Vagueness, Tolerance and Contextual Logic*

Haim Gaifman

The goal of this paper is a comprehensive analysis of basic reasoning patterns that are characteristic of vague predicates. The analysis leads to rigorous reconstructions of the phenomena within formal systems. Two basic features are dealt with: semantic indeterminacy—as exemplified by the existence of borderline cases—and tolerance: the insensitivity of a predicate to sufficiently small changes in the objects of predication (a one-foot increment of a walking distance is a walking distance); tolerance is the source of the Sorites paradox. In each of these subjects, the philosophical analysis is followed by formal implementation. I have tried to organize and present the material, so as to make the main philosophical points accessible, without requiring technical sophistication.

The considerable literature on vagueness has bundled together the existence of borderlines and the Sorites paradox. This is understandable, given that most standard examples display both indeterminateness and tolerance, and there is, indeed, a connection. Yet, I will argue, these are two separate aspects. It is not difficult to establish by direct analysis, and by using examples from natural language, that vagueness *per se* (that is, indeterminacy that is manifested in borderline cases) does not imply tolerance. The non-implication from tolerance to vagueness is more difficult, since, in natural language, tolerance is manifested in predicates that are vague. Nonetheless, an analysis, aided by a specially constructed example, will show that, in principle, tolerance need not imply vagueness. Tolerance turns out to be a contextual phenomenon, whose rigorous reconstruction requires a framework of *contextual logic*, one that provides explicit representation of certain contexts attaching to declarative sentences. Within this logic we can fully maintain tolerance without being led to the Sorites contradiction. I will propose such a formal system: a semantics and a sound and complete deductive system, which extends first-order logic and fully preserves the classical rules, and I will show how to apply it to vocabularies that include tolerant predicates, such as ‘poor’, ‘walking distance’, ‘noonish’, etc.[1] [2]

---

* Most of this work represents research that has been carried on since 1996. The first versions of contextual logic were presented in two meetings of a joint workshop on vagueness, held in 1996 at NYU and Columbia University. It was also presented in the 1997 fall meeting of the New York Conference on Science and Methods. The system, in its present revised and simplified form that fully preserves classical logic, was presented in an invited talk at the 2001 annual meeting of the Association of Symbolic Logic, and its ideas are sketched in [Gaifman 2001]. The results on borderlines, higher order vagueness, and *KTB* represent more recent research done in the summer of 2001 and discussed in my seminar on vagueness, in the fall of 2001 at Columbia University. I wish to thank my colleague Achille Varzi and the participants of this seminar for illuminating discussions.-

[1] D. Raffman [1994] has proposed a contextual approach to the Sorites. My paper is independent of her work, which came to my attention after the presentations of my preliminary results in 1996. Raffman’s approach is psychologistic, rather than semantic, and does not involve a logical system. See section 5 for further comments on the difference.

[2] I should also mention here Kamp’s work [1981], of which I became aware in 1997 and which is an unfinished attempt at a logical system for handling the Sorites. His approach would not have lead to the present system. I think it may succeed on the sentential level but would fail when it comes to quantifiers.

The full system incorporates two machineries: one, which takes care of borderline phenomena, including higher order vagueness, and another, which handles tolerance by means of contextual logic. In both of these, the material presented here is, philosophically and technically, new. The treatment of tolerance via the framework of contextual logic is the part on which I have been working longer (cf. the footnote to the title). The handling of borderlines and higher order vagueness follows the known strategy of using a modal definiteness operator; but the interpretation of the operator and the general approach are new. This topic needs further elaboration and the working out of further technical details; yet, the outline presented here gives quite an adequate idea of the full system.

Underlying the treatment of both aspects is a positive conception of the phenomena in question: Borderline vagueness is not mere absence of semantic determination, but recognition, shared by competent speakers, that a certain divergence in usage is *legitimate and to be expected*. In the same vein, tolerance is not merely the overlooking of small differences, but a rule that is constitutive of meaning, which mandates that *small differences should not matter*.

