of speakers of the language determine conditions of application for the word *F*, and the facts about *a* determine that these conditions are met'. Statements are vague, accordingly, when the thoughts and practices of speakers in some sense under-determine what their conditions of correct application are. This might of course in the end be the correct view of vagueness, it is simply that it is incompatible with the (standard) epistemic conception of vague language whereby the thoughts and practices of speakers do fully determine the conditions under which vague terms are true or false—it is just that in borderline cases we are unable to tell whether or not these conditions obtain.

Perhaps, then, determinacy (or definiteness) is a *sui generis* notion? Field (1994, p. 111; 2001, p. 227) has offered the view that 'definitely' is a primitive expression whose meaning is to be grasped in the same way in which speakers might be said to grasp the meaning of the standard logical operators—via their introduction and elimination rules.[15] On this model, we can implicitly define what determinacy is by the operational rules for the 'definitely' operator, and then on that basis offer a constitutive definition of vagueness via some appropriate characteristic sentence. Field's conception looks to be of little use to the minimalist. A *sui generis* conception of definiteness precludes the possibility that 'definitely' or 'determinately' can be explicated in terms of such epistemic notions as knowledge, clarity, or knowability. Again, since we do not want our minimal theory to represent a proto-substantive theory of vagueness, one which disqualifies the epistemic conception from the beginning, then Field's conception of determinacy can form no part of the minimal theory.

Might there then be an *epistemic* reading of 'determinately'/'definitely' which is compatible with all conceptions of vagueness? Wright (1995, pp. 144–6) has suggested that 'determinately'/ 'definitely' might best receive a *quasi-epistemic* reading: roughly, when a statement *P* is determinately/definitely true then for a speaker *s* to judge that not-*P* means that *s*'s verdict is 'cognitively misbegotten'—the lighting might be bad, *s* might be drunk, tired, distracted, or forgetful. In borderline cases, cases where *P* is neither determinately true nor determinately false, Wright envisages that there can be 'faultlessly generated—cognitively un-misbegotten—conflict': subjects may permissibly disagree about the borderline cases, where the notion of permissible disagreement is, for Wright, of the very essence of vagueness (Wright 1987, p. 277, 1995, p. 138). Moreover, Wright urges that this interpretation is compatible with both standard epistemic and standard non-epistemic

[15] Though, Hyde does not explicitly say that the meaning of 'determinacy' is given by the introduction and elimination rules governing the 'determinately' operator, he does seem to sponsor a *sui generis* reading of 'determinacy'.

conceptions of what it is to be a borderline case.[16] While that may be so, Wright's quasi-epistemic reading of 'determinately' requires the insight that the thesis of permissible disagreement is of the very essence of vagueness. Yet the thesis of permissible disagreement has proved hard to stabilize (see Wright 2001, pp. 55–62 for the most relevant evaluation). Even if it is stable, it's not at all clear that one can stabilize it given uncontroversial resources. Consequently, Wright's quasi-epistemic reading can form no part of the minimal theory.

Might 'determinately' receive a more overtly epistemic reading? Williamson has famously argued that it is not just possible to give 'definitely' or 'determinately' some (overt) epistemic reading, it is the *only* illuminating and coherent reading we can give to these adverbs in the contexts of vagueness (Williamson 1994, pp. 194–5; and especially his 1995 paper). He suggests that 'definitely'/'determinately' may in effect mean something like 'knowably'. But such a reading is not available from the perspective of the minimal theory of vagueness because it rules out the intuitionistic conception of vagueness offered by Wright (2001) whereby if a borderline statements has a truth-value it is in principle to possible to find out (via some method or other) what that truth value is. A minimal epistemic reading of 'determinately'/'definitely' must not entail that the truth-values of vague sentences are verification-transcendent, for this is just to preclude Wright's conception from the start.

It now begins to look that that the only way in which the dialectic of the vagueness debate can begin at a mutually agreed point is to leave the requisite notion of determinacy or definiteness employed in some characteristic sentence such as **DT2** unspecified. Each substantive conception of vagueness would then sanction the minimal indeterminist conception only in so far as it is permitted to interpret 'determinately' according to considerations local to that conception. If this were so, then Wright's dictum that it is essential to have the expressive resources afforded by some notion of definiteness or determinacy begins to look

[16] To elaborate: in the epistemic case, suppose *P* is true, but unknowable in borderline cases, then a (non-inferential) verdict that not-*P* is *not* cognitively misbegotten, according to Wright, as there is no sense in which one can blame the speaker for their mistaken verdict when the truth of *P* is undetectable. On an indeterminist conception of vagueness where the adverb 'determinately' is non-epistemic and not strictly redundant (for example, on Field's view, but not Dummett's) an object *a* is borderline for a predicate *F* just in case *a* is neither determinately *F* nor determinately not-*F*. Thus, the thought goes it would be not be quite right and not be quite wrong to assert that *a* is *F* (likewise for an assertion that *a* is not *F*) for the matter is unsettled. And so, Wright envisages that subjects can be represented to permissibly differ in their verdicts on such a conception (or at the very least that such a conception is compatible with the thesis of permissible disagreement). It's worth noting that in his new approach to vagueness, Wright (2001) dispenses with the idea that permissible disagreement is of the very essence of vagueness.

of little use in specifying the sort of minimal theory we would ideally like to have. It is theoretically unsatisfactory to allow the minimal indeterminist conception to command universal assent simply in virtue of employing a multiply ambiguous conception of determinacy—for in that case we can hardly be said to have given a minimal definition of vagueness at all.

One (perhaps obvious) option remains: might we simply take the expression 'It is definitely/determinately the case that' to mean 'It is known that'? In so doing, for one thing, there is no implication that truth is potentially verification-transcendent. Arguably, every conception of vagueness entails that a vague declarative sentence is neither known to be true nor known to be false, that an object *a* is a borderline case for some predicate *F* when neither *F* nor not-*F* are known to be true of *a*. This suggestion has the immediate advantage that we can simply dispense altogether with the any notion of definiteness and determinacy in giving our minimal theory of vagueness, and in so doing, we can rid the minimal theory of the misleading non-epistemic overtones that the notions of determinacy or definiteness inevitably carry. It thus looks as though we must leave behind the characteristic sentence **DT2**, and instead endeavour to specify how a notion of knowledge can feature in a minimal characterization of vagueness.

6. Minimal vagueness *qua* borderline cases: the minimal epistemic conception

It is not too far wrong to say that there is an emergent consensus that the proper minimal theory of vagueness ought to be stated in epistemic rather than non-epistemic terms. In particular, that one can usefully employ a notion of ignorance to ground an uncontroversial definition of vagueness. (Though it has to be said that those who make this claim are not explicitly interested in developing a minimal theory of vagueness as such.) Sainsbury, for one, has urged that:

> All theorists can agree that a certain kind of ignorance is a sign of vagueness. We do not know whether or not some people are tall, not because we do not know how tall they are, but because we do not know whether being that tall counts as being tall ... We do not know whether we are still on Snowdon, not because we do not know where we are (we might know our map reference, or our precise distance from the summit) but because we do not know whether being here counts as being on Snowdon. Let us call cases which do or would give rise to such ignorance borderline cases (1995, p. 64).

Likewise, Williamson (1997, p. 921) agrees that vagueness 'is the phenomenon of borderline cases', and that we can at least *ostensively* define what it is to be a borderline case by giving examples:

> At some times, it was unclear whether Rembrandt was old. He was neither clearly old nor clearly not old. The unclarity resulted from vagueness in the statement that Rembrandt was old. We can even use such examples to define the notion of vagueness. An expression or concept is vague if and only if it can result in unclarity of the kind just exemplified. Such a definition does not intend to display the underlying nature of the phenomenon. In particular, it does not specify whether the unclarity results from the failure of the statement to be true or false, or simply from our inability to find out which. The definition is neutral on such points of theory. (1994, p. 2; see also p. 202)[17]

Any conception of vagueness which sanctions such a characterization we may call an *epistemic* minimal conception of vagueness *qua* borderline cases. Though it is tempting to read the above remarks as providing the materials for a rigorous minimal definition of vagueness, (one which supplies necessary and sufficient conditions for when an expression counts as vague), Sainsbury and Williamson both resist offering such a definition. Is such resistance justified? This will be the question which will preoccupy us in the following sections.

Suppose we now offer the following characteristic sentence:

(**K1**) $\exists\alpha$ In α, it is not known that *S* is true and it is not known that not-*S* is true

where we not only require that an acceptable substitution of *S* must be the sort of sentence whose truth is determined by the degree of variation in one or more graded or continuous underlying parameters $\nu_1, \ldots \nu_n$, but, crucially, that the source of ignorance must issue from features of the substituted sentence (or from features of the *use* of that sentence) and not from any ignorance as to the underlying ν-facts. This is to say that even when a speaker is apprised of all the relevant ν-facts, it will

[17] Tye (1995, p. 1) offers a similar neutral characterization. Keefe and Smith (1996, p. 2) follow suit, but then proceed to find it immediately plausible that our unclarity or ignorance in borderline cases is due to there being no fact of the matter to be clear about. It is noteworthy that Williamson (see his 1994, p. 16) assumes that, in the context of vagueness at least, that clarity and knowledge (and unclarity and ignorance) are freely interchangeable notions. This is to say, the operators 'it is clearly the case that' and 'it is known that' are fully interchangeable. There are various issues attending such an identification, not least that there is sense of 'clearly' in which an object *a* may be clearly *F*, despite the fact that *a* has never been seen by anyone. For this reason alone, it proves more convenient to develop the minimal theory of vagueness using the notion of knowledge rather than the notion of clarity.

still remain unknown whether or not the sentence *S* is true.[18] Is the constitutive definition using **K1** at all compelling?

