Australasian Journal of Philosophy
Vol. 83, No. 2, pp. 157–183; June 2005

Routledge
Taylor & Francis Group

VAGUENESS AS CLOSENESS

Nicholas J. J. Smith

This paper presents and defends a definition of vagueness, compares it favourably with alternative definitions, and draws out some consequences of accepting this definition for the project of offering a substantive theory of vagueness. The definition is roughly this: a predicate '*F*' is vague just in case for any objects *a* and *b*, if *a* and *b* are very close in respects relevant to the possession of *F*, then '*Fa*' and '*Fb*' are very close in respect of truth. The definition is extended to cover vagueness of many-place predicates, of properties and relations, and of objects. Some of the most important advantages of the definition are that it captures the intuitions which motivate the thought that vague predicates are tolerant, without leading to contradiction, and that it yields a clear understanding of the relationships between higher-order vagueness, sorites susceptibility, blurred boundaries, and borderline cases. The most notable consequence of the definition is that the correct theory of vagueness must countenance degrees of truth.

In this paper I present and defend a definition of vagueness, and draw out some consequences of accepting this definition.

Section I asks what we should properly expect of a definition of vagueness, and Section II reviews existing definitions. Section III explains a key background notion necessary to an understanding of the definition to be presented here. Section IV presents this definition and Section V explains its advantages. Section VI draws out the consequences of accepting this definition for the project of offering a substantive theory of vagueness and Section VII considers and rejects a variation on the definition offered here.

I. What Should We Demand of a Definition of Vagueness?

I take a *definition* of a property, object or phenomenon *P* to be a statement about what it is to be *P* that is *true, useful,* and *fundamental*.[1] The first and third criteria are obvious: a definition of *P* must capture the *fundamental fact(s)* about *P*; it must get to the heart of the matter and capture not just any truths about *P*, but the fundamental truths, from which others follow. The second criterion covers at least the following: a definition must not be circular ('human beings are human beings' may capture a fundamental fact about human beings, but it is not an adequate definition of *human being*); a definition must be clear and rigorous ('vague predicates draw blurred

[1]Compare the notion of *real definition* in Rosen and Smith [2004: 191–2].

Australasian Journal of Philosophy
ISSN 0004-8402 print/ISSN 1470-6828 online © 2005 Australasian Association of Philosophy
http://www.tandf.co.uk/journals DOI: 10.1080/00048400500110826

boundaries' captures a fundamental fact about vague predicates, but it is not sufficiently clear or perspicuous to be an adequate definition of vagueness—see Section II); a definition should not (unless this is shown to be unavoidable) define some problematic notion in terms of even more problematic notions; and a definition of *P* must link in usefully with the project of offering a substantive theory of *P* (the definition tells us which are the fundamental facts about *P*, and the theory of *P* then needs to account for these facts).

Some feel that a definition of *P* must meet a further condition: it must be something on which all the major theories of *P* can agree. The rationale here is that if proponents of theories *A* and *B* disagree over the definition of *P*, then they do not have a genuine disagreement about the nature of *P*: rather, theories *A* and *B* are really theories about *different things*, and their proponents are *talking past one another*. A recent example of someone who insists on this point is Greenough [2003]. Greenough asserts that by characterizing vagueness from a neutral standpoint we 'can at least ensure that we are all talking about the same thing from the outset in our inquiry into the nature and source of vagueness' [2003: 235], and claims that if we failed to be neutral in our initial characterization of vagueness, then there would 'be a very real sense in which there would be no disagreement about the character of vague language at all since each partisan would mean something different by the predicate "is vague"' [2003: 238–9]. But is this rationale convincing? It rests on an old-fashioned form of descriptivism according to which the sole determinant of the reference of our terms is the definitions we would give of those terms. This sort of descriptivism would never be accepted in philosophy of language without extremely detailed, good arguments, and I see no reason to accept it in the present context either. If the correct definition of vagueness turns out to imply that (say) supervaluationism, but not epistemicism, is an incorrect theory of vagueness, then rather than say that the supervaluationists never disagreed with the epistemicists about vagueness at all, for all along they were each talking about different phenomena, we should say that, just as the flat-earthers were wrong *about the earth*—not *right* about some *other*, flat thing—so too the supervaluationists were simply wrong—*about vagueness*. Vagueness is a real phenomenon, out there in the world (at least in the part of the world that is human language, if not also in extra-linguistic reality), and just as those who thought that water was an element and those who thought it was H_2O were talking about the same stuff—the kind underlying the salient examples of water—independently of whether this stuff satisfied the descriptions they associated with their terms for water, so too all of us in the philosophical literature on vagueness are investigating this real phenomenon, to a large extent regardless of what we say *about* it.[2] We cannot insulate ourselves from the possibility of offering an incorrect theory of vagueness by offering a definition of vagueness which 'makes our theory

[2]Some writers in other disciplines—e.g., linguistics and engineering—sometimes talk about other phenomena—e.g., lack of specificity, or contextual reference-shifting—under the heading 'vagueness'; but philosophers easily recognize (see for example Greenough [2003: n.27]) that these authors are lumping distinct phenomena in with the one that we investigate under the heading 'vagueness'.

correct' (that would just make our definition wrong, as well as our theory); and likewise, we do not need a definition of vagueness to ensure that we are all talking about the same phenomenon.[3]

Rejecting the view that definitions must be acceptable to all partisans raises two questions. First, if a definition of vagueness is not needed to hold the vagueness debate together and prevent different parties talking past one another, why do we need a definition of vagueness at all? Second, if the fact that theory *A* is incompatible with a proposed definition of vagueness does not rule out that definition, then how do we judge proposed definitions: what should make us think that one definition is correct and another incorrect? Vague predicates seem (at least) to generate sorites paradoxes, to have borderline cases, to draw blurred boundaries, and to be subject to higher-order vagueness. Which of these phenomena, if any, is fundamental, is the essence of vagueness? What are the relationships between these phenomena? A correct definition of vagueness will yield answers to these questions, and anyone interested in vagueness would surely like such answers. That is why we should look for a definition of vagueness. As for how to assess a proposed definition, we should judge it by its clarity and rigour, its intuitive appeal, and by how satisfying and how unified an account it yields of the various aspects of vagueness and of the relationships between them.

II. Existing Attempts to Define Vagueness

Until very recently, the closest thing in the literature to a definition of vagueness was an amalgam of two or three ideas: the borderline case idea, the blurred boundaries idea, and the sorites-susceptibility idea.[4] The topic of vagueness is usually introduced via the borderline case idea, which can be traced at least as far as Peirce [1902]. We learn that vague predicates are those, such as 'tall', for which there exist objects to which it is unclear whether or not the predicate applies—that is, for which there exist borderline cases. This does not yield an adequate definition of vagueness, however, for the borderline case phenomenon is not sufficiently fundamental. The existence of borderline cases is consistent with there being clear divisions between the cases to which the predicate applies, the cases to which it does not apply, and the borderline cases—but this sort of clear demarcation of cases is something which the ordinary idea of vagueness

[3]There is another reason for rejecting the demand that a definition of *P* must be agreeable to all partisans: it would make the project of defining anything—and certainly vagueness—almost, if not totally, impossible. Since the appearance of Greenough [2003]—in which he rejects any definition of vagueness which is not compatible with all current theories of vagueness—Dorr [2003] has produced a theory of vagueness which is *not* compatible with Greenough's definition. In a hotly-contested area such as vagueness—where there seems to be *nothing* on which *everyone* agrees—even the most asinine definition is bound to conflict with one or another substantive theory.

[4]For example, Keefe [2000: 6] introduces her topic by giving a short list of vague predicates, and then saying that these predicates share three interrelated features: 'they admit borderline cases, they lack (or at least apparently lack) sharp boundaries and they are susceptible to sorites paradoxes'. The first two features are almost universally accepted as defining features of vague predicates, but some authors—for example Soames [1999: 217]—deny that all vague predicates are sorites-susceptible. See Section V.E below.

rules out. Thus we come to the second idea: the essence of vagueness is not borderline cases *per se*, but blurred boundaries—an idea found in Frege's statement that if we represent concepts in extension by areas on a plane, then vague concepts do not have sharp boundaries, but rather fade off into the background (*Grundgesetze* II, §56; in Beaney [1997: 259]). Unlike the borderline case idea, the blurred boundaries idea does seem to be truly fundamental—to capture what is essential to vagueness. However this second idea also fails to yield an adequate definition: for when it comes to getting a clear understanding of vagueness and its associated problems, Frege's metaphor is not perspicuous enough to be really useful.[5] The third idea is that vague predicates are those which are susceptible to sorites paradoxes. Whether or not it is *true* that all vague predicates give rise to sorites paradoxes will be discussed in Section V.E (I think it is true), but in any case, this idea is not sufficiently *fundamental* to form the basis of an adequate definition of vagueness: for surely if vague predicates give rise to sorites paradoxes, they do so *because they are vague*, rather than their giving rise to such paradoxes being *constitutive* of their vagueness.