Among other things, the topic illustrates how, taking off from familiar linguistic practices, philosophical reflection (and formal modeling) can create questions, for which there is no fact of the matter, or some "philosophical fact", to be found. With the full reinstitution of metaphysical questioning, the door has been opened to the application of concepts beyond their domain of significance and to the creation of "problems" out of conflicting metaphors, without sufficient backup. A critical reining in would do some good. This does not mean a return to the discredited doctrines of positivism or some ordinary-language school of philosophy. But the study of vagueness could do with a dose of down to earth behaviorism.

1. Overview

In ordinary usage, 'vagueness', is a broad term that covers an assortment of loosely connected linguistic phenomena: imprecision, fuzziness, ambiguity, obscurity, lack of specificity (hence the expression 'vague generalities'), and their like. Some of the overlap has crept into the early philosophical literature on the subject.[3] But nowadays, the term is used in philosophy to denote indeterminateness, exemplified cases where the semantic rules seem to leave it open whether some particular object falls under a given concept, or, in general, whether a given description applies to a certain situation. Apparently, the unclarity is not due to lack of factual knowledge, but to a semantic gap: the semantic rules do not (or do not seem to) yield a determinate answer. The indeterminateness, moreover, does not mean that the question is out of order (as would be a category mistake, like asking whether number 3 is happy); the question is appropriate, and unambiguous, but the semantics does seem to decide it.

Philosophers have also subsumed under vagueness the phenomena of *tolerance*: the insensitivity of a predicate to very small changes in the objects of which it can be meaningfully predicated.[4] Someone who is a child at any given time point is also a child

[3] E.g., Russell [1923] conflates vagueness with ambiguity. This, however, is not a confusion steming from loose usage, but a result of Russell's attempt to define vagueness within his metaphysical framework.

[4] The term 'tolerance' has been used in engineering to mark the amount of permissible deviation from sharply specified values. It was also used by Zeeman [1961] in a somewhat related sense (I thank Peter

one minute later; a heap of sand continues to be a heap, if one grain is removed, etc. The Sorites paradox, which is based on tolerance, has been classified in the literature as a "paradox of vagueness". The paradox, let us recall, consists in the existence of a finite sequence of objects, such that (i) the first clearly falls under *P*, (ii) the last clearly does not, but (iii) the difference between any two successive objects is so small that, by tolerance, if one falls under *P*, so does the other. Applying instances of (iii) sufficiently many times, we can deduce from (i) that the last object falls under *P* (e.g., a twenty years old person is a child). I shall henceforth refer to such sequences as *Sorites chains.*

Some philosophers have taken the Sorites contradiction as grounds for radical conclusions. Dummett [1975] has argued that tolerant predicates infect natural language with inconsistency, implying that tolerance cannot be accommodated within a coherent semantics. Unger [1979a, 1979b] has voiced even a more extreme position.

The first main part of the paper, consisting of sections 2 and 3, is devoted to borderline logic; the second part, sections 4–8, is devoted to tolerance and contextual logic. The two are independent and can be read separately. Section 9, which concludes the paper, outlines the combined system, containing the operators needed both for vagueness (and higher order vagueness) and tolerance. Let me go first over the second part, which is the larger part of the paper.

Section 4 presents criticisms of the common approaches to the Sorites, and the arguments for the view, advocated here, that tolerance is a norm governing tolerant predicates, which is constitutive of their meaning. We begin to see at that point how contexts enter the picture. Section 5 presents, in non-technical manner, the basic ideas underlying the contextual approach, including the crucial notion of *feasible context* (roughly, a context is feasible if it does not contain objects that form a Sorites chain). Since tolerance is imposed only in feasible contexts, we can maintain it without running into contradiction. This section contains the argument that, in principle, tolerance need not imply vagueness.

Section 6 presents **CL,** the formal system of first-order contextual logic, and some of its basic properties. Section 7 shows how the well-known examples of tolerant predicates are axiomatized in **CL,** via axioms that are universal sentences. Also expressible in **CL** is the notion of a feasible context, using which the tolerance condition is expressible as well. Section 8 introduces the notion of a *feasible sentence* **(**roughly, a sentence whose semantic evaluation requires only feasible contexts), and shows that these sentences constitute a deductively self-sufficient portion of the language. It is within this portion that ordinary reasoning involving tolerant predicates is conducted.