In section 4 above, we saw that the possibility of such artificial terms as 'oldster' ruled out using **DT1** to constitutively define vagueness. Effectively the same worry rules out the use of **K1**. The term 'oldster' is neither known to be true nor known to be false of the intermediate cases (since it is neither determinately true nor determinately false of those cases), and since this ignorance issues from features of the sentence '*x* is an oldster' (and not from any ignorance regarding the underlying ν-facts), then it counts as vague given the characteristic sentence **K1**. Again, we should be strongly inclined to say that the term is not vague since it draws a perfectly sharp and clearly identifiable three-fold division across its associated dimension of comparison. Hence, a constitutive minimal definition invoking **K1** does not allow us to distinguish between vagueness and various distinct but superficially similar phenomena, such as semantic incompleteness.[19] But might one rehabilitate **K1** by appealing to some form of higher-order vagueness?

We saw above that one way in which we can incorporate a notion a notion of higher-order vagueness into **DT1** is to draw on Hyde's thesis that the notion of 'borderline case' is ambiguous between the type of borderline cases that stem from such terms as 'oldster' and genuine borderline cases where higher-order vagueness 'is built in from the very start'. But now we are dealing with an epistemic notion of a borderline case. Hyde's purported ambiguity in the notion of borderline case is 'ultimately as a result of the ambiguity of "determinately"', where Hyde

[18] For example, I may know the exact temperature of my bath water but still not know whether or not my bath is hot. However, it must be granted that in such a case, the underlying ν-facts are not strictly speaking non-vague, for, presumably, the exact spatio-temporal extension of my bath water is unknown (think of the clouds of steam rising from the surface of the water). Hence, it might be that I do not know whether or not my bath is hot simply because I do not know exactly what object the term 'my bath' refers to. To get round this difficulty, we can seek to define 'is vague' via a range of cases whereby the underlying ν-facts do not admit of vagueness. Take Wang's paradox, whereby we have a vague predicate 'is small' as applied to the natural numbers such that it is unclear whether or not certain natural numbers are small, relative to some relevant comparison class (see Dummett 1975). In such a case, the relevant ν-facts (just which natural number we are talking about) are non-vague: if it's unclear, say, whether or not 23 is a small number, this unclarity must be due to the vagueness of 'is small' since there is no unclarity as to which number is being referred to. The idea then is that we can define vagueness by analogy to such cases: suppose there is no ignorance as to the temperature of my bath (suppose that I know the spatio-temporal extension of my bath and I know that the watery object which is thus extended has a temperature between 23° and 26°) but suppose also that I do not know (after due consideration) whether or not my bath is hot (despite understanding the meaning of the expression 'my bath is hot'). Then, the idea is that, given **K1**, we have established that 'my bath is hot' is vague, and that it's vagueness is entirely due to the vagueness of 'is hot'.

[19] One suspects it may have been partly for this reason that Sainsbury, Williamson, and Tye, all resist giving a constitutive definition in terms of unclarity or ignorance.

appears to give this operator a non-epistemic reading (see Hyde 1994, p. 40). Is the expression 'it is known that' systematically ambiguous in a similar way? If it were then we could offer the following characteristic sentence:

(K2) $\exists\alpha$ In α, it is not **known** that *S* is true and it is not **known** that not-*S* is true.

The hope is that while the sentence '*x* is an oldster' would not satisfy **K2**, a sentence such as '*x* is red', for example, would satisfy **K2** (where *x* is on the red–orange borderline).

It is not obvious that 'known' is ambiguous in the manner in which 'determinately'/'definitely' might be, that somehow our grasp of what it is to be ignorant in borderline cases due to vagueness is such that higher-order vagueness is built in from the very start. Of course, it is open for one to *explicate* that '**known**' is to be understood in the requisite way, but to do so is problematic. How else can we fix the truth-conditions of **K2** without explicitly adverting to the fact that there are (epistemic) border cases, and (epistemic) border cases of those of border cases, and so on? In this case, our grasp of **knowledge** would be secondary to our grasp of higher-order vagueness. This is perhaps bad news for Hyde, but not necessarily bad news for the characteristic sentence approach. It is still open for us to make an explicit reference to (epistemic) higher-order vagueness when setting forth some appropriate characteristic sentence. But, as mentioned above, the existence of higher-order vagueness (*qua* borderline cases) is a matter of controversy (see Kamp 1981; Wright 1987, 1992b; Sainsbury 1990, 1991; Burgess 1990, 1998; Koons 1994). This ought to make us withhold (if only temporarily) from employing **K2** to ground a rigorous definition of vagueness.

This leaves us with two main options: follow Williamson and Sainsbury, and rest content with an ostensive minimal definition of vagueness *qua* (epistemic) borderline cases and thereby concede that the constitution of vagueness can only be identified from within some substantive conception (if at all), or, rehabilitate the characteristic sentence approach by appealing not to the notion of a borderline case in the first instance, but by reference to some other salient (and perhaps deeper) feature of vague expressions. Given the general difficulty of rigorously defining any philosophically interesting concept, perhaps the promises of the minimal theory have been overstated, and thus we should rest content with the first option. But such pessimism is arguably unjustified. It is theoretically unsatisfactory that our minimal theory of vagueness should just rest content with defining vagueness via exemplars of vague language. In the next section, I hope to show that we can indeed offer a rigorous definition by reference to a further feature of vague expressions, namely, the phenomenon of blurred boundaries.

7. Vagueness *qua* epistemic tolerance

For some, vagueness is the phenomenon of borderline cases, and there's an end on't. Fine, for instance, takes vagueness to be, in essence, a one-dimensional phenomenon. Bluntly: 'a predicate is extensionally vague if it has borderline cases' (1975, p. 266). We have already seen the shortcomings of such a characterization given the possibility of such terms as 'oldster'. What is surprising is that Fine makes no reference to a feature of vague expressions which is prima facie far more basic—namely, that such expressions draw no known boundary across their range of signification.[20] This feature reflects the basic phenomenological datum that along some smooth or graded dimension of comparison governed by some predicate *F*, subjects characteristically do not cognize any boundary between the *F*'s and the not-*F*'s (cf. Burgess 1998, p. 233). More loosely, we can say that vague terms have blurred boundaries. It is tempting to read this as substantiating the further thesis that vague expressions draw no *sharp* boundary across their range of signification. This tempting conclusion is not one which is available to the minimal theory of vagueness.[21] From the fact that speakers are unable to locate a sharp cut-off between, *red* and *not-red*, for instance, in no way entails that there is no sharp cut-off. The minimal conception must allow that perhaps, after all, there is such a sharp cut-off but we are simply unable to determine its whereabouts. This point is analogous to the one made above in section 4; namely, that we do not wish to saddle ourselves from the outset with a proto-theory of vagueness, one which excludes

[20] When reading Fine (1975) one can be left with the disconcerting impression that he is not really talking about vagueness at all. Hence, the pressing need for an articulation of the conceptual relationship between vagueness *qua* borderline cases and the prima facie more basic feature of vague expressions, namely, the phenomenon of blurred boundaries. See section 8.

[21] Keefe and Smith (1996, pp. 2–3), in their introductory characterization of vagueness offer the view that 'vague predicates … apparently lack well-defined extensions', but then add that on 'a scale of heights, there is *no sharp boundary* between the tall people and the rest' (see also Keefe 2000, Ch. 1). From the introductory perspective, this is highly misleading. One should take care that the expression 'sharp' does not carry any epistemic overtones such that no *sharp* boundary comes to mean something like no *known* boundary. Once 'sharp' is free of epistemic connotation, then the expressions 'no sharp boundary' and 'no boundary' should be thought of as equivalent—there is, strictly speaking, no such thing as an unsharp boundary (cf. Sainsbury 1990, 1991).

the epistemic conception of vagueness (and the concomitant commitment to sharp boundaries) from the outset.

Since all partisans to the dispute can or do, for whatever reason, accept that vague terms draw no known boundary across their associated dimension of comparison, it seems right that the minimal conception can and should countenance a further dimension to vagueness: there is vagueness *qua* sorites-susceptibility, vagueness *qua* borderline cases, and vagueness *qua* lack of known boundaries. In the next section, we shall look at the conceptual connections between the latter two dimensions, for present purposes we need to find a characteristic sentence which appropriately unpacks the claim that vague expressions have blurred or unknown boundaries. To give a more rigorous characterization we must dispense with talk of 'cut-offs' and 'boundaries' and offer instead a characterization in terms of a what I call *epistemic tolerance*. What is meant by tolerance and epistemic tolerance in this context?