Recently Greenough [2003] has proposed a definition of vagueness as *epistemic tolerance*. According to Greenough, a sentence S is vague just in case it yields a truth when substituted into the following schema, where τ is a variable which ranges over the set of truth states {true, determinately true, not true, not determinately true, determinately determinately true, ...}, α and β are variables which range over actual and counterfactual cases, v is a function from actual and counterfactual cases to non-negative real numbers, such that the truth of S depends only on the value of v in actual or counterfactual cases, c is some small positive real number, and K_s abbreviates 'It is known (by a speaker s) that':

$$\forall\tau\forall\alpha\forall\beta \text{ if } |v(\beta) - v(\alpha)| < c \text{ and } K_s(S \text{ is } \tau) \text{ in } \alpha \text{ then } \neg K_s(S \text{ is not-}\tau) \text{ in } \beta.$$

Thus S is vague just in case, for any two cases which differ by at most a small amount in the one respect that makes a difference to the truth of S, if a speaker knows that S is (say) true in one case, then she does not know that S is not true in the other case. But of course as stated this cannot be correct—and so Greenough adds three conditions which 'ensure that a speaker's ignorance does not result from the wrong source but solely from the vagueness of the sentence S': speaker s knows the value of v in every case; s knows the meaning of S; and 'we restrict the range of α and β to "normal" cases of judgement conditions for the speaker s' [2003: 259–60].

There is an inherent tension between looking for a definition of vagueness which is acceptable to all partisans, and looking for a definition which is useful and fundamental—so it is not surprising that Greenough's definition is neither as useful nor as fundamental as we should want a definition of vagueness to be. On the first point: Greenough's definition makes use of two notions—the notion of a speaker *knowing that*

[5]Frege was fully aware of the limitations of his metaphor: 'this is admittedly a picture that may be used only with caution' [Beaney 1997: 259].

something is the case, and the notion of *normal observers in normal conditions*—which are at least as hard to define as vagueness.[6] His definition thus merely pushes back the problem of defining vagueness: we will not have a complete definition of vagueness until we have a definition of knowledge, and a definition of a normal observer in normal conditions—and (to put it mildly) it does not look as though we will have these any time soon.

On the second point: it is extremely natural to respond to Greenough's definition as follows. 'Ah, I see ... suppose you cannot know that S in α and that not-S in β... and suppose that this ignorance is not due to the fact that you do not know what S means, nor due to the fact that you do not know the value of v in α or β, nor due to the fact that you are perceptually impaired in α or β... then this ignorance arises *because S is vague*.' But Greenough must baulk here: he is not saying that your ignorance *arises from* the vagueness of S; he is saying that S's vagueness *consists in* the fact that you are ignorant and that this ignorance does not arise from various other possible sources. This seems wrong. It is not at all implausible that vagueness leads to the sort of ignorance Greenough describes; but it is not plausible at all that the existence of such ignorance is the *fundamental fact* about vagueness. Thus Greenough's definition is subject to the same sort of objection as the definition in terms of sorites-susceptibility: we want to say that vague predicates engender sorites paradoxes and certain sorts of ignorance *because* they are vague; but if we *define* vagueness as sorites-susceptibility or as epistemic tolerance, then we cannot do this—we miss out on this explanation.[7]

III. Closeness

I shall now explain a key background notion necessary to an understanding of the definition of vagueness to be presented in Section IV.

For any set S of objects, and any predicate 'F'—vague or precise—a competent user of 'F' can discern relationships of *closeness* or *nearness* or *similarity* amongst the members of S: closeness or nearness or similarity in the respects that are relevant to—or determine—whether something is F (for short, 'F-relevant respects'). For example, consider the term 'red', and the set of all visible objects in the room: the orange things are closer in red-relevant respects to the red things than are the green things. This is not

[6]Greenough mentions only the notion of normal judgemental conditions for a speaker; but he really needs to add that the speaker herself is normal.

[7]Compare Johnston's [1989; 1993] missing explanation argument against dispositional theories of value and colour. At this point Greenough may respond that his definition was merely intended to be extensionally correct—not to capture the essence of vagueness. But this just brings us back to the question of what we should demand of a definition of vagueness. Once we see that the usual rationale for wanting a definition acceptable to all partisans carries no weight; and once we note that Greenough's definition fails to meet his own standards (because it is not compatible with at least one current theory of vagueness), and that furthermore those standards are probably impossible to meet (see note 3); then we should demand that a definition be more than merely extensionally correct: it should capture the fundamental fact about vagueness—the essence of vagueness—and this is something that Greenough's definition fails to do.

to say that orange things are *more red* than are green things: neither orange things nor green things are red *at all*. Rather, it is to say that *in the respects that determine* whether something is red, orange things are more similar to red things than are green things. Think of a colour wheel (or better, a colour solid) and imagine locating each object on the point of the colour wheel that has the same colour as that object: the closeness relationships amongst objects in the room in red-relevant respects correspond to the relationships of *spatial* closeness amongst these objects when they are located on the colour wheel in this way. Exactly the same can be said for the green-relevant respects: the respects that are relevant to whether an object is green are the *same* as the respects that are relevant to whether an object is red, and thus the similarity relationships amongst objects in green-relevant respects are the same as the similarity relationships amongst objects in red-relevant respects (and similarly for all other colours). For example, two bright red objects are very similar in respects relevant to whether something is green, even though neither is green at all; and conversely, a dark olive green object and a bright pea green object are not *very* similar in respects relevant to whether something is green, even though both are clearly green.

Now consider a non-colour predicate, and indeed a non-vague predicate, for example 'weighs over one kilogram'. What determines whether this predicate applies to an object is the object's weight. We can visualize relationships of closeness in weight on the set of visible objects in the room by imagining locating every object in the room on a long line, with an object weighing x kilograms being placed x metres from the beginning of the line: the closeness relationships amongst these objects in respects relevant to whether a thing weighs over one kilogram correspond to the relationships of *spatial* closeness amongst them when they are located on the line in this way. Thus the black grand piano in the corner and the black fragment of ash on the window sill are very similar in red-relevant respects, but not so similar in respects relevant to whether a thing weighs over one kilogram; while the blue and red pen caps on the table are not very similar in red-relevant respects, although they are very similar in respects relevant to whether a thing weighs over one kilogram.

The similarity relationships corresponding to other words are often not as easy to visualize as the ones associated with colour terms and the ones associated with the predicate 'weighs over one kilogram', but they are always just as apparent to users of these terms. Consider, for example, the term 'heap'. A twenty-grain pile of sand is close to a twenty-one-grain pile of sand in respects relevant to whether something is a heap; a pile of ten olives is further away; but not as far away as an armchair, which itself is reasonably close (in heap-relevant respects) to a dishwasher—for in the respects relevant to whether something is a heap (for example: size; the extent to which the thing is composed of separate smaller things held together only by gravity; shape) the armchair and the dishwasher are not very different. These things are apparent to any competent user of the term 'heap', even though there is no familiar object (such as a colour wheel or a line) onto which we can map things in order to represent the relationships of similarity in heap-relevant respects amongst them as relationships of spatial closeness.

Thus far I have simply been stating obvious facts (obvious, at least, on reflection). The question now arises as to how to give a more precise account of them. We need to distinguish two sorts of similarity or closeness relationship that are apparent to competent speakers: relationships of *relative* closeness, and relationships of *absolute* closeness.

First, relative closeness. In respects relevant to whether something is red, the orange things are closer to the red things than are the green things; in respects relevant to whether something is a heap, the twenty-grain pile of sand is closer to the twenty-one-grain pile of sand than is the ten-grain pile of sand; and so on. Given a set of objects and a predicate, we wish to represent the relative closeness relationships on that set associated with that predicate. The simplest and most general way to do this is directly in terms of a three-place relation $x \overset{F}{\leqslant}_z y$: '$x$ is at least as close to z as y is, in F-relevant respects'. However, it should be noted that in a given case—that is, for a given predicate 'F'—it may be that what speakers directly discern is not a ternary similarity or closeness relation, but some other sort of structure which *yields* such a relation—for example a metric structure, or some other sort of topological structure. Consider for example the predicate 'nearly home'. Bill, Ben, and Bob share a house. All three are heading home: Bill is 10 kilometres away, Ben is 9.5 kilometres away, and Bob is a few streets away. In 'nearly home'-relevant respects, Ben is closer to Bill than Bob is (even though Bob is nearly home and the other two are not). However what competent users of 'nearly home' discern most directly is not such facts as this, but facts about the spatial *distances* between objects; they then calculate that Ben is closer to Bill than Bob is in 'nearly home'-relevant respects if the difference between Ben's and Bill's distances from home is less than the difference between Bob's and Bill's distances from home. In general, I do not suppose that there is a metric or a topology on a given set of objects associated with a given predicate (cf. note 12); I assume only that there is a structure of relative closeness relationships, represented by a three-place relation. However in many actual examples there *is* an associated metric—and often it is what is most salient to speakers.

The question arises as to what properties the three-place relation has. Extending terminology for binary relations in an obvious way, we may assume that it is *transitive*: $\forall x,y,z,w(x \leqslant_w y \wedge y \leqslant_w z \rightarrow x \leqslant_w z)$.[8] We may also assume that it is *reflexive*: $\forall x,y(x \leqslant_y x)$. But what about *antisymmetry*: $\forall x,y,z(x \leqslant_z y \wedge y \leqslant_z y \rightarrow x = y)$? What about *connectedness*: $\forall x,y,z(x \leqslant_z y \vee y \leqslant_z x)$? I do not suppose that there are any *general* answers to these questions—that is, answers that apply to *all* predicates 'F'. Whilst it is quite clear that for any predicate 'F', competent users of 'F' can discern on a given set of objects a structure of relative closeness

[8]Here and in what follows I omit the superscript 'F' on the relation symbol to indicate generality of the schematic sort; i.e., I am here asserting transitivity for all particular relations $\overset{F}{\leqslant}$.

relationships associated with 'F' (whether they discern this structure directly, or via some other structure), specifying this structure in detail for a given predicate—and hence answering questions of the sort just posed—involves a great deal of work. In some cases this work has been done: for example, the closeness relationships amongst objects in colour-relevant respects are codified in the colour solid. In other cases it has not been done, and very likely never will be: for example, we have no general theory that codifies our thoughts about closeness of objects in respects relevant to whether something is a chair.