Turning to the first part, section 2 takes a closer look at borderline phenomena and argues, using diverse examples, that there is no implication from vagueness to tolerance. Section 3 focuses on higher order vagueness. The proposed analysis shows that underlying this phenomenon are certain contextual factors, which are however of a different kind than the one underlying tolerance. Here, little or nothing is gained by explicit representation of the contextual elements. A systematic treatment of higher order vagueness is best achieved within a modal framework, based on a 'definitely' operator,

---

Freyd, of the mathematics department at Penn. for bringing this to my attention). In the context of vagueness the term was introduced in [Wright 1975].

which marks a kind of semantic necessity. The most reliable criteria for the suitability of a modal system for higher order vagueness are certain properties of iterated modalities. An analysis of these properties suggests *KTB* as a suitable modal framework—more suitable, in any case, than each of the other combinations of the standard axioms. Such a framework makes it possible to derive, in a purely modal setting, a scale of degrees for the predicate in question (a "degree" represents the extent to which the predicate is true of a given object).

2. Borderlines and Tolerance

Semantic indeterminateness can be manifested in two ways: (1) hesitation on the part of the speaker, which does not derive from lack of factual knowledge,[5] and (2) divergence in usage among competent speakers—in situations in which they are competent judges—including, possibly, the same speaker on different occasions. These are not intended as reductions of semantic indeterminacy to surface phenomena, but as concrete expressions of it. (1) and (2) are related: one's hesitation can signify one's recognition that divergent answers are legitimate. Having to decide, one answers this way or that, yet one is aware that a different answer cannot be ruled out on semantic grounds. Items that give rise to this indeterminacy constitute the *borderline* region ('penumbra' is another term current in the literature).

The very terminology—'borderline', 'penumbra'—indicates an additional element at play; these are geometrical metaphors that suggest distance. The problem of classifying borderline cases can be seen as the problem of determining whether an object is sufficiently near, or sufficiently similar (in the relevant respects) to some other object, so as to warrant their classification in the same class. And from this there seems to be a natural link to tolerance. It would thus appear that tolerance is an aspect of the borderline phenomenon, and that a suitable analysis of the latter should suffice to dissolve the Sorites.

That this is far from true will become clear when we uncover (in sections 4 and 5) the contextual element that underlies tolerance. For the moment let me elaborate a more obvious point: indeterminateness does not imply tolerance. That conceptually indeterminateness *per se* does not imply tolerance is clear. There need not be an implication even when the former involves some notion of distance. Here are some of the many examples that testify to the claim.

First, many cases of indeterminacy lack the kind of scaling that tolerance, in any appreciable sense, requires. The classification of a newly discovered object, or a new situation, can be indeterminate, because the item lacks sufficient paradigmatic features, or because it combines paradigmatic with anti-paradigmatic ones. Legal cases are often of this sort. To cite a well-known hypothetical example from Hart,[6] imagine a city ordinance that prohibits the operation of vehicles in a public park. It is not clear whether motorized skateboards fall under 'vehicle', in the sense used in the ordinance. Motorbikes obviously do, baby strollers obviously do not. Motorized skateboards share relevant

---

[5] By this I mean the usual kinds of facts that a competent speaker will find relevant; e.g., a person's age—if the predicate is 'old', or 'child', a person's height—if the predicate is 'tall', etc. It is not supposed to include the unknown "facts", which the epistemic view posits.

[6] H.L.A. Hart, *The Concept of Law* (1961). I am indebted to my student Robert T. Miller for this example.

features with the former, but fall short in respect of bulk and speed. And it is not clear whether they endanger or inconvenience people, to an extent that justifies their banning. A motorized skateboard can be seen as a borderline case: more of a "vehicle" than a baby stroller, less than a motorbike. Here there is no tolerance to speak of.[7] (That the vagueness in question will disappear upon the court's decision—if and when the court decides—is irrelevant for the issue.)