In Wright 1975, (p. 334), Wright provisionally suggested that vague predicates are 'tolerant'.[22] Suppose we have a (monadic) predicate F which governs some dimension of comparison Φ, then according to Wright

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

Say that a predicate F fails to draw a (sharp) boundary when F is tolerant in the sense just given. Though the notion of tolerance is a key concept in the vagueness debate it cannot form the basis of a characteristic sentence. In section 3, we saw that we could not employ the characteristic sentences **SS1** and **SS2** (nor the induction step **A2** or the premiss **B1** from the *B*-sorites) to define vagueness. We likewise cannot employ

[22] Since the discussion of tolerance in the literature has focussed on predicate-vagueness, I shall for convenience here follow suit. It should be noted that Wright (in his 1975 and 1976 papers at least) takes vague terms to be tolerant if a particular view of the language capacity is adopted—a view he calls the 'governing view'. Very roughly, the governing view is the view that linguistic understanding simply consists in the mastery of a set of rules (both syntactic and semantic). Wright rejects this conception in favour of a view of language which is not rule-governed in this way and which must give priority to behavioural data.

Wright's notion of tolerance as the basis for a characteristic sentence for tolerant expressions are subject to the sorites paradox (at least in many logical systems).[23]

This predicament is puzzling since, as we have seen with **SS1** and **SS2**, the claim that vague predicates are tolerant is highly seductive. One response to this puzzle—the very puzzle of vagueness—is to oppose the thought that vague predicates are tolerant in Wright's sense. On an epistemicist view of vagueness, for instance, vague predicates are represented to be intolerant. For Sorensen (1988), as for other epistemicists, these predicates have 'unlimited sensitivity': there is some degree of change in respect of Φ which does make a difference as to whether F correctly applies or not to a particular case—it is just that we are (irremediably) ignorant as to which particular Φ-difference effects the change. Can we characterize vagueness in such a way as to be neutral as to whether or not vague predicates are tolerant?

Say that

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

Arguably, every vague predicate is tolerant in this respect. Roughly, epistemically tolerant predicates do not draw known boundaries across their range of signification. When we add the observation that large changes in respect of Φ will affect whether F applies or not then, loosely speaking, we can say that across Φ there is a difference without a known distinction. That is just what is (or ought to be) meant by the thesis that vague terms draw blurred boundaries. Given this, vagueness *qua* no known boundary and vagueness *qua* epistemic tolerance can be thought of as effectively equivalent. (Henceforth, I will generally employ the latter terminology.)

To draw no boundary is to draw no known boundary, but not vice versa. Likewise to draw no knowable boundary is to draw no known boundary, but not vice versa. Nor does the equation of vagueness and epistemic tolerance entail that any small change with respect to Φ will fail to make a known difference to whether or not F applies to a particular case. There is a world of difference between saying that no small degree of change makes a known difference as to whether x is F and say-

[23] It was noted above that A2 does not give rise to paradox in the subvaluational system of Hyde (1997), a system in which *modus ponens* fails. Likewise, if we represent vague terms to be tolerant, this does not give rise to paradox in this system. If one thinks that a substantial conception of vagueness has merit insofar as it respects our naive intuitions concerning vagueness, this feature ought to count as a distinctive merit of Hyde's view.

ing that small degrees of change make no difference as to whether x is known to be F. Some formalization might help here.

Suppose we have open sentence 'x is hot' governing the dimension of comparison of temperature where x ranges over temperatures in the interval 0° to 100°. Let 'K_s' abbreviate the functor 'It is known (by a speaker s that) that'. To say that no small positive degree of change in x makes a known difference as to whether x is hot can be given as: $\neg\exists x$ (K_s(x is hot)) $\wedge$ K_s($x \pm c$ is not hot), where c is some suitably small value. In contrast, to say that small positive degrees of change make no difference as to whether it is known that x is hot can be given as: $\neg\exists x$ (K_s(x is hot)) $\wedge$ $\neg K_s$($x \pm c$ is hot). The difference relies on the fact that 'K_s' and negation do not commute one way: $\neg K_s$(x is F) does not entail K_s(x is $\neg F$). To say that small positive degrees of change make no difference as to whether it is known that x is hot is simply an instance of the problematic schema **SS2** from which we can run a form of the B-sorites (see fn. 8 above). **SS2** is paradox-inducing while the equation of vagueness and epistemic tolerance is not.

Such observations suggest how we might offer two (classically equivalent) characteristic sentences which exploit the notion of epistemic tolerance. Let's also be completely explicit about this and build in all the provisos we have encountered hitherto (as well as some provisos we have not previously discussed). Our characteristic sentences are

(**ET1**) $\forall\alpha\ \forall\beta$ if $|\nu(\beta)-\nu(\alpha)|<c$ and K_s(S is true) in α then $\neg K_s$(not-S is true) in β

which is classically equivalent to:

(**ET2**) $\neg\ \exists\alpha\ \exists\beta$ such that $|\nu(\beta)-\nu(\alpha)|<c$, and K_s(S is true) in α and K_s(not-S is true) in β

where for both **ET1** and **ET2**:

(a) whether or not S is true depends on the value ν in actual or counterfactual cases

(b) c is some small positive real number

(c) small ν-values are cumulative (a series of small ν-values forms a large ν-value)

(d) all the relevant ν-facts are known, i.e. $\forall\alpha\ \nu(\alpha)$ is known (by the speaker s) in α

(e) the meaning of S is known (by the speaker s)

(f) we restrict the range of α and β to 'normal' cases of judgement conditions for the speaker *s*

(g) where *c* is large then $\exists\alpha\ \exists\beta$ if $|\nu(\beta)-\nu(\alpha)|>c$ then *S* is known (by the speaker *s*) to be true/false in α and not-*S* is known (by the speaker *s*) to be true/false in β.

A sentence *S* is vague just in case when substituted in to either **ET1** and **ET2** these schemas are true—at least when all of the clauses (a)–(g) are satisfied. This may look rather cumbersome, but without each of these clauses we will not be able to properly distinguish the vague from the non-vague.

Clause (a) we have encountered already. A more sophisticated formulation of **ET1** and **ET2** would advert to the fact that whether or not a vague sentence is true may depend on the variation in more than one underlying graded or continuous parameters. This is compatible with the vagueness of *S* issuing exclusively from the subject term or terms contained in *S*. Clause (b) allows the underlying ν-facts to vary continuously or discretely. Clause (c) is surely unproblematic. Clauses (d)–(f) are crucial: they each ensure that a speaker's ignorance does not result from the wrong source but solely from the vagueness of the sentence *S*. Clause (d) ensures that all the relevant ν-facts are known by the speaker *s*. In the case of assessing whether the sentence 'My bath is hot' is vague, it is an prerequisite that I know the temperature of my bath (in actual and counterfactual cases). Likewise, nor must my ignorance issue from any misunderstanding or ignorance as to the meaning of the sentence *S*. Clause (e) ensures just that. Clause (f) is also vital for a speaker may know all the relevant ν-facts and know the meaning of *S* and yet the speaker may be ignorant for reasons other than vagueness. For example, the speaker may be drunk tired, distracted, or hallucinating such that **ET1** or **ET2** is satisfied. Indeed, external conditions may produce ignorance—the lighting might be bad, there might be smoke in the room, and so on. It may be a delicate matter to give a general characterization of normal judgement conditions, and I will not endeavour to do so here beyond the remarks already given. Lastly, clause (g) ensures that large differences in the value taken by ν in α and the value taken by ν in β will always entail a known difference in the truth-value of S/not-S. Any conception of vagueness which does or can constitutively define vagueness via **ET1** (or **ET2**) together with clauses (a)–(g) we may call a conception of vagueness *qua* epistemic tolerance. But can **ET1** and **ET2** avoid the problem of 'oldster' and cognate problems?

For simplicity I shall focus on **ET2**. As stated, **ET2** is in fact insufficiently general to ensure that the artificial term 'oldster' fails to count as vague. Why? Take the open sentence 'a person of age x is an oldster'. Given **ET2**, it is sufficient for this sentence to be vague that there is no x such that it is known that this sentence is true *and* known that the sentence 'a person of age x–c is *not* a oldster' is likewise true (given all the other provisos, and where c takes some small value). When the values of both x and x–c fall in the stipulated penumbral area (that is, when $65 \leq x \leq 68$ and $65 \leq x-c \leq 68$) then both sentences are neither determinately true nor determinately false (according to the dictates of the stipulation), and so both are neither known to be true nor known to be false on just that basis (where we assume knowledge requires determinate truth). Hence, for those values the schema is satisfied. When the values of both x and $x-c$ fall outside the stipulated penumbral area, then one (and only one) of the sentences is false, and so for these values the schema is satisfied. It is only when $x-c < 65 \leq x$, or when $x-c \leq 68 < x$ that we might—at first blush—expect these sentences to be known to be true.

Take the 'higher' of the two cut-offs. The statement 'a person of age x is an oldster' is known to be true, but the statement 'a person of age $x-c$ is *not* a oldster' is not known to be true since it is penumbral—it is not determinately true (and not determinately false). Hence, the schema **ET2** is satisfied for the higher cut-off. Take the 'lower' cut-off. The statement 'a person of age x–c is *not* a oldster' is known to be true, but the statement 'a person of age x is an oldster' is not known to be true since it is penumbral. Hence, the schema is satisfied for the lower cut-off. Consequently, there are no two neighbouring values across the dimension of comparison of age for which the schema **ET2** fails: 'oldster' satisfies **ET2** together with (a)–(g), and thus counts as vague. But we have seen that 'x is an oldster' is intuitively not vague: it draws a perfectly sharp and clearly identifiable three-fold division across its range of significance.