Second, absolute closeness. In respects relevant to whether something is tall, $5'$ Bill and $5'{}^{1}/_{64}''$ Ben are very similar; in respects relevant to whether something is a heap, a 10,000 grain pile of sand and a 10,001 grain pile of sand (of a very similar shape) are very similar; and so on. Given a set of objects and a predicate, and an associated structure of relative closeness relationships (represented, in general, by a reflexive, transitive three-place relation) how are we to represent this additional structure of absolute closeness relationships? The obvious thought is by a two-place relation $x \overset{F}{\approx} y$: 'x is very close to y, in F-relevant respects'. Of course, this notion of 'very close' is undoubtedly vague, and thus a completely adequate representation of it will have to wait until we have an adequate account of vagueness. For example, if the epistemicist is right about vagueness, then the representation of this relation by an ordinary set of ordered pairs is adequate; however, if the fuzzy approach to vagueness is right, then this relation will have to be represented by a fuzzy set of ordered pairs. For present purposes, however, all that matters is that for a given predicate 'F', and certain given sets of objects, there will be *clear cases* of pairs of objects which are, in an absolute sense, very close in respects relevant to whether something is F. In particular, in a sorites series for a predicate 'F', adjacent members of the series are very close in F-relevant respects (this is how the paradox gets off the ground—see Section V.D).

IV. Vagueness as Closeness

I shall now present a first version of my proposed definition of vagueness (the final version will be presented under point (5) below). A predicate 'F' is vague just in case it satisfies the following condition, for any objects a and b:

> **Closeness** If a and b are very close in F-relevant respects, then 'Fa' and 'Fb' are very close in respect of truth.[9]

The Closeness condition can be stated more precisely if we assume the following framework. Suppose that we have a domain of discourse D, and a

[9] More accurately (but less readably): For any objects a and b, and any singular terms '***a***' and '***b***', if a and b are very close in F-relevant respects, and '***a***' refers to a and '***b***' refers to b, then '***Fa***' and '***Fb***' are very close in respect of truth.

set T of truth values. Consider a function from D to T which assigns to each object x in the domain the truth value of the sentence 'Fx' (where 'x' denotes x). Call this the *characteristic function* of the predicate 'F'. Let $[Fx]$ be the value of the characteristic function for 'F' at the object x. Let $\overset{F}{\approx}$ be the relation on D of being very close in F-relevant respects, and let $\overset{T}{\approx}$ be the relation on T of being very close in respect of truth (see below for an explanation of the latter notion). Then the Closeness condition may be stated thus:

$$x \overset{F}{\approx} y \Rightarrow [Fx] \overset{T}{\approx} [Fy].$$

Precise terms need not satisfy Closeness: two objects a and b might be very close in respects relevant to whether something (say) has a mass of at least one kilogram (a's mass is one kilogram, while b's is a milligram less), and yet it be *true* that a has a mass of at least one kilogram and *false* that b has a mass of at least one kilogram.

Several points of clarification and qualification are required:

(1) What is this notion of closeness or similarity of sentences in respect of truth? Well, it is just a special case of the notion of closeness or similarity of objects in respect of a property, and I take this to be a familiar notion. Suppose I give you a bunch of objects, from which I single out one object a. I then ask you to select, from the rest of the objects, the ones that are very similar to a in respect of possession of some property. Suppose the objects are apples and the property is *being ripe*, or *being red*, or *being heavy*, or *being an apple*, or *being a prime number*—this should not cause you any difficulty. Suppose the objects are sentences and the property is *being in English*, or *being in German*, or *being grammatical*, or *being long*, or *having poetic merit*—this should not cause you any difficulty.[10] Nor should it cause you any more difficulty if the property is *being true*. In particular, one's understanding of the notion of similarity of sentences in respect of truth is entirely separate from one's views about whether truth comes in degrees. Suppose one thinks that truth does not come in degrees: then one will think that two sentences are very similar in respect of truth just in case they are both true *simpliciter*, or both false *simpliciter*. There is no problem here: one understands the idea of finding the sentences which are similar to a in respect of truth—if a is true one finds the true ones, and if a is false one finds the false ones. Thus it is certainly *not* the case (as some readers of earlier drafts of this paper have thought) that there is no way to *understand* the idea of two sentences being very similar in respect of truth—and hence to understand Closeness—unless one accepts a degree-theoretic framework. The idea of similarity of sentences in respect of truth is just a special case of the idea of

[10] Of course, in all these cases you might have problems *determining* whether some object is indeed similar to a in the relevant respect—but you should not have any difficulty knowing what it is you are trying to determine, i.e., no problem understanding the *idea* of locating the objects which are very similar to a in some given respect.

similarity of two objects in respect of possession of some property, and we can make sense of this idea without reference to any theory of whether property-possession is an all-or-nothing matter or a matter of degree.[11] Once we have the idea of two sentences being very similar in respect of truth, we can derive the relation $\overset{T}{\approx}$ of two truth values being very similar in respect of truth: the values are similar in respect of truth just in case the sentences which have those values are similar in respect of truth.

(2) In many semantic frameworks, the same objects serve as membership values and as truth values. Given a background set S, (the characteristic function of) a subset B of S is a function from S to the set of values. Suppose this function assigns the value v to the object b: the idea is that v represents the extent to which b is a member of B, that is, the extent to which b possesses the property B (this property being thought of as represented by the set B). Now consider the sentence '*Pa*'. An interpretation of the language consists in (among other things) a domain (a set of objects), an assignment of a member of the domain to each name, and an assignment of a subset of the domain to each predicate. Consider the interpretation with domain S, on which '*a*' is assigned b as its denotation, and '*P*' is assigned B as its extension. The truth value of the sentence '*Pa*' on this interpretation is the object v: that is, '*Pa*' is true to precisely the extent to which b is B. Thus the values represent both amounts of property-possession, and amounts of truth. For example both classical model theory and fuzzy model theory fit the description just given (their sets of values are $\{0,1\}$ and $[0,1]$ respectively), but the supervaluationist theory does not. In any such system, what I have non-standardly called the characteristic function of the predicate '*F*' is identical to what is standardly called the characteristic function of the set of *F*'s (i.e., the set which is the extension of '*F*'). If we were working in such a system, we could also state Closeness this way: 'If *a* and *b* are very close in *F*-relevant respects, then they are very close in respect of *F*', or in other words: 'If two objects are very close in the respects that determine whether something is *F*, then they are very close in respect of *F* itself.' In general, however, I do not assume we are working in such a system, which is why I officially state Closeness in the formal mode.

(3) There is a close affinity between Closeness and the idea of *continuity*. The intuitive idea behind the notion of a continuous function is that a small change in input always produces a small change in the value of the function. Thus it might seem that we could say that a predicate is vague just in case its characteristic function is

[11]This will be important in Section VI, where I shall argue that—given the existence of a sorites series for the predicate '*F*'—one cannot *accept* the claim that '*F*' conforms to Closeness unless one countenances degrees of truth. The difference between non-degree-theorists not being able to *understand* Closeness, and their not being able to *accept* that any predicate conforms to Closeness, is crucial.

continuous. However the details get complex, and we run into a problem which the Closeness definition does not face—and for these reasons I have chosen to present my definition of vagueness in terms of Closeness rather than continuity.[12]

(4) Closeness does not explicitly deal with the possibility of truth value gaps. To deal with this possibility, we stipulate that if $x \overset{F}{\approx} y$, then Closeness is satisfied if neither [*Fx*] nor [*Fy*] exists, but not if one exists and the other does not. This is not an arbitrary stipulation: for motivation, see the discussion of truth value gaps in Section VI.

(5) Now for an important qualification. For a start, note that we would not want to say that a predicate '*F*' is vague if it satisfies Closeness *trivially*—that is, either because there are *no a* and *b* such that *a* and *b* are very close in *F*-relevant respects, or because for *every a* and *b*, '*Fa*' and '*Fb*' are very close in respect of truth. Furthermore, note that we would not want to say that in order for a predicate to be vague, it must satisfy Closeness non-trivially *across its entire domain of application*. For consider the predicate 'is tall, or exactly four feet in height',[13] abbreviated '*E*': if Bob is four feet in height, and Bill is one nanometre taller, then Bob and Bill are very close in respects relevant to whether a thing is *E*, yet 'Bob is *E*' is true, while 'Bill is *E*' is false—and hence these two sentences are *not* very similar in respect of truth—and yet intuitively, '*E*' is vague. Consider also the predicate 'is tall, or greater-than-or-equal-to exactly four feet in height', abbreviated '*O*': this predicate also fails to satisfy Closeness, but intuitively it is *not* vague. So what is the difference between '*E*' and '*O*'?