Other staple examples come from taxonomy. An organism may baffle its classifier, by possessing combinations of "incompatible" features. The platypus is a celebrated case; like a bird it has a bill, and like birds and a reptiles, it lays eggs; it has also other reptilian features. But the females lactate and suckle the newly hatched "cubs"—a fact that has determined its present classification as a mammal. Here the court has decided. But originally it formed a borderline case, and even nowadays it is put in a special subclass of mammals, the *monotreme*, which it shares with two other living, similarly "strange", species of *echidna*.[8] One should also note that numerous everyday adjectives (e.g., 'generous', 'courageous', 'smart', 'loyal', 'lazy', 'depressed', 'happy') which have borderlines, have only some loose tolerance that does not lead to well-defined Sorites chains.

Most important, many vague predicates, which are completely determined by numerical values and which can display higher order vagueness, lack any appreciable tolerance. Consider 'large number of fingers', as in 'You can perform the trick, without using a large number of fingers', or 'A large number of John's fingers were infected by the fungus.' Unlike 'most fingers', this predicate (abbreviated henceforth as LNF) is vague. Obviously, 9 and 10 are LNFs; 1 and 2 are not; but is 6? is 5? Given that the range is from 0 to 10 and that changes are in discrete units, there is hardly space for tolerance; plausibly, one finger can make a difference: n+1 can be an LNF, while n is not. Or consider, 'Only a small fraction of the committee is Republican', where the committee consists of eight people; 1 qualifies as a small fraction, 4 obviously does not, but 3 is a borderline case. Even this limited example can give rise, as we shall later see, to intricate patterns of higher order vagueness. Or consider 'Large number of siblings', as in 'Mary has a large number of siblings'; taking western, college-educated, present day parents as a reference group, the range is quite limited. Unlike the previous examples, the range does not have a precise upper bound; but this is of no concern: the predicate is vague but not tolerant. 'A large number of divorces for one person' (7 certainly qualifies, does 4? does 5?), 'A graduate, philosophy seminar with many enrolled students' (the range at Columbia is roughly 3 to 15), 'A large family group of gray wolves (rough range, 2-10), or 'A large pride of lions' (rough range, 4-24). The last may perhaps display some tolerance, and tolerance is certainly appreciable in 'A large community of chimpanzees' (rough range 35-65). As the range becomes larger with respect to the unit of change, there is more scope for tolerance; not because differences become less "discernible"—these are not perceptual predicates and the difference between 40 and 41 chimpanzees is as clear-cut as that between 5 and 6 fingers—but because in the contexts in which the predicate is used the difference can be practically ignored. I shall have more to say on this when I take a closer look at tolerance. Let me focus here on borderline phenomena.

[7] A physical configuration of midsize bodies can be transformed into any other through a sequence of many small changes. In particular, the continuous removal and addition of small chunks of matter can transform a skateboard into a bicycle. Obviously, this is not the kind of Sorites chain that is relevant to our example..

[8] I thank Laura Franklin for calling my attention to these biological details.

3. Semantic Modalities, Borderlines and Higher Order Vagueness

Philosophers have considered a "definiteness operator", which is supposed to characterize the clear-cut cases of a vague predicate. Let $\Delta(P(a))$—which reads: 'definitely $P(a)$'—mean that a is a clear-cut case of P. Then, as Fine has suggested,[9] borderline cases can be characterized as neither definitely P, nor definitely non-P:

(B) $$\neg\, \Delta(P(a)) \wedge \neg\, \Delta(\neg P(a))$$

The interpretation of the operator, however, has been a matter of dispute and far from clear. Some have objected to the distinction between Φ and definitely-Φ, claiming that the prefixing of 'definitely' does not change the content of a statement. Sometimes the argument is made that being definitely true is nothing more than being true.[10] This is not a happy way of presenting the issue; the difference, primarily, is not between truth and definite truth, but between the truth of 'Φ' and the truth of 'definitely Φ'; e.g., the truth of '*a* is white' and that of '*a* is definitely white'.[11] Conceivably, 'definitely' can be a sort of emphasis, or a mark of confidence, which does not change the informative content. But surely this is not the only function of the term. The test for 'definitely *P*' can be stricter than the test for '*P*': pointing to a patch of unspoiled snow I assert, "this is definitely white, while that"—pointing to a trodden patch—"is not definitely white, but still white." Since one's hesitation may be due to semantic indeterminacy, an expression of certainty can mean that this case does <u>not</u> fall under semantic indeterminacy, i.e., is not a borderline case. And this *does* add informative content. Note that we are not concerned here with analyzing the natural-language meaning of 'definitely'—a highly vague term that can serve to emphasize, to express confidence, and also to narrow a concept. We can appropriate the word as a theoretical term, without undue awkwardness, given that its specific sense in our framework is one of its actual senses. Yet, we still owe an explanation of the meaning of 'definitely' in our system, which is more than an appeal to pre-theoretic intuitions about definiteness. (If these intuitions are our only guide, we are likely to be bogged down in futile arguments.)