To meet this worry one might simply seek to offer the following schemas:

(**ET3**) $\forall\alpha\ \forall\beta$ if $|\nu(\beta)-\nu(\alpha)|<c$ and K_s(S is determinately true) in α then $\neg K_s$(S is not determinately true) in β

which is classically equivalent to:

(**ET4**) $\neg\exists\alpha\ \exists\beta$ such that $|\nu(\beta)-\nu(\alpha)|<c$ and K_s(S is determinately true) in α and K_s(S is not determinately true) in β

(where all the other clauses (a)–(g) must be met in order for S to count as vague).[24] The characteristic sentence **ET4** and **ET3** correctly identify that the sentence 'a person of age x is an oldster' to be non-vague. In more detail, when $x-c \leq 68 < x$ then 'a person of age $x-c$ is an oldster' is *not* determinately true, and it is, moreover, known that this is so, while the statement 'a person of age x is an oldster' is determinately true, and it is, moreover, known that this is so. Thus, for these values **ET4** is not satisfied.

But while **ET3** and **ET4** correctly predict that 'oldster' is non-vague, they are not able to identify the non-vagueness of any term which is stipulated to admit of second-order borderline cases. To illustrate: suppose we offer the following stipulation for a new term 'oldster*':

(i) If $x > 70$ then 'x is an oldster*' is determinately determinately true,

(ii) If $70 \geq x > 67$ then 'x is an oldster*' is neither determinately determinately true nor determinately not determinately true,

(iii) If $67 \geq x \geq 66$ then 'x is an oldster*' is determinately not determinately true and determinately not determinately not true,

(iv) If $66 > x \geq 63$ then 'x is an oldster*' is neither determinately determinately not true nor determinately not determinately not true,

(v) If $63 > x$ then 'x is an oldster*' is determinately determinately not true.

Again, there are no two neighbouring values on the dimension of comparison for which clause **ET4** (or **ET3**) fails. The point generalizes. If we modify **ET3** and **ET4** to cope with this second-order artificial stipulation a third-order counter-example (with nine mutually exclusive and exhaustive truth-states) can be gerrymandered. In the limiting case, it is presumably possible to stipulate that some term is radically higher-

[24] In specifying **ET3** and **ET4**, I have of course employed a non-epistemic 'determinately' operator. There might of course be no such notion, in which case one would simply not be able to stipulate the meaning of 'oldster' as was done above and **ET1** and **ET2** would be entirely satisfactory as they stand. (Indeed one may have very general worries about the possibility of such stipulations—see Williamson 1997.) However, our minimal theory must be free from controversy and must be generous enough to allow that there may well be such a non-epistemic notion of determinacy/definiteness which permits us to stipulate such terms as 'oldster'. Any epistemicist who rejects the idea that there is any workable non-epistemic notion of determinacy may simply replace 'determinately true' with 'true' in the characteristic sentences **ET3** and **ET4** in order to employ these sentences to adequately define vagueness.

order vague—that the borderline cases to borderline cases are non-terminating. Is there a response?

The first thing to note is that counterexamples to **ET1–4** can only be generated within some suitable logical framework. For instance, the stipulation governing 'oldster*' requires a logic for the predicate 'determinately true' in which the following schemas are invalid:

(**DD**) If *S* is determinately true then *S* is determinately determinately true

(**D¬D**) If *S* is not determinately true then *S* is determinately not determinately true

(these are just formal mode analogues of the S4 and S5 principles, where 'determinately' takes the place of 'necessarily'). In a modal system which lacks these schemas (such as KT or KTB) there are no so-called reduction laws; hence there are an infinity of distinct modalities. Whatever logical framework is exploited in order to gerrymander some borderline term which nonetheless has sharp boundaries, we can define a variable τ which ranges over all possible modalities or 'truth-states' in this logical system. In KT and KTB, τ ranges over: {true, determinately true, not true, not determinately true, determinately determinately true, determinately not determinately true, ...}[25] In systems which contain reduction laws, τ will range over a finite number of modalities.[26] Given this, we can offer the following replacements for **ET3**, **ET4**:

(**ET5**) $\forall\tau\ \forall\alpha\ \forall\beta$ if $|\nu(\beta)-\nu(\alpha)|<c$ and $K_s(S \text{ is } \tau)$ in α then $\neg K_s(S \text{ is not-}\tau)$ in β

(**ET6**) $\forall\tau\ \neg\exists\alpha\ \exists\beta$ such that $|\nu(\beta)-\nu(\alpha)|<c$ and $K_s(S \text{ is } \tau)$ in α and $K_s(S \text{ is not-}\tau)$ in β.

Schema **ET6** effectively says that there are no close cases in which it is known that a sentence takes a certain truth-state in one case and known that this sentence takes the complementary truth-state in the other close case. Schema **ET5** effectively says that if it is known that a sentence takes a certain truth-state then in nearby cases it is not known that this sentence lacks this truth-state. (Given classical logic, both schemas are inter-derivable.)

[25] Note that the minimal theory does not entail that there are such (non-reducible) truth-states. The idea here is to combat the possibility that if there are such states then vagueness cannot be constitutively defined.

[26] In the modal system KT5 there are four reduction laws (see Chellas 1980, pp. 147–54).

The idea is that however far up the hierarchy of truth-states one goes in order to gerrymander some counterexample to **ET5** and **ET6**, one's stipulation will always invoke a sharp and clearly identifiable cut-off between at least one truth state τ and its complementary truth-state not-τ. The point also holds for any term which might be stipulated to be radically higher-order vague in the relevant sense. Thus such sentences as '*x* is an oldster', '*x* is an oldster*', and higher-order analogues, are correctly identified as non-vague. We have, at last, distinguished vagueness from superficially similar phenomena such as semantic incompleteness or underspecificity.[27] Sentences which satisfy **ET5** and **ET6**, together with clauses (a)–(g), are vague, and conversely. Vagueness, from the perspective of the minimal theory is the phenomenon of epistemic tolerance: *S* is vague just in case there is no small change in respect of Φ (actual or counterfactual) which makes a known difference as to whether or not *S* is τ (where τ ranges over whatever truth-states *S* could possibly take over Φ).

We should thus be relatively satisfied that the characteristic sentence approach can after all offer a reasonably rigorous definition of vagueness which distinguishes the vague from the non-vague and which is acceptable to all parties to the vagueness debate. Moreover we have succeeded in giving a definition of sentential vagueness which does not make an explicit reference to higher-order vagueness but which is rich enough to predict that terms (like 'oldster') which are stipulated to be have borderline cases are non-vague. It looks as if there is no need to offer a substantive semantic story in order to constitutively define vagueness as Sainsbury (1991, p. 174) has surmised.[28] The reluctance of Williamson and Sainsbury, (and many others) to offer anything more than an ostensive definition of vagueness now seems overly modest. What, then, are the conceptual and explanatory relationships between vagueness *qua* epistemic tolerance and vagueness *qua* borderline cases?

[27] Some commentators (for example, Channell 1994, p. 2, *passim*) confuse vagueness *qua* tolerance with, what we may term, vagueness *qua* generality or 'underspecificity'. The predicate 'is between five and six hundred miles' is underspecific (in certain contexts it does not carry enough information), but it is not an example of vagueness proper in that it draws clear boundaries over its range of signification. Our characteristic sentences allow us to distinguish between these distinct species of vagueness.

[28] Indeed, Sainsbury in his 1990 and 1991 works, takes the hallmark of vagueness to be *boundarylessness*—a feature which surely entails that vague terms draw no *given* or clear boundary across their associated dimension and which ought therefore sanction the veracity of **ET5** and **ET6**.

8. Which dimension is more basic?

At the end of section 3, it was suggested that vagueness *qua*-sorites susceptibility is secondary in the explanatory order—sentences are sorites-susceptible because they are vague, and not vice versa. At the end of section 6, it was suggested that vagueness *qua* epistemic tolerance is more basic than vagueness *qua* borderline cases. Epistemic tolerance seems more basic than the epistemic notion of a borderline case (both explanatorily and conceptually): sentences give rise to borderline cases *because* they are epistemically tolerant and being epistemically tolerant is constitutive of vagueness in a way in which being a borderline case is not. It thus looks as though we have already established how the three dimensions of vagueness are related to each other (both conceptually and explanatorily). But this is too quick. While it's fair to say that vagueness *qua* sorites-susceptibility is the less basic dimension of vagueness, in this section we will assess whether vagueness *qua* epistemic tolerance and vagueness *qua* borderline cases are in fact conceptually equivalent dimensions despite the possibility of such terms as 'oldster'.