A first thought is that there is *some* subset of the overall domain of discourse[14] such that for any objects *a* and *b in that subset*, it is (non-trivially) the case that if *a* and *b* are very similar in *E*-relevant respects, then '*Ea*' and '*Eb*' are very similar in respect of truth, while there is *no* subset of the overall domain of discourse such that for any objects *a* and *b* in that subset, it is (non-trivially) the case that if *a* and *b* are very

[12]Very briefly on the details and the problem: in order for the notion of continuity to be well-defined in this context, we need to suppose that there are topologies defined on the domain and codomain of the characteristic function for '*F*'. The topology on the codomain (which is the set of truth values of our logical framework, whatever that is) codifies the notion of closeness in respect of truth. On the domain we need not just one topology, but one topology for each predicate '*F*', with the topology associated with '*F*' codifying the notion of closeness in *F*-relevant respects. We then say that '*F*' is vague if its characteristic function—from the domain of discourse *endowed with the F topology* to the topological space of truth values—is continuous. The *problem* concerns discrete domains (for example, a finite set of colour patches each distinguishable in colour from the others, as opposed to a continuously varying spectrum of colours). *Any* function from a set endowed with the discrete topology is continuous, which would yield the result that *every* predicate is vague. (Note that it is *not* the case that every predicate '*F*' automatically satisfies Closeness relative to a discrete domain of discourse. Without going into the details, the reason for this boils down to the fact that in the official definition—as opposed to the intuitive idea—of continuity, the idea of *absolute* closeness drops out, the whole thing ultimately being done in terms of *relative* similarity; whereas this is not the case with Closeness, which makes use of the absolute notion of closeness.) We could avoid this by stipulating that the domain of discourse is the set of all possible objects, and that for each predicate '*F*' which is intuitively vague, the *F* topology on the domain of all possible objects is not discrete. But of course the latter stipulation may simply not be true, and in any case it would be best to avoid burdening our definition of vagueness with such heavy metaphysical commitments.

[13]Thanks to Kit Fine for this example. A similar problem is posed by Weatherson's example 'is in his early thirties' [2003: 2]; the comments below apply to all examples of this sort.

[14]That is, the domain we quantify over when we say 'absolutely *everything* there is' (with our quantifiers 'wide open', as it is sometimes put). I assume this domain includes persons of all sorts of different heights.

similar in *O*-relevant respects, then '*Oa*' and '*Ob*' are very similar in respect of truth. However this is not correct: there *is* a subset of the domain of discourse over which '*O*'—and other intuitively precise predicates such as 'is greater-than-or-equal-to exactly four feet in height'—non-trivially satisfies Closeness: the subset consisting of people less than three feet in height or more than seven feet in height.

We can avoid this problem (and avoid the need for a separate non-triviality clause) as follows. Say that a set *S* is *F-connected* iff for any two objects in *S*, either they are very close in *F*-relevant respects, or they can be linked together by a chain of objects—all of which are in *S*—with adjacent members of the chain being very close in *F*-relevant respects. Say that a set *S* is *F-uniform* iff for every *a* and *b* in *S*, '*Fa*' and '*Fb*' are very similar in respect of truth; if a set is *not* *F*-uniform it is *F-diverse*. Say that a predicate satisfies Closeness *over a set S* iff it satisfies Closeness as stated above when the initial quantifiers 'for any objects *a* and *b*' are taken as ranging over *S*. Now we arrive at our final definition of vagueness:

> **Vagueness as Closeness** A predicate '*F*' is vague iff there is some *F*-connected, *F*-diverse set *S* of objects such that '*F*' satisfies Closeness over *S*.

We can classify vague predicates as follows. A predicate is *fully vague* iff it is vague, and furthermore satisfies Closeness over *every* *F*-connected, *F*-diverse subset of the overall domain of discourse. A predicate is *partly vague* iff it is vague but not fully vague. Thus 'tall' is fully vague, '*E*' is partly vague, and 'is exactly four feet in height' and '*O*' are not vague. Intuitively, these classifications are exactly the right ones.

Finally, two further points of clarification:

(6) What of predicates '*F*' ascribing *fundamental* properties—for example 'spin up'—where there *are no* *F*-relevant respects? We have two choices: we can say that *every* x and y are such that $x \overset{F}{\approx} y$, or that *no* x and y are such that $x \overset{F}{\approx} y$. It makes no difference: either way, any such predicate '*F*' will count as non-vague by my definition.[15] This is a consequence I am happy to accept. I am also happy to accept the consequence of my definition that any predicate which applies to everything to exactly the same extent is non-vague.[16]

(7) My definition of vagueness is framed in terms of one-place predicates. The generalization to an n-place predicate '*R*' is straightforward: if the n-tuples $\langle x_1, \ldots, x_n\rangle$ and $\langle y_1, \ldots, y_n\rangle$ are very close in *R*-relevant respects, then '$Rx_1, \ldots, x_n$' and '$Ry_1, \ldots, y_n$' are very close in respect of truth. For example, if ⟨Bill,Ben⟩ and ⟨Bob,Maisy⟩ are very close in respects relevant to whether the first-mentioned person loves the second-mentioned person, then 'Bill loves Ben' and 'Bob loves Maisy'

[15]Case (i): no x and y are such that $x \overset{F}{\approx} y$. Then there are no *F*-connected sets. Case (ii): every x and y are such that $x \overset{F}{\approx} y$. So every set is *F*-connected. Now pick an arbitrary *F*-diverse set of objects. By definition of *F*-diversity, there are *a* and *b* in this set such that '*Fa*' and '*Fb*' are not very similar in respect of truth. But by hypothesis of this case, *a* and *b* *are* very similar in *F*-relevant respects. So '*F*' does not satisfy Closeness over this set.

[16]If '*F*' applies to everything to exactly the same extent then there are no *F*-diverse sets.

are very close in respect of truth. Furthermore, the account applies in an obvious way to the worldly counterparts of predicates, namely properties and relations: the property F is vague just in case the predicate 'F' which picks it out is vague.[17] What about the idea that there might be vague *objects*? I argue in Smith [in preparation] that for there to be vague objects is for certain properties or relations to be vague: for example the identity relation, the part-whole relation, and the occupation relation (which holds between an object and the spacetime points it occupies). As my definition applies to vague properties and relations, it therefore applies to vague objects as well.

V. The Merits of Closeness

Having presented a definition of vagueness, I shall now present the reasons why we should accept it.

A. Intuitive Appeal

The characterization of vagueness in terms of Closeness has significant raw intuitive appeal. Consider the two predicates 'is at least six feet in height' (precise) and 'is tall' (vague). Suppose that Bill is exactly six feet tall, and Ben is just under six feet tall. Bill and Ben are very close in respects relevant to whether a thing is at least six feet in height; yet 'Bill is at least six feet in height' is true, while 'Ben is at least six feet in height' is false; hence these two sentences are *not* very similar in respect of truth. Closeness is violated here—and this seems right. Now consider the vague predicate 'is tall'. Suppose you take two persons who are not very similar in respects relevant to whether a thing is tall: for example Bob and Bill, who differ in height by two feet. Might it be the case, intuitively, that the claims that Bob is tall and that Bill is tall are not similar in respect of truth? Certainly it might: a significant difference in height can make a significant difference to whether a person is tall. What if Bob and Bill are very close in height: for example they differ by less than one millimetre? Could it be that the claims that Bob is tall and that Bill is tall are not very close in respect of truth (e.g., one is true *simpliciter* and the other false *simpliciter*)? Intuitively not: *an insignificant difference in height cannot make a significant difference to whether a person is tall.*

B. Tolerance Intuitions without Incoherence

Crispin Wright has identified a certain position in the philosophy of language which he calls the *governing view*. It has two parts. First, there is

[17]This is compatible with the supervaluationist view that there are no vague properties, only vague predicates, because these predicates do not pick out particular properties.

the thesis that mastery of a language consists in the internalization of a set of semantic and syntactic rules that are definitive of that language. Second, there is the thesis that masters of a language can gain an *explicit* knowledge of the rules of which—according to the first thesis—they have an implicit understanding, by reflecting on such things as: speakers' known limitations, for example of perception and memory; standardly accepted criteria of misunderstanding a given expression; the generally accepted purpose of certain expressions; and the standard ways in which new users are trained to use certain terms. Wright argues that if the governing view is correct, then vague predicates are *tolerant*, where a predicate '*F*' is tolerant with respect to ϕ if there is some positive degree of change in respect of ϕ that things may undergo, which is 'insufficient ever to affect the justice with which *F* is applied to a particular case' [Wright 1975: 334]. In present terms, we can express a close relative of this idea as follows:

> **Tolerance** If *a* and *b* are very close in *F*-relevant respects, then '*Fa*' and '*Fb*' are identical in respect of truth.[18]

The great problem with Tolerance is that, when conjoined with the claim that we can construct a sorites series for the predicate '*F*' (a series which begins with an object which *is F* and ends with an object which is *not F*, and in which adjacent items are very close in *F*-relevant respects) it leads to contradiction (in particular, to the claim that each object in the sorites series both is and is not *F*). Nevertheless, Tolerance itself—considered apart from its unpalatable consequences—has been found very appealing: whether or not they accept the full-blown governing view, many have been strongly inclined to accept Wright's use of this view in deriving the conclusion that vague predicates are tolerant.

In fact, all the considerations in favour of the idea that vague predicates conform to Tolerance are, equally, considerations in favour of the idea that vague predicates conform merely to Closeness (and not Tolerance as well);[19] and Closeness (without Tolerance) generates no contradictions.[20] Thus one of the great advantages of Closeness is that it gives us tolerance intuitions *without* incoherence. Let us see how this works.