We come by the explanation through the characterization of 'definitely' in terms of borderlines: Definitely $P(a)$, just when (i) $P(a)$, and (ii) a is <u>not</u> a borderline case of P. Now a is a borderline case, just when the semantics does not provide grounds for ruling out either '$P(a)$', or '$\neg P(a)$'. The possibility we have to recognize here is that of a competent judge who classifies a as a P, yet recognizes the legitimacy of classifying a as a non-P. In terms of divergence, one can express this possibility by appealing to an idealized community of competent judges: Some competent judges (to whom I belong)

---

[9] [Fine 1975] p, 140.

[10] Williamson [1994, p. 194] argues that this is so, unless 'definitely' is interpreted epistemically..

[11] The "undeniable" intuition that being definitely true is the same as being true, is far from clear. Suppose *a* is a borderline case of 'white'. We can retain the Tarski biconditional:
True ('*a* is white') $\leftrightarrow$ *a* is white, by construing 'True' so that the left-hand side inherits the vagueness of the right-hand side; in this case being true is not the same as being definitely true. Other construals weaken the biconditional and endow 'true' with a "definitely" effect; such is the case, if, subscribing to supervaluations, we construe 'true' as 'supertrue'. In any case, we would do better to separate the treatment of 'true' from that of 'definitely'..

classify *a* as a *P*, but others (whom I acknowledge) classify it as a non-*P*. It is this possibility that 'definitely' is supposed to rule out; definitely $P(a)$ means that $P(a)$ and there are no legitimate semantic grounds for its negation. Wright [1994] took a similar line (if I understand him correctly), but seems to have misconceived an important aspect.[12]

Acquaintance with semantic indeterminacy is part of what it takes to be a fully competent speaker. It is not as if the speaker is simply at a loss; rather, one recognizes (explicitly or tacitly) that the difficulty does not stem from one's limitations. The vagueness of the predicate, as manifested in the existence of borderline cases, is common knowledge. We know that in some cases there can, and will be, legitimate disagreements. This does not imply that we are agreed on *which* exactly are the borderline cases—a question that raises the issue of higher order vagueness—but it is agreed that there are such cases.

Note that we are employing throughout classical two-valued logic. This means that the agent is called upon to make a yes/no decision. The addition of borderlines does not exempt one from this requirement; there is no switch to many-valued logic, or to degrees of truth (the connections with degrees will come out in a short while). But it makes place for classifying certain cases as those in which the speaker recognizes that —for reasons of semantic latitude—a different answer is conceivable. The same applies generally: To say that a sentence is a borderline case is to say that, the relevant facts being known, neither Φ nor $\neg\Phi$ can be ruled out on the basis of the existing semantics. And to say that definitely Φ is to say that Φ and that it is not a borderline case. Note that although the definiteness operator was suggested by Fine in the context of supervaluations,[13] the present approach is definitely not supervaluationist.

Formally, Δ, is a modal operator, which belongs to the family of necessity operators. We can conveniently speak here of *semantic necessity*. There is danger of reading into the term unintended meanings. My use of it should be understood solely in terms of the above explanations; no metaphysical or ontological overtones are intended. Let ∇ be the dual possibility operator, which we can conveniently regard as expressing *semantic possibility* (again, subject to the caution about unintended meanings). Evans [1978] uses '∇' for 'indefinitely', which is different; notational change is in general undesirable, but there is a very good reason for my notation. We need simple suggestive notations for basic dual pairs, such as those of modal logic ($\Box$ and $\Diamond$), or of sentential logic ($\wedge$ and $\vee$).