Generally speaking, most philosophers have been somewhat cavalier about what conceptual or explanatory relationships hold between the dimensions of vagueness *qua* borderline cases and vagueness *qua* tolerance. It is worth quoting in full what has been said about this matter. Black—who in general tends to characterize vagueness in terms of the former dimension—is an early exception; he says:

> The finite area of the field of application of the word is a sign of its *generality*, while its vagueness is indicated by the finite [borderline] area and lack of specification of its boundary. It is because *small* variations in character are unimportant … that it is possible, by successive small variations, in any respect, ultimately to produce 'borderline cases'. (1937, p. 31)

For Black, then, vagueness *qua* tolerance is explanatorily (and presumably conceptually) prior to vagueness *qua* borderline cases.[29] Sainsbury, in speaking on behalf of what he calls the classical conception—the conception which 'characterizes vagueness in terms of its allowing for borderline cases' (p. 179)—also takes borderline cases to result from tolerance. He says:

[29] Black seems to have in mind something like Wright's (semantic) notion of tolerance, when he should really have adverted to our notion of epistemic tolerance and said that it is because there is no small variation in character that is *known* to be important to the applicability of *F* that it is possible, by successive small variations, in any respect, ultimately to produce 'borderline cases' for *F*.

> a very small difference in shade cannot make the difference between something being green and being blue, so we need a class of borderlines; a very small difference in age cannot make the difference between childhood and adulthood, so we need a class of borderlines. (1991, p. 168)[30]

While Hyde in discussing what he calls the 'paradigmatic conception'—the conception which aims to characterize vagueness in terms of borderline cases and borderline cases to those borderline cases—says:

> The paradigmatic concept we have been discussing initially attempts to accommodate the intuition that there is no apparent sharp boundary between the positive and negative extension of a predicate in terms of the presence of a penumbra or border region (or border cases). So, for example with the predicate 'red' the absence of any apparent sharp boundary between the red and the non-red is initially described by reference to borderline cases. (1994, p. 36)[31]

In this case, Hyde seems to think that on this conception the notion of borderline case is more basic.[32] In contrast to Black, and the conceptions outlined by Sainsbury and Hyde, one might assume that the two dimensions are conceptually equivalent and such that there is no proper explanatory priority to either dimension—we can equally well explain what vagueness amounts to by reference to either facet. This seems to be the view of Keefe and Smith:

> Clearly having fuzzy boundaries is closely related to having borderline cases. It might be argued, for example, that for there to be no sharp boundary between the *F*'s and the not-*F*'s just *is* for there to be a region of borderline cases of *F*. Our 'two features' would then be thought of as the same central feature of vague predicates seen from two different slants. (Keefe and Smith 1996, p. 3, fn. 3, see also Keefe 2000, p. 7)[33]

[30] In this paper Sainsbury aims to show that the classical conception generates an implausible model of higher-order vagueness. In his later paper of 1995, as we have seen above, Sainsbury is happy to define vagueness ostensively via reference to borderline cases, though the focus of this paper lies elsewhere and there is no discussion of higher-order vagueness in this later paper.

[31] What Hyde means by the 'paradigmatic conception' in many ways coincides with what Sainsbury means by the 'classical conception'. Like Sainsbury (in his 1990 and 1991 papers at least), Hyde is intent on undermining the grip which this conception has had upon the vagueness debate.

[32] There are other passages in Hyde (1994) which suggest that vagueness *qua* epistemic tolerance is more basic.

[33] Keefe and Smith fall foul of the confusion between 'no sharp boundary' (tolerance) and 'fuzzy boundary' (which they ought to have read as the feature of epistemic tolerance). Just as with Black, it is easy to adjust their comments and employ an epistemic notion of tolerance and borderline case.

Consider, then, the following two-part claim:

> (**ET** $\Rightarrow$ **BC**) *From epistemic tolerance to borderline cases:* if a term is epistemically tolerant then this will entail that it will fail to draw a known boundary across its range of signification and so there will be cases such that it is not known whether or not this term applies.
>
> (**BC** $\Rightarrow$ **ET**) *From borderline cases to epistemic tolerance:* if it is not known whether or not a vague term applies to certain cases then this will entail that it is not known where the boundary lies between such cases, and in the absence of a known boundary no small change in the world will make a known difference as to whether or not the term applies.

Both Sainsbury's classicist and Black endorse **ET** $\Rightarrow$ **BC**, while Hyde's advocate of the paradigmatic conception advocates **BC** $\Rightarrow$ **ET**, and Keefe and Smith (in the quote given at least) endorse both claims. How can we adjudicate?

On an intuitive level at least, **ET** $\Rightarrow$ **BC** seems absolutely right. It is the route from (epistemic) borderline cases to (epistemic) tolerance that is suspect. A term (and its complement) can fail to apply (and so fail to be known to apply) to certain cases and yet nonetheless fail to count as epistemically tolerant—that was surely the lesson of the term 'oldster'.[34] So it looks as though vagueness *qua* epistemic tolerance is more basic. But is this right?

Firstly, let's vindicate the intuition that vagueness *qua* epistemic tolerance entails vagueness *qua* borderline cases. Suppose some open sentence '*Fx*' is epistemically tolerant such that it satisfies the schema **ET6** and clauses (a)–(g). Let '*Fx*' abbreviates the sentence 'a person of *x* years of age is old' where '*F*', let us say, ranges over the series of natural numbers from 0 to 120. The relevant epistemic tolerance of '*Fx*' is such that no small drop in age makes a known difference to whether or not '*Fx*' is τ. As before, let 'K_s' abbreviate the functor 'it is known by a speaker *s* that', and let τ range over all possible truth-states which the sentence 'a person of *x* years of age is old' can possibly take over the dimension of comparison. Given the epistemic tolerance of '*Fx*' then we have

1 (1) $\neg\exists x\, K_s(\text{`}Fx\text{' is } \tau) \wedge K_s(\text{`}Fx{-}1\text{' is not } \tau)$.

[34] Indeed Keefe and Smith, in the same footnote quoted above, go on to recognize the possibility of terms (like 'oldster') which admit of borderline cases but where the borderline is sharply bounded. On p. 15 they add that 'merely having borderline cases is not sufficient for vagueness: rather, with a genuinely vague predicate, the sets of clearly positive, clearly negative and borderline cases will each be fuzzily bounded'.

Since F also satisfies clause (g), it follows that there is some large change in respect of Φ will make a known difference as to whether or not '*Fx*' is τ, which we may conveniently give as follows:

2 (2) $\exists x\, K_s(\text{'}Fx\text{' is } \tau) \wedge K_s(\text{'}Fx{-}10\text{' is not } \tau)$.

And for the sake of *reductio*, let us suppose that for all x, either it is known that '*Fx*' is τ or known that '*Fx*' is not τ:

3 (3) $\forall x\, K_s(\text{'}Fx\text{' is } \tau) \vee K_s(\text{'}Fx\text{' is not } \tau)$.

We now need to assume that for an arbitrary x that:

4 (4) $K_s(\text{'}Fx\text{' is } \tau) \wedge K_s(\text{'}Fx{-}10\text{' is not } \tau)$

(that is, line 4 is the typical disjunct for line 2). Given $\wedge$-E we can thus infer:

4 (5) $K_s(\text{'}Fx\text{' is } \tau)$

4 (6) $K_s(\text{'}Fx{-}10\text{' is not } \tau)$.

Now let us assume for the sake of absurdity that

7 (7) $K_s(\text{'}Fx{-}9\text{' is } \tau)$.

By $\wedge$-I on lines 6 and 7 we infer:

4,7 (8) $K_s(\text{'}Fx{-}9\text{' is } \tau) \wedge K_s(\text{'}Fx{-}10\text{' is not } \tau)$.

Given $\exists$-I this yields:

4,7 (9) $\exists x\, K_s(\text{'}Fx\text{' is } \tau) \wedge K_s(\text{'}Fx{-}1\text{' is not } \tau)$.

which contradicts line 1 and so we thus infer:

1,4 (10) $\neg K_s(\text{'}Fx{-}9\text{' is } \tau)$.

Now an instance of line 3 is:

3 (11) $K_s(\text{'}Fx{-}9\text{' is } \tau) \vee K_s(\text{'}Fx{-}9\text{' is not } \tau)$

and by disjunctive syllogism on lines 11 and 12 allows us to infer:

1,3,4 (12) $K_s(\text{'}Fx{-}9\text{' is not } \tau)$.

If we then assume for the sake of absurdity that $K_s(\text{'}Fx{-}8\text{' is } \tau)$ by the same pattern of inference we can derive that $K_s(\text{'}Fx{-}8\text{' is not } \tau)$. If we repeat this pattern of inference ten times, then we can thus derive:

1,3,4 (13) $K_s(\text{'}Fx\text{' is not } \tau)$

and given the factivity of the functor 'it is known by a speaker s that' this entails

1,3,4 (14) 'Fx' is not τ.

Given factivity, line 5 entails

4 (15) 'Fx' is τ.

Contradiction. So reject 3 to infer:

1,4 (16) $\neg\forall x\, K_s(\text{'}Fx\text{' is } \tau) \vee K_s(\text{'}Fx\text{' is not } \tau)$

which given the quantifier equivalences entails:

1,4 (17) $\exists x\, \neg(K_s(\text{'}Fx\text{' is } \tau) \vee K_s(\text{'}Fx\text{' is not } \tau))$.

Which via de Morgan gives:

1,4 (18) $\exists x\, \neg K_s(\text{'}Fx\text{' is } \tau) \wedge \neg K_s(\text{'}Fx\text{' is not } \tau)$

and by $\exists$-E on lines 2,4 and 18 this gives:

1,2 (19) $\exists x\, \neg K_s(\text{'}Fx\text{' is } \tau) \wedge \neg K_s(\text{'}Fx\text{' is not } \tau)$.