Wright's argument that, given the governing view, vague predicates are tolerant, proceeds by way of examples. First we have the case of 'heap'. If we look at the *occasions of use* of this predicate, it seems it must be tolerant with respect to small changes in number of grains. For 'heap' is 'essentially a coarse predicate': we use it on the basis of 'rough and ready judgement', and in these contexts would have no use for 'a precisely demarcated analogue'. In short, 'Our conception of the conditions which justify calling something a heap of sand is such that the justice of the description will be unaffected by

[18]That is: then they have the same truth value, or they both lack a truth value. Note that where Wright is concerned with *applying* predicates, Tolerance is concerned with the *truth* of sentences which express such applications; this difference has no significance for what follows.

[19]Note that Tolerance is a special case of Closeness.

[20]If a sorites series consists of objects $x_1, \ldots, x_n$, then Closeness tells us that 'Fx_i' and 'Fx_{i+1}' must always be very similar in respect of truth—which is quite compatible with 'Fx_1' being true *simpliciter* and 'Fx_n' being false *simpliciter*. (Note however that *given bivalence*, Closeness reduces to Tolerance. See Section VI below.)

any change which cannot be detected by *casual* observation' [1975: 335]. Second we have the case of predicates such as 'infant', 'child', 'adolescent', and 'adult'. If we look at the *social importance* of these predicates, we see that they must be tolerant with respect to small changes in maturity. For 'it would be irrational and unfair to base substantial distinctions of right and duty on marginal—or even non-existent—such differences' [1975: 336–7]. Third we have the case of colour terms. If we look at the *means of acquisition* of these predicates, we see that they must be tolerant with respect to small changes of shade. We learn such terms by ostension, so it must be the case that 'changes too slight for us to remember ... never transform a case to which such a predicate applies into one where such is not definitely correctly the right description. The character of ... basic colour training ... presupposes the *total memorability* of the distinctions expressed by our basic colour predicates' [1975: 336]. Wright sums up the lesson of the examples as follows: 'Our embarrassment about where to 'draw the line' with these examples is ... a reflection ... of the tolerance of the predicates in question' [1975: 337].

However, in all cases, the facts cited support not only the view that vague predicates are tolerant, but also the view that vague predicates conform merely to Closeness (and not Tolerance). Consider the first example. If 'heap' is a predicate of casual observation, then certainly there cannot be a difference of just one grain between a thing to which 'heap' clearly applies and a thing to which 'heap' clearly does not apply: for such a difference would not be noticeable to casual observers. It could be the case, however, that a negligible or insignificant difference (say, of one grain) between two objects makes a *negligible or insignificant* difference (i.e., one which we are entitled to *ignore* for all practical purposes) to the applicability of the word 'heap', and also that many insignificant differences add up to a significant one: this does not conflict with casual observationality, because many insignificant differences put together *are* noticeable to casual observers. The crucial point here is this: in order for a predicate to be usable in a context of casual observation, there does not need to be *no* difference in application of the predicate to objects that cannot be told apart by casual observation—there just needs to be no such difference in application that *cannot be safely ignored* in the context of casual observation. Of course, negligible differences in application add up and can*not* be ignored—but that is not a problem, because negligible differences in the objects to which they apply also add up, and *can* be detected by casual observation. Similarly, the social importance of predicates of degree of maturity is incompatible with there being a heartbeat that makes the difference between someone to whom 'child' clearly applies and someone to whom 'adult' clearly applies. It is not, however, incompatible with a negligible or insignificant difference (say, of one heartbeat) between two persons making a *negligible or insignificant* difference (i.e., one which we are entitled to *ignore* for all practical purposes) to the applicability of the word 'child', nor with *many* heartbeats making a *significant* difference. Thus in these first two cases Wright is unwarranted in saying that '[o]ur embarrassment about where to 'draw the line' with these examples is ... a reflection ... of the tolerance of the predicates in question'—for this embarrassment could just as well be a reflection of the fact that these predicates conform merely to Closeness (and not Tolerance).

What about the third example: could the use of colour predicates be taught ostensively if these predicates were not tolerant, but merely conformed to Closeness? Yes. Wright notes that we could not learn colour words ostensively if we did not have the capacity to remember how objects look—but he ignores another of our capacities, which is equally crucial to our being able to learn colour words ostensively: we are able to discern a structure of relationships of relative similarity or nearness or closeness in respect of colour on a given set of objects.[21] The ostensive teaching of colour words proceeds by the indication of paradigms and foils for each colour, with the implicit instruction that 'red', for example, applies to an object in proportion to its nearness to the red paradigms. Thus if an object is very similar to a red paradigm, it is red, or as good as for all practical purposes; if it is very similar to an orange paradigm, it is orange, or as good as for all practical purposes; and if it is somewhere between red and orange, then it is to some extent red and to some extent orange. This account of how we learn colour words applies to vague predicates quite generally: for each such predicate, we can discern an associated structure of similarity relationships on a given set of objects; paradigms for the predicate are indicated; and we learn that the predicate applies to an object in proportion to that object's similarity (in the relevant respects) to the paradigms.

Wright's examples do not, then, establish the tolerance of vague predicates. For all Wright has shown, these predicates could be intolerant (small changes need *not* make *no* difference), but conform to Closeness (small changes *never* make a *big* difference). Given also that we believe that some things are red and some not, that some men are bald and some not, and so on, one conclusion we might draw is that vague predicates are tolerant and incoherent. Of course the more reasonable conclusion to draw is that vague predicates are not subject to inconsistent requirements of use: they do *not* conform to Tolerance; they simply conform to Closeness. Accepting Closeness takes us far enough along the road to Tolerance to capture the intuitions which Wright uses to motivate Tolerance, without taking us so far as to run into contradiction. Closeness is everything that is right about Tolerance, and nothing more.

C. Higher-Order Vagueness

A central topic in the vagueness literature is that of *higher-order vagueness*. The term 'higher-order vagueness' is sometimes used rather freely, but one central conception can be distinguished. This conception is based upon the view that a predicate is vague if it has borderline cases. We begin by saying

[21]Children, for example, will quite naturally sort their coloured pencils into a rainbow pattern, without instruction, and without an actual rainbow pattern to copy: they can just see that red and orange are closer together than red and green (that is, they are closer together in respect of colour, and hence belong closer together in space, when the pencils are sorted in their tin).

that a vague predicate divides objects into three sets: the positive cases, the negative cases, and the borderline cases. But then, as Sainsbury notes, 'it soon appears that the idea that there is a sharp division between the positive cases and the borderline ones, and between the borderline cases and the negative ones, can no more be sustained than can the idea that there is a sharp division between positive and negative cases' [1991: 168]. So we now posit two new sorts of borderline case, between the original borderline cases and the positive cases, and between the original borderline cases and the negative cases—thus dividing objects into five sets. We may now generalize in an obvious way: in Sainsbury's terminology, a predicate is vague$_n$ if it divides objects into $2^n + 1$ sets; a vague predicate is vague$_n$ for some $n > 0$; a higher-order vague predicate is vague$_n$ for some $n > 1$; and a radically higher-order vague predicate is vague$_n$ for all n.

As already noted [159], merely possessing borderline cases is not enough to make a predicate vague. Hence Sainsbury's comment quoted in the previous paragraph, and later on in the same paper: 'it is a theoretical possibility that there be predicates which are vague$_1$ without being higher-order vague. Arguably, some predicates meet this condition, but, intuitively, meeting it is inconsistent with being a paradigm of vagueness' [1991: 169; see also 1991: 173]. On a similar note, Williamson and Graff write: 'For those unwilling to accept epistemicism, it might seem that vagueness *just is* higher-order vagueness' [2002: 17; cf. Keefe 2000: 31; Burns 1995: 29]. It seems then that higher-order vagueness as just described is inaptly named: a predicate which is—according to the conception just outlined—vague without being higher-order vague is not, intuitively, vague at all.

What is going on here is that a bad characterization of vagueness is accepted (i.e., vagueness as possession of borderline cases), and then our intuitive reservations about the characterization are given outlet in the positing of an additional phenomenon, over and above mere vagueness—i.e., higher-order vagueness. The proper thing to do would be to accommodate these residual intuitions within the characterization of vagueness itself. This is what the Closeness characterization does. The demand for higher and higher orders of borderline cases is just the demand for a gradual transition from the cases where the predicate clearly applies to the cases to where it clearly does not apply—rather than a transition marked by a series of jumps, whether one big jump or two smaller jumps or four even smaller jumps, and so on—and this idea of a gradual transition is captured in Closeness: as we proceed along the sorites series for '*F*', taking small steps in *F*-relevant respects as we move from one object *x* to the next (e.g., a difference of one hair in the case of 'bald', a difference of one nanometre in the case of 'tall', etc.), the sentence '*Fx*' takes correspondingly small steps in respect of truth. Thus the phenomenon of higher-order vagueness is really part of plain old vagueness—and the only thing that has stood in the way of wider acceptance of this view is adherence to a bad definition of vagueness, in terms of possession of borderline cases. The third advantage of the Closeness characterization of vagueness is, then, that it gives us a clear understanding of the intimate relationship between vagueness and higher-order vagueness.

D. Sorites Paradoxes

When it comes to sorites paradoxes, one can offer a negative solution, or a positive solution. The negative solution consists in arguing that there is nothing wrong with the paradox: vague predicates are incoherent, and the paradox simply brings this out [e.g., Dummett 1975]. A positive solution involves *two* tasks: one must say what is wrong with the paradox; and one must also explain why we did not reject the paradox as obviously mistaken when we first heard it. The fourth advantage of the Closeness characterization is that it provides the resources for accomplishing these latter tasks: it allows us to explain both how sorites paradoxes are mistaken, and why they are nevertheless compelling.