We can use $\neg\Delta\neg\Phi$ (just as we can rewrite $\Diamond\Phi$ as $\neg\Box\neg\Phi$, or $\Phi\vee\Psi$ as $\neg(\neg\Phi\wedge\neg\Psi)$), but this puts obstacles in the path of our thinking, since it obscures certain basic patterns.

---

[12] He proposes (page. 145) that for a sentence Φ "to be definitely true is for any appropriately generated opinion that [not-Φ] to be cognitively misbegotten". On the other hand he argues (p. 138) that "wherever a stable consensus can be elicited that something is on the borderline between two concepts, that is merely an indication that we could, if we wished, employ a concept intermediate between them and they are not really complementary". This is a wrong picture; the borderline is not "between two concepts", but the region in which legitimate disagreement can arise. If John divides the domain into *P*'s and non-*P*'s, then the borderline, from John's view, consists of those items, for which he grants, as legitimate, the possibility of a different classification. There can be full consensus on what the borderline is, without this impinging in any way the vagueness of *P*. Sorensen's reply to Wright, in the same volume, has the same misconception..

[13] Cf. footnote 9. I do not know if this is the first use of the operator in the context of vagueness.

Let $\mathbb{B}$ be the borderline operator. Each of the operators can serve as a basis for defining the other two, via standard equivalences of modal logic:

$$\text{(M)} \qquad \nabla\Phi \Leftrightarrow \neg\Delta\neg\Phi \qquad \mathbb{B}\Phi \Leftrightarrow \nabla\Phi \wedge \nabla\neg\Phi \qquad \Delta\Phi \Leftrightarrow \Phi \wedge \neg\, \mathbb{B}\Phi$$

If the vagueness of *P* is displayed by *a* such that $\mathbb{B}P(a)$ holds, second order vagueness is displayed by the holding of $\mathbb{B}^2P(a)$, i.e., $\mathbb{B}\mathbb{B}P(a)$. In general, higher order vagueness can be expressed through iterations of the semantic modalities. But there is little chance of getting insight into this matter, if we have nothing to go by, besides some philosophical pictures about the nature of "definiteness" or "borderline" (the fruits of gazing into panoramas of dissolving borders). Let us consider a concrete example.

A teacher has to grade a bunch of exams on a pass/fail basis; when the task is accomplished, she will have two piles: the "passes" and the "fails".[14] Here the negation of "pass" is "fail". She puts some exams, over which she hesitates, in a third, temporary pile, for further deliberation. We can take the third file as the borderline region. Note however that this is the borderline when the only alternatives are "pass" and "fail", not when "borderline" itself is added as a possible grade. Suppose that someone in authority adds "borderline" as a third grade. The teacher is now called upon to produce a classification into three categories. Does it follow that the three previous piles fit smoothly into these categories? Not at all. She might have classified, in the first round, an exam, *a*, as a "pass", because she did not doubt that the alternative, "fail", was less appropriate; but now that she has a third option, she may choose to classify *a* as a "borderline", or hesitate between "borderline" and (the new) "pass". Also, cases that gave her pause in the two-category situation may give her pause also in the three-category one: she hesitated between "pass" and "fail" when these were the only options, but now she wonders whether the third option suits the exam better than, say "pass". Nothing mandates that hesitation in the two-category situation be always translated into clear-cut "borderlines" in the three-category one.

The phenomenon is clearly contextual: the import of classifying an item under a certain category is determined by the other categories available for classification. This is such a trivial truth that I see no need to go into it further. What is not so obvious, and what has been ignored in the philosophical discussions of the topic, is that this trivial truth underlies the phenomenon of higher order vagueness.