Result: **ET** $\Rightarrow$ **BC**. From the fact no small Φ-change makes a known difference as to whether or not x is F and the fact that some large Φ-change does make a known difference as to whether or not x is F, shows that 'Fx' gives rise to borderline cases *at least given the resources of classical logic*. This qualification is important. In the proof above, we have used principles (negation-introduction, the de Morgan's laws, disjunctive syllogism, and so on) which have all been brought into doubt on certain approaches to the sorites paradox.[35] However, recall that in section 2 it was argued that the minimal theory of vagueness is entitled to exploit classical resources until such point as this generates tangible controversy. Would any theorist (classical or otherwise) seriously seek to doubt the entailment from epistemic tolerance to epistemic borderline cases? Just because one rejects certain classical principles in order to combat the sorites does not entail that those principles fail throughout one's theory of vagueness. It seems right to say that the proof above is available to all partisans.

[35] It's perhaps worth mentioning at this point that the step from line 16 to 17 is intuitionistically invalid. This is significant. The intuitionist can say that a vague term F draws no known boundary across its Φ-dimension without saying that there is an object a for which it is not known that a is F and not known that a is *not-F*. For the intuitionist at least, vagueness may well *not* be the phenomenon of borderline cases, a consequence which may well be important for what account the intuitionist gives of higher-order vagueness.

Of perhaps more immediate interest is that fact that it looks as though line 19 can be generalized so as to form the basis of characteristic sentence for vagueness *qua* borderline cases as follows:

(**K3**) $\forall\tau\ \exists\alpha$ In α, $\neg K_s(S \text{ is } \tau)$ and $\neg K_s(S \text{ is not } \tau)$.

This schema effectively says that for any truth-state τ that S may take over Φ, there is at least one case such that it is not known whether or not S is τ. (What is notable about **K3** in contrast to our earlier formulation **K2** is that there is no (implicit) reference to higher-order vagueness.) The key question now is: does vagueness *qua* borderline cases as codified in **K3** entail vagueness *qua* epistemic tolerance?

Suppose that our sentence '*Fx*' satisfies **K3**; thus for a particular number *m* it follows that:

1 (1) $\neg K_s(\text{'}Fm\text{' is } \tau) \wedge \neg K_s(\text{'}Fm\text{' is not } \tau)$

and suppose for the sake of *reductio* that '*Fx*' is intolerant in the relevant respect, that is:

2 (2) $\exists x\, K_s(\text{'}Fx\text{' is } \tau) \wedge K_s(\text{'}Fx{-}1\text{' is not } \tau)$

from line (1) by $\wedge$-E we get:

1 (3) $\neg K_s(\text{'}Fm\text{' is } \tau)$

1 (4) $\neg K_s(\text{'}Fm\text{' is not } \tau)$

Now let us assume for arbitrary *x* that:

5 (5) $K_s(\text{'}Fx\text{' is } \tau) \wedge K_s(\text{'}Fx{-}1\text{' is not } \tau)$

(that is, line 5 is the typical disjunct for line 2). Given $\wedge$-E, this yields:

5 (6) $K_s(\text{'}Fx\text{' is } \tau)$

5 (7) $K_s(\text{'}Fx{-}1\text{' is not } \tau)$.

We know a priori that:

(8) $x{=}m$ or $x{>}m$ or $x{<}m$.

If we suppose that $x{=}m$ then we can immediately derive a contradiction on lines 3 and 6; so instead just suppose that

9 (9) $x > m$.

But since it is plain that if it is known that *y* is not old then it is known that *y*−1 is not old, which is to say:

10 (10) $\forall y\, K_s(\text{‘}Fy\text{’ is not } \tau) \rightarrow K_s(\text{‘}Fy{-}1\text{’ is not } \tau)$.

So, given lines 7, 9 and 10 we can prove given successive applications of universal instantiation and *modus ponens* that:

5, 9, 10 (11) $K_s(\text{‘}m\text{’ is not } \tau)$

which contradicts line 4 so we have:

1, 5, 9, 10 (12) $\perp$.

So suppose that:

13 (13) $x < m$.

But since we know that if it is known that *y* is old then it is known that *y*+1 is also old, and so we have:

14 (14) $\forall y\, K_s(\text{‘}Fy\text{’ is } \tau) \rightarrow K_s(\text{‘}Fy{+}1\text{’ is } \tau)$.

So given lines 6, 13, and 14 we can infer:

5, 13, 14 (15) $K_s(\text{‘}Fm\text{’ is } \tau)$

which contradicts line 3 and so we have:

1, 5, 13, 14 (16) $\perp$.

and so by $\vee$-E on 8, 9, 12, 13, 16 we have:

1, 5, 10, 14 (17) $\perp$

and by $\exists$-E on 2, 5, 17 we get:

1, 2, 10, 14 (18) $\perp$

and by ¬-introduction on 2, 18 we conclude:

1, 10, 14 (19) $\neg\exists x\, K_s(\text{‘}Fx\text{’ is } \tau) \wedge K_s(\text{‘}Fx{-}1\text{’ is not } \tau)$

Result: **BC** $\Rightarrow$ **ET**. Thus we have shown that vagueness *qua* borderline cases (in the guise of **K3**) entails vagueness *qua* minimal tolerance (on condition that the above rules of inference are valid in this context, of course). What conclusions can we draw?

It now looks as though once we have properly isolated the phenomenon of vagueness *qua* borderline cases via the characteristic sentence **K3**, then neither dimension is conceptually more basic: vagueness *qua* epistemic tolerance and vagueness *qua* borderline cases are indeed two facets of the same phenomenon.[36] Our exposure to the problem of such terms as 'oldster' was in many ways a red-herring. Once we have isolated characteristic sentences which predicted the non-vagueness of such terms then any temptation to think that vagueness *qua* epistemic tolerance is the more basic phenomenon is lost. So while Fine (1975) may not have explicitly alluded to vagueness *qua* epistemic tolerance in setting forth his (indeterminist) account of vagueness *qua* borderline cases, it seems he (together with Williamson and Sainsbury) was nonetheless endeavouring to investigate the same phenomenon as Hyde, Burgess, Wright, and all those who have tended to focus on the phenomenon of blurred boundaries. Our partisans have (more or less) been talking about the same thing all along. Our minimal theory can thus not only constitutively define vagueness it can also both ensure that the dialectic of the vagueness debate can begin from a mutually agreeable point. In fact it can do something more than this, it can ensure that the vagueness debate does not start off on the wrong foot. To this issue I will now briefly turn.

9. A level playing-field

At the end of section 3, the question was raised as to what deeper feature of vague expressions might incline us to accept the (paradoxical) schemas **SS1** and **SS2**. Another way of putting that question runs: just why do we find the induction step of the sorites paradox so seductive—just why, hitherto, have we been so seduced by the sorites paradox? The answer ought now to be clear: the confusion of tolerance with epistemic tolerance, the confusion of the claim that vague terms lack sharp boundaries with the claim that vague terms have blurred or unknown boundaries. This confusion is easily made if one gives an epistemic reading to 'sharp' (see fn. 21 above). Arguably, our naïve intuitions concerning vagueness are not sophisticated enough to make the distinction between tolerance and epistemic tolerance, between lacking sharp boundaries and lacking known boundaries (having blurred boundaries). Since our naïve intuitions generate paradox, once we grasp that tolerance, unlike epistemic tolerance, has *nothing* whatsoever to do with the constitution of vagueness (at least as we experience this phenomenon, namely the experience of blurred boundaries) then we can

[36] Recall that 'oldster' in fact satisfied the simple characteristic sentence **ET2**, prompting us to formulate **ET6** in its stead. It is for this reason that one can establish a parallel interderivability result between **ET2** and **K1**. But once we have excluded the possibility of such terms as 'oldster' then the simple-minded rationale given above for equating these dimensions ought to be valid.

discard our naïve intuitions without fear of recrimination. Nothing is lost in severing the connection between vagueness and tolerance. Moreover, something is gained: one can reject the induction step of the *A*-sorites with impunity. Of course, most substantive theories of vagueness do indeed reject the induction step, but of those which do, the majority do so with a certain reluctance (as if they were betraying something about the nature of vagueness in doing so). Typically such theories reject this premiss but feel an obligation to mitigate the impact of this rejection by adopting some non-classical logic. But if this premiss has nothing to do with the phenomenon of vagueness (as we experience it at least) then why put oneself under such an obligation from the outset?

This again shows why the minimal theory is dialectically valuable. There is a real sense in which the vagueness debate has in many ways got off to the wrong start. There ought to be no reason for any theorist of vagueness to *feel guilty* about rejecting tolerance intuitions (at least if one keeps one's epistemic tolerance intuitions intact). Thus there is no anterior reason to think that vagueness requires some non-classical formal system which allows one to mitigate the impact of rejecting the induction step of the sorites. Of itself, this is not an argument in favour of the standard epistemic approach to vagueness (whereby the induction step is rejected *guilt free*, as it were, and classical logic and semantics are retained), for one might find reasons to lessen the impact of rejecting the induction step which are not motivated in any way by tolerance intuitions. It is rather that we have removed an unhelpful and long-standing prejudice against finding the right view of vagueness. The substantive vagueness debate can now proceed from a more level playing field. Indeed it is clear what the key explanatory burden of any substantive approach now amounts to: find some way of defusing the paradox which does justice to epistemic tolerance intuitions. It is also clear what the key explanatory burden of any non-epistemic approach to vagueness (in which the induction step is rejected) amounts to: find some way of rejecting the induction step whereby any moves to mitigate the impact of such a rejection must *in no way* be motivated by tolerance intuitions.