Consider a standard version of the sorites paradox:

1. This 10,000 grain pile of sand is a heap.
2. If we remove one grain of sand from a heap, we still have a heap.
3. So even if we removed 10,000 grains from this heap, we would still have a heap.[22]

Distinguish two readings of the second premise: the Closeness reading and the Tolerance reading. On the Closeness reading, the second premise expresses the claim that 'heap' conforms to Closeness. Two piles of sand *a* and *b* which differ by just a grain are very similar in heap-relevant respects, thus the two claims '*a* is a heap' and '*b* is a heap' must be *very similar* in respect of truth. So if *a* is a heap, then for all practical purposes we can just say that *b* is a heap too. On the Tolerance reading, the second premise expresses the claim that 'heap' conforms to Tolerance. Two piles of sand *a* and *b* which differ by just a grain are very similar in heap-relevant respects, thus the two claims '*a* is a heap' and '*b* is a heap' must be *exactly the same* in respect of truth. So if *a* is a heap, then *b* is a heap too: not just for all practical purposes, but absolutely and without qualification.

On the Tolerance reading, the conclusion follows from the premises. On the Closeness reading, it does not: each successive statement 'this is a heap' (said after removing one grain) must be very similar in respect of truth to the one before, but need not be *exactly the same* in respect of truth; and by the end, the final statement might certainly be false, i.e., the final object in the series might certainly be not a heap. No grain makes a significant difference to the truth of the claim that the thing before us is a heap, but enough grains might eventually make enough difference to render the claim false.

Suppose we accept Closeness (and not Tolerance). Why is the paradox compelling? Because (a) there is a natural reading of the second premise on which it expresses the fact that 'heap' conforms to Closeness, and the first premise is almost universally acceptable,[23] so we accept the starting-point of

[22]In general (given a sorites series for '*F*'): The first object in the series is *F*; If an object in the series is *F*, then so is the next object; Therefore the last object in the series is *F*.
[23]Of course some philosophers will reject it, e.g., Unger [1979a; 1979b].

the paradox,[24] and (b) there is *a* reading (the Tolerance reading) of the premises on which the argument is valid. Why is the paradox mistaken? Because reading the second premise in the way required to make the argument valid, we will *reject* that premise: for (*ex hypothesi*) we accept Closeness (and not Tolerance).[25]

I have argued that *if* we accept Closeness (and not Tolerance), we can explain why the sorites paradox is compelling, and also why it is mistaken. Given that we want to explain these two things, we should accept Closeness. I also believe that ordinary speakers *do* accept Closeness (and not Tolerance), and that this fact accounts for their reactions to sorites paradoxes. As one presents the paradox, a typical audience agrees that removing a grain of sand from a heap will not result in a non-heap (or that removing a hair from a hirsute man will not render him bald, etc.), but when one then says 'But then, by your reasoning, one grain of sand is a heap!', the audience baulks: they say we can remove one grain, or two grains, or even quite a few grains, but we cannot go on removing grains indefinitely and still have a heap—not because there will be a sudden change to non-heaps, but because all the insignificant changes add up to a significant change. These reactions fit perfectly with the hypothesis that the audience accepts Closeness (but not Tolerance): they accept the major premise of the sorites paradox, but they deny that they have thereby committed themselves to the claim that all men are bald (etc.). This is further evidence that Closeness captures ordinary intuitions about vagueness.[26]

Weatherson [2003: 6] claims (plausibly) that a good explanation of why the sorites premise 'If x_i is a heap then x_{i+1} is a heap' is seductive should *also* explain why the following claim is *not* attractive: 'Either x_i is not a heap or x_{i+1} is.' The Closeness explanation does this. The first sentence is attractive because it sounds like an expression of the fact that 'heap' conforms to Closeness; the second sentence is not at all attractive because it cannot plausibly be read as such an expression.

E. Other Characterizations Subsumed

A fifth advantage of the Closeness definition is that it subsumes characterizations of vagueness in terms of blurred boundaries, borderline cases, and sorites-susceptibility.

[24]Note carefully: I am saying that the second premise is attractive because it sounds like an expression of Closeness, *not* because (as some have thought) it is very nearly true. I have expressed no view at all on how true the second premise is. It might be flat-out false; this does not affect the claim that it is attractive because it *sounds like* an expression of the fact that 'heap' conforms to Closeness.

[25]It is worth stressing how different this response to the paradox is from simply saying that vague predicates are not tolerant. The latter response identifies a mistake in the paradox, but it provides no answer at all to the question why we find the paradox compelling—in particular why we are so strongly inclined to accept the second premise.

[26]Some will want to claim that ordinary persons find the major premise of the paradox compelling even if it is made absolutely clear that this premise is intended to express the fact that the vague predicate involved conforms to Tolerance rather than merely Closeness. In my experience this is simply false. Everyone knows that however many times you add zero to itself, you get zero—and thus sees in an instant that (for example) if adding a hair to a bald man makes absolutely no difference to the truth of the claim that he is bald, then adding a thousand hairs makes no more difference—and thus, given that the latter claim is absurd, rejects the former claim.

First, the Closeness characterization is a spelling-out of Frege's blurred boundary metaphor. The Closeness characterization consists in a constraint upon the relationship between the extension of a vague predicate and the absolute similarity relationships associated with that predicate: the extension cannot impose a big difference in *F*-ness between two objects which are not very different in *F*-relevant respects. Now assuming that for any vague predicate and any possible object, there is another possible object that is very close to the first in *F*-relevant respects, consider the set of all possible objects, structured by relationships of closeness in *F*-relevant respects. Given Closeness, the extension of '*F*' amongst this set cannot consist in a sharp line between the *F*'s and the non-*F*'s: rather, *F*-ness must gradually fade away as one travels further from the full-fledged *F* objects. To take a concrete example, consider the term 'red'. This term does not cut a sharp band out of the rainbow: as one moves across the points of the rainbow, small steps in red-relevant respects—which in this case correspond to small steps in space—can never, given Closeness, make for big changes in the truth of the claim that the point one is considering is red. By small steps one can move from full-fledged red points to full-fledged non-red points: but there is no sharp boundary between them that can be crossed in one small step. Thus Closeness gives a precise spelling-out of the blurred boundaries idea, which is at the heart of the intuitive conception of vagueness.

Second, the Closeness characterization captures the thought that vague predicates admit of borderline cases. Consider a vague predicate '*F*', and a sorites series $x_1, \ldots, x_n$ for '*F*'. 'Fx_1' is true *simpliciter*, and 'Fx_n' is false *simpliciter*; but given Closeness, it cannot be that there is an i such that 'Fx_i' is true *simpliciter* and 'Fx_{i+1}' is false *simpliciter*. There must then be sentences 'Fx_i' which are neither true *simpliciter* nor false *simpliciter*; and the corresponding objects x_i are borderline cases for '*F*'. Thus, if we characterize vagueness in terms of Closeness, we can see why vague predicates have borderline cases, without being committed to the false converse claim that every predicate which admits of borderline cases is vague.

Third, it has been claimed that vague predicates are sorites-susceptible. What is the relationship between vagueness and sorites-susceptibility (cf. note 4 above)? Well, suppose that we are dealing with a vague predicate '*F*', and we have a sorites series for '*F*'. In this case we can easily generate a sorites argument; and given Closeness, we can explain why this argument is compelling (and mistaken). Thus there is a close connection between vagueness and sorites-susceptibility. But sometimes it is difficult to imagine a sorites series for '*F*': a series of (possible) things ranging from one which is *F* to one which is non-*F*, with adjacent items in the series being very close in *F*-relevant respects. This is the possibility that Soames [1999: 217] has in mind when he denies that all vague predicates are sorites predicates, and as he says, in this case 'compelling versions of the sorites paradox are much harder to construct'. Nevertheless, they can usually be constructed by artificial devices. For example, begin with a paradigm *F* object, and remove one microscopic speck of matter at a time until eventually you reach an object that is non-*F* [Unger 1979a]. Alternatively, begin with a paradigm *F* object *a* and a paradigm non-*F* object *b*, and take as the intermediate objects in the series the things depicted by the frames of a piece of movie footage which consists in a slow motion morph from *a* to *b*. These ideas do not cover every possible case, but there are other tricks that can be employed. Thus I think that vagueness and sorites-susceptibility are closely bound together. Nevertheless, sorites-susceptibility should not be included as a defining mark of vagueness: the Closeness characterization is all we need. For if a sorites series exists for a predicate '*F*', then supposing that '*F*' conforms to Closeness, we can explain why the associated sorites argument is compelling. However, if '*F*' is *not* susceptible to a sorites paradox, this does not mean that '*F*' is not vague: it might be that '*F*' *is* vague, but there is no readily imaginable sorites series for '*F*'.[27]

F. Formal Properties

Last but not least, the Closeness definition is clear and easy to grasp, and does not define vagueness in terms of other contentious concepts. In order to understand the definition, one need only grasp two perfectly ordinary notions: the notion of two objects being very similar in respects relevant to the possession of some property; and the notion of two objects being very similar in respect of possession of some property.[28]

The Closeness definition thus has advantages which no other definition of vagueness can match. It captures fundamental intuitions about vagueness. It allows us to see how vague predicates can be usable in contexts of casual observation, have great social importance, and be taught by ostension, and yet be perfectly coherent. It yields a clear understanding of the relationships between the various phenomena of vagueness: higher-order vagueness, sorites paradoxes, blurred boundaries, and borderline cases. And it is clear, perspicuous, and easy to grasp, and does not define vagueness in terms of other contentious notions. Given these advantages, we should accept Closeness as the correct definition of vagueness. In the next section I explore the consequences of doing so.