It is clear that the new "pass" and "fail" are not the same as the old ones. Since they result by subtracting "borderline", they should be construed as "definite pass" and "definite fail". The initial borderline (the third pile) constitutes the borderline between (the old) "pass" and "fail", when "borderline" was not an option. Higher-order vagueness is vagueness in the context of extending the language by introducing the old borderlines as additional categories. The phenomenon described in [Sainsbury, 1990] as "concepts without boundaries" is the visual effect of a fast moving frame, a repeated-extension process in which the addition of borderline predicates creates new borderlines. Note that this makes it possible for a fixed finite domain of objects to support an infinite order of

[14] I owe Achille Varzi the idea of using a grading situation in order to analyze vagueness.

vagueness; the new borderlines can consist of previous definite cases (as *a* in the example above), and also of previous (lower order) borderlines.

Having noted the discontinuity that results from the addition of options, let us also note the continuities. The new judgments should cohere with old ones, for we are not concerned here with revisions but with an extension of a previous view to a new situation. The teacher cannot classify *a* as a "definite pass" unless she has classified it, in the pass/fail situation, as a "pass"; and she cannot judge *a* as being possibly a "definite fail", unless she has judged it as being possibly a "fail" (i.e., it should be either in the old "fail" pile or in the "borderline" pile). This restriction implies that the modal system should be at least *KT*; i.e., it should satisfy the usual requirements for a *K* system, as well as: $\Delta\Phi \rightarrow \Phi$.

The teacher's recognition that different grade assignments are legitimate can be expressed by imagining other graders, whom she considers competent, who assign different grades. This brings us back to a group-based modeling. If a borderline case is one about which competent speakers can disagree, then the vagueness of 'borderline' can be derived from the vagueness of 'competent speaker', which is vagueness in the metalanguage. Such an interpretation accords with our construal of second order vagueness as one that obtains within an extended language. But practically no insight is gained by analyzing higher order borderlines by means of the metalinguistic hierarchy. The same is also true of the analogous set-theoretic hierarchy proposed by Fine [1975 p. 146].[15] Also dealing explicitly with the contextual parameter (the classification options) is rather tedious and not illuminating. The most perspicuous way of handling higher order vagueness is to encompass, within one modal language, all the generated categories, expressing them in terms of iterated modalities.

In the same way that 'definitely' can narrow a concept by ruling out borderline cases, 'definitely definitely' can effect an additional narrowing: Something is definitely definitely *P*, just when it is definitely *P* and is not a borderline case of 'definitely *P*'. And, in principle, we can iterate further. As a rule, vague predicates are associated with some notion of degree (very often, they are correlated with well defined quantitative scales: age, height, distance, etc.). Degree theories propose to take such a notion as primitive. The idea underlying the present approach is to use modalities in order to get degrees.

[15] The straightforward modal corollary of supervaluations is an *S*5 system, in which there is no place for higher order vagueness. In order to express the higher order phenomenon, Fine uses an iterated supervaluation construction. He defines a 0-order space as a set of (admissible) sharpenings of a vague predicate, a 1-order space—as a set of 0-order spaces, and so on. An *n*-order boundary is defined as a sequence $s^0, \ldots, s^n$, such that each s^i is an *i*-order space and $s^i \in s^{i+1}$. An ω-order boundary is such an infinite sequence of length ω. (We can then go on to define an ω-order space as a set of ω-boundaries, and continue so, through all ordinals.) Boundaries of a given order can be identified with possible worlds. On Fine's proposal, a boundary, $b^0, b^1, \ldots,$ is accessible from another, $a^0, a^1, \ldots,$ iff $b^i \in a^{i+1}$ for all $i <$ the given order. This does not provide a fruitful grip on higher order vagueness. There is no direct relation between the order of the boundary and the order of vagueness. The accessibility relation can turn out to be an equivalence relation, in which case ω-order boundaries do not yield any vagueness of order > 1.