There is also a second important, and related, reason why the minimal theory as given is dialectically valuable: it counsels us to be more cautious when identifying the essence of vagueness from within some substantive conception. Sainsbury (1991), for instance, takes the essence of vagueness to be *boundarylessness*. He says

> To convince you that boundaryless classification is possible, I would ask you to think of the colour spectrum. It contains bands but no boundaries. The different colours stand out clearly, as distinct and exclusive, yet close inspection shows that there is no boundary between them. The spectrum provides a paradigm of classification, yet it is boundaryless. (p. 180)

But arguably the phenomenological data merely supports a thesis of epistemic tolerance and not a thesis of boundarylessness: close inspection simply shows that there is no clear or known boundary between the bands, not that there is no boundary. Of course it might in the end be that boundarylessness (best) explains why terms are epistemically tolerant; it's simply that the phenomenological data *per se* is no evidence for boundarylessness. The constitution of vagueness (as we experience it at least) is exhausted by epistemic tolerance. What, then, does this minimal theory have to say about higher-order vagueness? Can it tell us that if there is first-order vagueness then there is *n*th-order vagueness at every order *n*?

10. Must there be higher-order vagueness?

Russell (1923) and Black (1939) can be credited with putting the topic of vagueness back on the philosophical agenda. But notably these two philosophers disagreed about whether or not there is higher-order vagueness (*qua* borderline cases). While Russell says that the 'penumbra itself is not accurately definable' (1923, p. 86), Black says it is 'impossible to accept Russell's suggestion that the fringe itself is ill-defined' (1939, p. 37).[37] Though, most commentators have sided with Russell (though typically from within different conceptions of vagueness), some have sided with Black (though not always for quite the same reasons). Wright (1987, 1992b), for instance, has posed the challenge that independently of any worries concerning the sorites paradox, the notion of higher-order vagueness is in itself paradoxical. Kamp (1981) and Sainsbury (1990, 1991) meanwhile have raised doubts concerning any model of higher-order vagueness couched in set-theoretical terms. Koons (1994), in the context of defending a particular hybrid conception of vagueness, has said that the

> desire to avoid sharp boundaries is no reason to postulate higher-order vagueness. I will not postulate such second- or higher-order vagueness un-

[37] Black's doubts stem from the thought that classical negation rules out the possibility of an indeterminist conception of borderline cases. These doubts relate to those raised by Williamson (1992, 1994, Ch. 7) concerning the stability of truth-value gaps.

> less some independent argument can be made for doing so. (Koons 1994, p. 447)[38]

Burgess (1990, 1998), in contrast, grants the existence of lower-order vagueness (*n*th-order vagueness for small *n*-values) but has questioned the existence of *n*th-order vagueness for all *n*. For Burgess,

> it is far too early in the day to claim with confidence that higher-order vagueness fails to terminate, either as a matter of logic or as a matter of fact. (1998, p. 240)

And indeed, like Koons, he defends this view from within a kind of hybrid conception of vagueness, whereby

> for each [vague] concept, at some point in the ascending orders of vagueness, higher-order vagueness will terminate for it ... Ordinary speakers could not know where this order is, still less could they know the exact location of these lines. (*ibid*., pp. 249–50)

For Burgess, vague expressions give rise to non-epistemic indeterminacy in the guise of borderline cases, but this indeterminacy is fairly shallow, as it were, since it does not generate a non-terminating hierarchy of borderlines cases, but rather a partial hierarchy which terminates at some unknowable point. Thus, at some level in the hierarchy the borderline cases to the borderline cases will have sharp but unknowable boundaries.[39]

In this last section, we will thus be concerned with two challenges: that there is first-order vagueness but no *n*th-order vagueness for $n>1$ (Koons and Wright), and that there is *n*th-order vagueness for small *n*-values (what we may loosely term *lower-order* vagueness) but no radical higher-order vagueness—*n*th-order vagueness for all *n* (Burgess).

Consider then the following iterativity principle for knowledge, namely a version of the so-called KK principle:

(KK) $\forall\alpha$ If $K^{n-1}(A)$ in α then $K^{n}(A)$ in α, (for $n>1$)

where 'K^{n-1}' abbreviates $n-1$ iterations of 'It is known that' (for $n>1$), and where 'A' schematizes sentences which are not themselves prefixed with the *K*-operator, and where for convenience the relativization of

[38] Koons thinks that there is vagueness *qua* non-epistemic borderline cases at first-order such that vagueness gives rise to truth-value gaps. However, he maintains that the limits of the gap are unknowable. It may in the end be that Koons is merely reluctant to postulate *non-epistemic* but not epistemic higher-order vagueness.

[39] I hasten to add that Koons and Burgess sponsor quite different hybrid conceptions of vagueness. While there may be advantages to be had from adopting a hybrid view of vagueness, the most natural point for the orders of non-epistemic borderline cases to terminate is at first-order. Consequently, the conception of vagueness offered by Burgess is unmotivated.

knowledge to a speaker *s* has been left out but can be taken to remain implicit in what follows. Thus, where $n=2$, and 'A' says that *a* is hot, for example, we have: $\forall\alpha$ If it is known that *a* is hot in α then it is known that it is known that *a* is hot in α. If this principle fails for $n=2$, for example, then there must at least be second-order vagueness; if it fails for $n=3$, then there must at least be third-order vagueness, and so on. If this principle fails for all *n*, then there must be *radical* higher-order vagueness, that is, a non-terminating hierarchy of borderline cases. Burgess (1990, 1998) effectively argues that while there is lower-order vagueness (that is, **KK** fails for small *n* values), the orders of vagueness terminate at some unknowable point (that is, there is some largish and unknowable value for *n* for which **KK** is valid for all values $>n$.)

Hyde, in contrast, claims that *radical* higher-order vagueness *qua* borderline cases arises

> because vague predicates typically fail to draw *any* apparent sharp boundaries within their range of signification. (Hyde 1994, p. 36)

Is he right to do so? Arguably, yes. If there is epistemic tolerance at each order *n* (which is just shorthand for saying that a sentence of the form 'K^n(A)' is epistemically tolerant for any *n*-value) then **KK** will fail whatever value we take for *n*. Loosely, if epistemic tolerance goes all the way up then so should genuine higher-order vagueness *qua* borderline cases. But can we rigorously show Hyde's claim to be correct?

One way to do so, is to reconfigure Wright's so-called paradox of higher-order vagueness in epistemic terms (Wright's original paradox is given in terms of a non-epistemic 'definitely' operator; see Wright 1987, 1992b.) Once we do that we can see that it is no paradox at all as Wright alleges, but rather it merely shows that the knowledge operator is absolutely non-iterative, that is, **KK** fails *for all n values.* We can certainly say that if epistemic tolerance goes all the way up then the following schema ought to hold for $n>1$:

(1) $K\neg\exists\alpha\ \exists\beta\ KK^{n-1}(\mathrm{A})$ in α and $K\neg K^{n-1}(\mathrm{A})$ in β (where β is close to α).

Given the factivity of 'it is known that' then this entails:

(2) $\neg\exists\alpha\ \exists\beta\ KK^{n-1}(\mathrm{A})$ in α and $K\neg K^{n-1}(\mathrm{A})$ in β (where β is close to α).

For $n>1$, assume that

(3) $K\neg K^{n-1}(\mathrm{A})$ in β

and also assume for *reductio* that

(4) K^{n-1}(A) in α (where β is close to α).

Given the **KK** principle, then from line 4 we can infer:

(5) KK^{n-1}(A) in α

and by $\wedge$-I on 3 and 5, and two applications of $\exists$-I, this yields:

(6) $\exists\alpha\ \exists\beta\ KK^{n-1}$(A) in α and $K\neg K^{n-1}$ (A) in β (where β is close to α, and n>1)

which contradicts line (2) and so we reject 4 to infer:

(7) $\neg K^{n-1}$(A) in α.

Since 7 depends upon assumptions all of which are known to be true (that is, lines 1 and 3) then 7 is likewise known to be true; hence we infer:[40]

(8) $K\neg K^{n-1}$(A) in α

and by a step of conditional proof together with two steps of $\forall$-I, this yields:

(9) $(\forall\alpha)\ (\forall\beta)K\neg K^{n-1}$(A) in $\beta \rightarrow K\neg K^{n-1}$(A) in α

which depends only on line 1. But 9 is disastrous for it allows us to infer that if $K\neg K^{n-1}$(A) in some case α, then $K\neg K^{n-1}$(A) in all cases. We have four basic options:

(a) Retain **KK** for $n>1$, but reject line 1, and thus say that 'K^n(A)' is not epistemically tolerant for $n>1$.

(b) Retain **KK** for $n>m$, where m is some small value, and thus say that 'K^n(A)' is not epistemically tolerant for $n>m$.

(c) Retain **KK** for all $n>1$, but reject one or more of the other rules of inference

(d) Reject **KK** for all n, (and retain all the other rules of inference).