VI. Capturing Closeness

Suppose that we have a sorites series $x_1, \ldots, x_n$ for the predicate '*F*'. If bivalence (thought of as the claim that not only are there only two truth values, but also, every declarative sentence has exactly one of them) is true, then Closeness must be violated here. 'Fx_1' is true *simpliciter*. x_1 and x_2 are

[27]Note that if a predicate '*F*' is vague according to the Closeness definition, then there is an *F*-connected, *F*-diverse set *S* of objects such that '*F*' satisfies Closeness over *S*. But an *F*-connected, *F*-diverse set *S* of objects does *not* automatically yield a sorites series for '*F*'. An *F*-diverse set is one for which it is *not* the case that for every *a* and *b* in *S*, '*Fa*' and '*Fb*' are very similar in respect of truth. Thus an *F*-diverse set *need not* contain an object which is clearly *F* and an object which is clearly not *F*—whereas a sorites series begins with an object which is clearly *F* and ends with an object which is clearly not *F*.

[28]Recall that these are different: 5′ Bill and $5'{}^{1}/_{64}''$ Ben are very similar in respects *relevant to* whether something possesses the property *being exactly 5′ tall*; but they are *not* similar in respect of possession of this property itself (one of them possesses it and the other does not).

very similar in F-relevant respects, so if 'F' conforms to Closeness, then 'Fx_1' and 'Fx_2' must be very similar in respect of truth. Given bivalence, the only way this can be the case is if 'Fx_2' is true *simpliciter*. Similarly on down the series. But towards the end, when we get to x_n, it must be the case that 'Fx_n' is false *simpliciter*. Thus if Closeness is to be respected, we need to abandon bivalence. The point can be made more directly: given bivalence, Closeness reduces to Tolerance: the only way in which two sentences can be *very similar* in respect of truth is by having the *same* truth value. So to capture Closeness without Tolerance, we will have to reject bivalence.

What if we have just two truth values, but we also allow truth-value gaps? This does not help; for now, as we move along our sorites series, there will come a point at which 'Fx_i' is true and 'Fx_{i+1}' lacks a truth value—and thus these sentences are *not* very similar in respect of truth, even though x_i and x_{i+1} *are* very close in F-relevant respects. The problem would be the same if we had a third value, Indeterminate, in place of the gap: for if sentence S is True [False] and sentence T is Indeterminate, then intuitively, it is *not* the case that S and T are very similar in respect of truth. It is important to be clear about what I am claiming here. Suppose I give you a bunch of sentences, some of which are true, some false, and some lacking in truth value (or possessing a third value, Indeterminate). I then give you a true sentence a and ask you to select the sentences which are very similar to a in respect of truth. In order to complete this task, you don't need to look at what the sentences *say*; you only need to look at their *truth-values*. Having done this, surely you will select the true sentences, and no others?[29] Remember, it is completely irrelevant in this context that sentence a says that Bill is bald, and sentence b says that Ben is bald, and Bill and Ben differ by just a hair. The only thing that matters at present is that sentence a is true, and sentence b lacks a truth value—and this surely makes it the case that these sentences are not very similar in respect of truth (no matter in what *other* respects the sentences *are* very similar). The situation is analogous to one where I give you a bunch of objects, some of which are green, some red, and some colourless. I then select a green object a, and ask you to select the objects which are very similar to a in respect of colour. Surely you will select only the green objects—even if one of the red or colourless objects is exactly like a in every respect other than colour.

Can the fuzzy theory accommodate Closeness? Yes. The fuzzy theory takes as its set of degrees of truth the real numbers between 0 and 1 inclusive. Now there are certainly pairs x and y of reals in [0,1] such that if sentence S's truth value is x and sentence T's truth value is y, then intuitively, S and T are very similar in respect of truth (for example, let $|x-y| = 0.00001$). Thus, the fuzzy framework can easily accommodate Closeness: it has a sufficiently rich structure of truth values to allow arbitrarily small steps in F-relevant respects to correspond to arbitrarily small steps in truth—while also allowing insignificant steps to add up to significant ones.

[29]Or at least, surely that's what you will try to do—of course you may *think* a sentence is true when in fact it is not, but that is beside the point here.

Can the supervaluationist theory accommodate Closeness? Not in its standard form—in which sentences are assigned the values (super) true or (super) false, or no value at all (or the value Indeterminate)—for the reasons given in the paragraph before last. But wait: someone might think that even if 'Bill is bald' is (super) true [false] and 'Ben is bald' lacks a truth value (or has the value Indeterminate), still they are close in respect of truth, because they are (classically) true/false on *almost* exactly the same admissible interpretations. But this is to confuse similarity in *respects relevant to* truth, with similarity in respect of truth itself. Assuming the supervaluationist framework, if two sentences are classically true/false on almost exactly the same admissible interpretations, then they are very similar in respects relevant to whether a sentence is (super) true; but if one sentence is (super) true [false] and the other lacks a truth value altogether (or has the value Indeterminate), then they are *not* very similar *in respect of truth*. To claim otherwise would be analogous to claiming that Al Gore is very similar to George Bush *in respect of being President of the USA*—which is false, because Bush is (definitely, totally) President, and Gore is not (at all, in any way)—when all that is true is that Gore is (or was) very similar to Bush *in the respects that determine* who is President. In the above situation, 'Bill is bald' and 'Ben is bald' are very close in the respects that determine whether a sentence is (super) true; but in respect of truth itself, they are not close at all—for just the reasons given in the paragraph before last.

The standard form of supervaluationism cannot, then, accommodate Closeness. However the degree-theoretic form of supervaluationism *can* accommodate Closeness. The degrees-of-truth supervaluationist countenances continuum-many truth values, represented by real numbers in the interval [0,1]: the degree of truth of a sentence is the size, as given by a normalized measure function, of the set of its admissible interpretations on which it is classically true [Kamp 1975; Kamp 1981: 234–5; Lewis 1983a: 228–9; Lewis 1983b: 69–70]. Like the fuzzy framework, this view has a sufficiently rich structure of truth values to allow arbitrarily small steps in F-relevant respects to correspond to arbitrarily small steps in truth—while also allowing insignificant steps to add up to significant ones.

So, if we have two or three truth values, we cannot accommodate Closeness, while if we have a continuum of degrees of truth, we can. Will anything less than a continuum of degrees of truth do the trick? This is an open question: the point for now is that we need a significant number of degrees of truth, in between truth *simpliciter* and falsity *simpliciter*, if we are to make room for Closeness. In our sorites series $x_1, \ldots, x_n$ for the predicate 'F', we need it to be the case that 'Fx_1' is true *simpliciter* and 'Fx_n' is false *simpliciter*, and that 'Fx_i' and 'Fx_{i+1}' are always very similar in respect of truth (because x_i and x_{i+1} are always very similar in F-relevant respects). This can only happen if there are gradations or degrees of truth in between truth *simpliciter* and falsity *simpliciter*, in such a way that two sentences can have *different* gradations of truth and yet it still be the case that the two sentences are *very similar* in respect of truth.

At this point in the dialectic it has sometimes been objected to me that there is no substance in the claim that we need degrees of truth in order to accommodate Closeness, because there is no way to understand the idea of two sentences being very similar in respect of truth, unless one accepts a degree-theoretic framework. However this is incorrect. We need to carefully distinguish two claims: (i) There is no way to understand the idea of similarity in respect of truth without accepting the idea that truth comes in degrees. This claim is simply false [165]. (ii) Given the existence of a sorites series for the predicate '*F*', there is no way to accommodate the claim that '*F*' conforms to Closeness without accepting the idea that truth comes in degrees. *This* is my central claim in this section of the paper. Everyone can *understand* the claim that vague predicates conform to Closeness, no matter what their views about truth; but only those who countenance degrees of truth can *accept* this claim (given the existence of sorites series for vague predicates).

Accepting the Closeness definition of vagueness thus has important implications for the project of offering a substantive theory of vagueness. But accepting Closeness does not point us to one particular theory: it only points us to degree theories in general. Within the class of such theories there is a huge gulf between supervaluationist-type theories—which see vagueness as a purely linguistic matter, as a matter of a loose fit between language and a world which is inherently perfectly precise—and fuzzy-type theories—which view vagueness as a worldly matter;[30] and then there are further divisions within each of these types as well. Furthermore, even epistemicists, and others who cannot accommodate vagueness as Closeness, do not simply have to pack up their toys and go home: they can re-enter the fray repackaged as error theories, which agree that for a predicate to be vague would be for it to satisfy Closeness, but argue that there are no such predicates. Such theories would start on the back foot, however—for it seems axiomatic that there *are* vague predicates. (For further discussion, see the end of Section VII.)

VII. Alternatives to Closeness

Proponents of non-degree-theoretic treatments of vagueness might try to argue that the Closeness characterization of vagueness is (subtly) incorrect. In this section I shall look at a possible response along these lines from the epistemicist. The epistemicist might claim that Closeness is wrong, and that the correct way to define vagueness is in terms of the following principle:

> **Pragmatic Closeness** If *a* and *b* are very close in *F*-relevant respects, then '*Fa*' and '*Fb*' are very similar in respect of *acceptability*.

The idea expressed in the consequent is that to whatever extent it is justifiable or reasonable to believe or to assert that *Fa*, it is to much the same

[30] For explanation of this distinction see Smith [2001; 2004: n.17; in preparation] and Rosen and Smith [2004].

extent justifiable or reasonable to believe or to assert that *Fb*. Thus the overall point is that the true idea to be captured by a theory of vagueness is an epistemological/pragmatic one, rather than—as the Closeness characterization would have it—an alethic/metaphysical one.

The problem with defining vagueness in terms of Pragmatic Closeness is that this definition does not capture the blurred boundaries idea, which is essential to the ordinary conception of vagueness. Suppose that we have a sorites series for '*F*', and someone asserts that Fx_i and not Fx_{i+1}, where x_i and x_{i+1} are adjacent items in the series. Intuitively there is a problem here. The epistemicist tells us that the problem is that this person has violated epistemic or pragmatic norms: she has unjustifiable beliefs, or she has asserted something which she does not believe. But intuitively the problem is deeper: there is a problem with the very idea that there *is* a sharp boundary between the *F*'s and the non-*F*'s, and thus someone who asserts that such a boundary does exist (in such and such place) seems not merely to have violated epistemic or pragmatic norms, but to have violated the norm of *truth*. This comes out particularly clearly if we imagine someone *guessing* that the boundary between the *F*'s and the non-*F*'s comes between x_i and x_{i+1}. Someone who makes such a guess seems, intuitively, to have misunderstood the nature of vagueness just as much as someone who makes the corresponding assertion. Yet the guesser—unlike the asserter—has not violated Pragmatic Closeness. Thus Pragmatic Closeness does not capture the ordinary conception of vagueness. The intuitive problem with the guess is that it *could not be true*: because vague predicates have blurred boundaries, not sharp but unknowable ones. Phenomenologically, there is an enormous difference between a proposition such as 'The boundary between baldness and non-baldness comes at the four hundredth hair' and a proposition such as 'The least upper bound of velocities reached by polar bears on 11th January 2004 was 31.35 kilometres per hour', which may well be true—and could certainly be *guessed* to be true, if one wanted to—but should not be asserted or accepted, as no-one could justifiably believe it to be true. Thus it would seem to be quite clear that our fundamental intuitions about vagueness are not merely, as the Pragmatic Closeness characterization would have it, epistemological/pragmatic: they are, as the blurred boundaries idea and the Closeness characterization would have it, alethic/metaphysical. Satisfaction of Pragmatic Closeness is not, then, of the essence of vagueness.

In opposition to the foregoing, Greenough [2003: 272, 274] writes:

> 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 ... 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 ... not that there is no boundary.

The upshot would be that I cannot claim—as I did in the previous paragraph—that intuition sides definitively with Closeness rather than

Pragmatic Closeness.[31] But if, as Greenough claims, there is an argument to back up his position here, it would seem that the argument must rest on a conflation of the content of an utterance and its assertibility conditions. Suppose a witness tells us that there was no-one there. This is incompatible with the claim that there was someone there. The barrister responds 'Ah, you are saying that from where you were positioned you could not see anybody'—a claim which is compatible with there having been someone there. This is clearly an insidious move: the witness said what he said because, from where he was positioned, he could not see anybody; but this is not what he said, which was that there was nobody there. Likewise in the case of vagueness. We can all agree that the ordinary speaker says that there is no sharp boundary between red and orange because she cannot see a clear boundary; but it is illegitimate to move from this to the claim that what she says is that there is no clearly visible boundary between red and orange. What any ordinary speaker, when faced (say) with an unobscured, not-too-distant rainbow—and whilst wearing her glasses or contact lenses and not under the influence of drugs or alcohol, etc.—would say is not the absurdly timid 'I cannot perceive a clear boundary between red and orange', but the stronger claim that there is no sharp boundary. This should not be controversial. If ordinary speakers made only the weaker claim, then epistemicism would not have been met with howls of disbelief (from both the opponents and some proponents of the view)[32] when it first came onto the philosophical scene—and Frege (to whom the blurred boundaries idea can be traced) would not have thought that vagueness has no place in a logically perfect language.

A note on the dialectic here: I am not objecting to *epistemicism*, on the grounds that it is counter-intuitive; I am objecting to the proposed definition of vagueness in terms of Pragmatic Closeness, on the grounds that it does not capture an essential feature of the ordinary conception of vagueness. The Pragmatic Closeness definition is simply unacceptable. Epistemicism might still be acceptable however—as an error theory. The idea would be that the correct way to codify our ordinary conception of vagueness is in terms of the Closeness definition; but lo and behold, there are no predicates which satisfy Closeness—i.e., no vague predicates. There are only predicates which satisfy Pragmatic Closeness, and of which the epistemicist theory is correct—i.e., which draw sharp but unknowable boundaries. Whether to accept an epistemicist error theory over a degree-theoretic treatment of vagueness would depend on a host of issues currently being thrashed out in the debate between different substantive theories of vagueness. The details of this debate are beyond the scope of this paper. Accepting the Closeness definition does not *end* this debate—but it does have important implications for it. One such implication would be that the epistemicist theory could only be correct if recast as an error theory.

[31]Note that Greenough uses the term 'blurred boundary' to mean no *known* boundary, whereas I use it to mean no *sharp* boundary. This terminological difference does not mitigate our substantive disagreement here.
[32]E.g. Williamson [1994: xi]: 'This book originated in my attempts to refute its main thesis . . . For years I took this epistemic view of vagueness to be obviously false, as most philosophers do.'

VIII. Conclusion

I have presented a definition of vagueness, set out its merits, and argued that it is superior to other definitions. I have also argued that we should not expect a definition of vagueness to have no impact on the debate between different substantive theories of vagueness—and indeed accepting the definition proposed here does have important consequences for this debate. The main such consequence is that unless we adopt an error theory of vagueness—that is, unless we accept that really there are no vague predicates—we need to adopt a theory of vagueness that countenances degrees of truth.[33]

University of Sydney

Received: February 2004
Revised: August 2004

References

Beaney, Michael, ed. 1997. *The Frege Reader*, Oxford: Blackwell.

Burns, Linda 1995. Something To Do With Vagueness, in *Vagueness*, ed. Terence Horgan, *Southern Journal of Philosophy* 33 (supplement): 23–47.

Dorr, Cian 2003. Vagueness Without Ignorance, *Philosophical Perspectives* 17: 83–114.

Dummett, Michael 1997 (1975). Wang's Paradox, in *Vagueness: A Reader*, ed. R. Keefe and P. Smith, Cambridge MA: MIT Press: 99–118.

Greenough, Patrick 2003. Vagueness: A Minimal Theory, *Mind* 112: 235–81.

Johnston, Mark 1989. Dispositional Theories of Value, *Proceedings of the Aristotelian Society* supp. vol. 63: 139–74.

Johnston, Mark 1993. Objectivity Refigured: Pragmatism Without Verificationism, in *Reality, Representation and Projection*, ed. J. Haldane and C. Wright, New York: Oxford University Press: 85–130.

Kamp, Hans 1975. Two Theories About Adjectives, in *Formal Semantics of Natural Language*, ed. Edward L. Keenan, Cambridge: Cambridge University Press: 123–55.

Kamp, Hans 1981. The Paradox of the Heap, in *Aspects of Philosophical Logic*, ed. Uwe Mönnich, Dordrecht: D. Reidel: 225–77.

Keefe, Rosanna 2000. *Theories of Vagueness*, Cambridge: Cambridge University Press.

Lewis, David 1983a (1970). General Semantics, in *Philosophical Papers 1*, New York: Oxford University Press: 189–232.

Lewis, David 1983b (1976). Survival and Identity, in *Philosophical Papers 1*, New York: Oxford University Press: 55–77.

Peirce, C. S. 1902. Vague, in *Dictionary of Philosophy and Psychology*, ed. J. M. Baldwin, New York: Macmillan: 748.

Rosen, Gideon and Nicholas J. J. Smith 2004. Worldly Indeterminacy: A Rough Guide, *Australasian Journal of Philosophy* 82: 185–98.

Sainsbury, Mark 1991. Is There Higher-Order Vagueness?, *Philosophical Quarterly* 41: 167–82.

Soames, Scott 1999. *Understanding Truth*, New York: Oxford University Press.

Smith, Nicholas J. J. 2001. *Vagueness*, Ph.D. thesis, Department of Philosophy, Princeton University. UMI ProQuest Digital Dissertations no. AAT 3021982, http://wwwlib.umi.com/dissertations/fullcit/3021982.

Smith, Nicholas J. J. 2004. Vagueness and Blurry Sets, *Journal of Philosophical Logic* 33: 165–235.

Smith, Nicholas J. J. in preparation. *Worldly Vagueness*.

Unger, Peter 1979a. There Are No Ordinary Things, *Synthese* 41: 117–54.

Unger, Peter 1979b. Why There are No People, *Midwest Studies in Philosophy* 4: 177–222.

Weatherson, Brian 2003. Vague Terms with Sharp Boundaries: Reply to Greenough, typescript.

Williamson, Timothy 1994. *Vagueness*, London: Routledge.

Williamson, Timothy and Delia Graff 2002. Introduction, in *Vagueness*, ed. Timothy Williamson and Delia Graff, Aldershot: Ashgate: xi–xxviii.

Wright, Crispin 1975. On the Coherence of Vague Predicates, *Synthese* 30: 325–65.

[33]For helpful comments and conversations, I thank John P. Burgess, Matti Eklund, Kit Fine, Amitavo Islam, Ed Mares, Gideon Rosen, Scott Soames, Marcelo Vasconez, Brian Weatherson, the anonymous referees, and audiences at Princeton University (October 2000), Victoria University of Wellington (May 2002) and the University of Otago (August 2002).