Let's briefly take each of these in turn. Option (a) is surely the least attractive. It ought to be entirely uncontroversial that the sentence '*x* is

[40] The principle employed here (namely, the principle which says that if a sentence depends upon a set of assumptions all of which are known to be true then this sentence is known to be true) allows us to prove, given further uncontroversial rules, the following distributivity principle: $K(\mathrm{A}\rightarrow\mathrm{B})\rightarrow(K\mathrm{A}\rightarrow K\mathrm{B})$. In the context of vagueness at least, such a distributivity principle is entirely plausible, as is the more general principle upon which it depends.

known to be hot' is vague (on our three dimensions). Option (b) is perhaps the one which Burgess would endorse. It amounts to saying that not only does **KK** hold for n-values greater than m (where m is, say, greater than 10) but that, as the above proof shows, this means that 'K^n(A)' ceases to be epistemically tolerant at level $m+1$, though the exact value for m remains unknown (and perhaps unknowable). But given the connection between epistemic tolerance and sorites-susceptibility (expressions which are epistemically tolerant can be utilized in a sorites paradox, and vice versa), then Burgess must argue that 'K^n(A)' is *not* sorites-susceptible for some n-value. But we can always employ a sentence of the form K^n(A) to generate a sorites paradox (if the sentence 'A' is itself sorites-susceptible). Consequently, (b) is no option either.

Option (c) is perhaps more plausible still. In endeavouring to show that there is genuine radical higher-order vagueness we have made use of rules of inference which have been brought into question when dealing with vagueness—*reductio ad absurdum* and conditional proof, being the two most obvious ones. So it is certainly true then one might seek to reject the import of the above proof by questioning the use of these rules from within some substantive non-classical theory of vagueness. But Black seems to accept classical logic, and Koons (1994) appears to sponsor classical logic in the meta-language. Even if one does think that the sorites paradox is to be defused from within some non-classical logic that does not entail that classical rules of inference should thereby fail in the context of showing that there must be higher-order vagueness. So at the very least the above derivation represents a challenge to find a 'relevant failing' in one the rules of inference employed. Thus we are left with the last option. **KK** fails for all n: there must be radical nth-order vagueness. If that is right, then from axioms which are uncontroversial (roughly, vagueness is epistemic tolerance) we can derive an important and controversial theorem from within our minimal theory of vagueness.

What has been achieved? The promises of the minimal theory of vagueness have been satisfied. We have found a way to give a relatively neutral and reasonably rigorous characterization of vagueness in terms of the phenomenon of epistemic tolerance. Vagueness (as we experience it at least) just *is* epistemic tolerance. Hence, we have found a way to distinguish the vague from the non-vague. In so doing, we have ensured that the vagueness debate is not skewed in favour of indeterminist over epistemic conceptions of vagueness. Moreover, we have rigorously shown (at least given classical logic) that vagueness *qua* epistemic tolerance and vagueness *qua* epistemic borderline cases is just the same phenomenon, contrary to what might initially be expected. Lastly, we have given just the sort of independent argument Koons has called for which shows that there must be radical higher-order vagueness. Radical higher-order vagueness is not an illusion as some have thought but a phenomenon that we are all beholden to admit.

Department of Logic and Metaphysics PATRICK GREENOUGH
University of St. Andrews
Edgecliffe
The Scores
St. Andrews
Fife KY16 9AL
Scotland
pmg2@st-andrews.ac.uk

References

Alston, William P. 1967: 'Vagueness'. *The Encyclopaedia of Philosophy*, Vol. 8, New York: Macmillan, pp. 218–221.

Black, Max 1939: 'Vagueness: an Exercise in Logical Analysis'. *Philosophy of Science*, 4, pp. 427–55.

Burgess, John 1990: 'The Sorites Paradox and Higher-Order Vagueness'. *Synthese,* 85, pp. 417–74.

——1998: 'In Defence of an Indeterminist Theory of Vagueness'. *Monist*, 81, pp. 233–52.

Channell, Joanna 1994: *Vague Language*. Oxford: Oxford University Press.

Chellas, Brian 1980: *Modal Logic: An Introduction*. Cambridge: Cambridge University Press.

Dummett, Michael 1973: *Frege: Philosophy of Language*. London: Duckworth.

——1975: 'Wang's Paradox'. *Synthese*, 30, pp. 301–24.

——1978: *Truth and Other Enigmas*. London: Duckworth.

——1991: *The Logical Basis of Metaphysics*. London: Duckworth.

Field, Hartry 1994: 'Disquotational Truth and Factually Defective Discourse'. *Philosophical Review*, 103, pp. 405–452.

——2001: *Truth and the Absence of Fact*. New York: Oxford University Press.

Fine, Kit 1975: 'Vagueness, Truth, and Logic'. *Synthese*, 30, pp. 265–300.

Frege, Gottlob 1979: *Posthumous Writings*, trans. Peter Long and Roger White. Oxford: Blackwells.

Horwich, Paul 1997: 'The Nature of Vagueness'. *Philosophical and Phenomenological Research*, 57, pp. 929–35.
Hyde, Dominic 1994: 'Why Higher-Order Vagueness is a Pseudo-Problem'. *Mind*, 103, pp. 35–41.
——1997: 'From Heaps of Gaps to Heaps of Gluts'. *Mind*, 106, pp. 641–60.
Kamp, Hans 1981: 'The Paradox of the Heap', in U. Monnich (ed.), *Aspects of Philosophical Logic*, Dordrecht: Reidel.
Keefe, Rosanna and Smith, Peter (eds) 1996: *Vagueness: A reader*. Cambridge, MA: MIT Press.
Keefe, Rosanna 2000: *Theories of Vagueness*. Cambridge: Cambridge University Press.
Koons, Robert 1994: 'A New Solution to the Sorites Problem'. *Mind,* 103, pp. 439–449.
McGee, Vann 1991: *Truth, Vagueness, and Paradox*. Indianapolis: Hackett.
McGee, Vann and McLaughlin, Brian 1995: 'Distinctions Without a Difference'. *Southern Journal of Philosophy,* 33, (supplement), pp. 203–51.
Peirce, Charles S. 1902: 'Vague', in J. M. Baldwin (ed.), *Dictionary of Philosophy and Psychology*, New York: Macmillan, p. 748.
Putnam, Hilary 1983: 'Vagueness and Alternative Logic'. *Erkenntnis*, 19, pp. 297–314.
Russell, Bertrand 1923: 'Vagueness'. *Australasian Journal of Philosophy and Psychology*, 1, pp. 84–92.
Sainsbury, Mark 1989: 'Tolerating Vagueness'. *Proceedings of the Aristotelian Society*, 89, pp. 33–48.
——1990: 'Concepts Without Boundaries', inaugural lecture, published by King's College London, Department of Philosophy.
——1991: 'Is There Higher-Order Vagueness?'. *Philosophical Quarterly*, 41, pp. 67–82.
——1995: 'Why the World Cannot be Vague'. *Southern Journal of Philosophy*, 33, (supplement), pp. 63–81.
Sorensen, Roy 1985: 'An Argument for the Vagueness of "vague"'. *Analysis*, 45, pp. 134–7.
——1988: *Blindspots*. Oxford: Oxford University Press.
Tye, Michael 1995: 'Vagueness: Welcome to the Quicksand'. *Southern Journal of Philosophy*, 33, (supplement), pp. 1–22.
Unger, Peter 1979: 'There are no Ordinary Things'. *Synthese*, 41, pp. 117–54.
Wheeler, Samuel 1979: 'On That Which is Not'. *Synthese*, 41, pp. 155–94.

Williamson, Timothy 1992: 'Vagueness and Ignorance'. *Proceedings of the Aristotelian Society*, (supplementary volume), 66, pp. 145–62.
——1994: *Vagueness*. London/New York: Routledge.
——1995: 'Definiteness and Knowability'. *Southern Journal of Philosophy*, 33, (supplement), pp. 171–91.
——1996: 'Wright on the Epistemic Conception of Vagueness'. *Analysis*, pp. 39–45.
——1997: 'Imagination, Stipulation, and Vagueness', in E. Villanueva (ed.) *Truth: Philosophical Issues,* vol. 8, Atascadero, CA: Ridgeview.
——1999: 'On the Structure of Higher-Order Vagueness'. *Mind*, 108, pp. 127–43.
——2000: *Knowledge and its Limits.* Oxford: Oxford University Press.
Wright, Crispin 1975: 'On the Coherence of Vague Predicates'. *Synthese*, 30, pp. 325–65.
——1976: 'Language Mastery and the Sorites Paradox', in Gareth Evans and John McDowell (eds), *Truth and Meaning: Essays in Semantics*. Oxford: Clarendon Press, pp. 223–47.
——1987: 'Further Reflections on the Sorites Paradox'. *Philosophical Topics*, 15, pp. 227–90.
——1992a: *Truth and Objectivity.* Cambridge, MA: Harvard University Press.
——1992b: 'Is Higher-Order Vagueness Coherent?'. *Analysis*, 52, pp. 129–39.
——1995: 'The Epistemic Conception of Vagueness'. *Southern Journal of Philosophy*, 33 (supplement), pp. 133–59.
——2001: 'On Being in a Quandary'. *Mind*, 110, pp. 45–98.