In terms of the possible world semantics, ∇Φ is true in world *w* just when Φ is true in some world accessible from *w*. In terms of semantic modality, the truth of ∇Φ means that holding Φ true is semantically legitimate. It is therefore natural to take possible worlds as representations of semantic views: ways of applying the vague predicates when "yes/no" decisions are required; for the present, they amount to sharpenings. (Later, when tolerance is modeled in contextual logic, they will be models of a different kind.) If *w* represents Mary's view—her way of partitioning the relevant domain into *P*'*s* and non-*P*'s—and *w*′ represents John's view. Then *w*′ is *accessible* from *w*, just when Mary acknowledges John's view as a legitimate way of applying *P*. Mary's acknowledgement of John's view does not, of course, mean that she subscribes to it. In particular, it does not mean that she acknowledges what John acknowledges. Say Mary acknowledges John, and John acknowledges Bill; it need not follow that Mary acknowledges Bill. If Bill holds Φ true, then John holds ∇Φ true, and Mary holds ∇Φ true, but not necessarily ∇Φ (that is, she can hold true ∇Φ and Δ¬Φ). It is also possible that Mary acknowledges John's views with respect to the non-modal part of the language, without acknowledging his views about sentences containing modal operators; in this case, w′ is not accessible from *w*, but a copy of it, say *w*′′, is; the difference between *w*′ and *w*′′ is in the worlds that are accessible from them. In general, the possible worlds represent some hypothetical semantic views, not necessarily views held by actual persons. The model is chosen so as to include the views that achieve certain effects. Later we shall see some examples.

It is well known that the popular system, *S5*, does not accommodate higher order vagueness. In models of *S5* there is a division of the relevant domain into definite *P*'s, definite non-*P*'s and a borderline region; each of these coincides with its definite core, implying (it is not difficult to see) that they are sharp predicates. In standard examples, 'walking distance', 'young', 'tall', and their like, this outcome is counterintuitive (a sharp cutoff for 'definitely young' appears no less arbitrary than a cutoff for 'young'). Yet, in some cases *S5* yields acceptable results. Consider the example of section 2, 'a small fraction of the committee', where the committee consists of 8 members. Let '*P*(*x*)' read as: '*x* is a small fraction of 8', where '*x*' ranges over the nine integers: 0,…,8. Consider three possible worlds, w_1, w_2, and w_3, such that, in w_1, *P*'s extension is {0,1}, in w_2 it is {0,1,2}, and in w_3 it is {0,1,2,3}. If every world is accessible from every world, then, in each of these worlds, Δ*P*, 𝔹*P*, and Δ¬*P* have, respectively, the extensions {0,1}, {2,3}, and {4,…,8}, and each is sharp (i.e., has an empty borderline). Vagueness disappears, if instead of using the predicate *P*, and its negation, ¬*P*, we use Δ*P*, 𝔹*P*, and Δ¬*P*. But we cannot translate statements involving '*P*' into this language; the vagueness of *P*, which consists in the existence of semantically indeterminate cases, is not diminished by the fact that 𝔹*P* is sharp (see also footnote 12). There are others plausible modeling of 'a small fraction of 8', in which, we shall see, the predicate has vagueness of infinite order; there are also, for each $n \geq 1$, modelings in which its vagueness is of order *n*, but not $n+1$.

Higher order vagueness can exist in models of *S4*. But if we adopt the above view that 'definitely definitely' can, in principle, be stronger than 'definitely', then *S4*, which blocks this possibility, should not be adopted. This point is reinforced by the fact that *S4* limits severely the possibilities of higher order vagueness. It can be shown that the following is provable in *S4*, where '$\mathbb{B}^k$' stands for *k* iterations of 𝔹:

$$\mathbb{B}^{2+n}\,\Phi \leftrightarrow \mathbb{B}^2\Phi.$$

This implies that in *S4* a vague predicate either does not have second order vagueness, or has vagueness of infinite order in which all higher-order borderline regions of order ≥ 2 are the same. (I will later indicate how the order of vagueness is defined.)

By itself, modal logic cannot help us in choosing a system for semantic modality. And there is not much intuition to guide us, since the whole setup is a rarefied philosophical exercise; we do not, as a rule, extend our stock of predicates by adding, repeatedly, borderlines of borderlines of borderlines…, or by iterating 'definitely'. There is nonetheless a certain large-scale feature of iterated modalities, which we may find attractive and which suggests *KTB* as a good system for modeling vagueness. *KTB*, recall, is obtained by adding to *KT* the scheme (known as B): $\Psi \rightarrow \Delta\nabla\Psi$, or the equivalent one:

$$\nabla\Delta\Phi \rightarrow \Phi$$

In terms of possible worlds semantics, *KTB* is characterized by an accessibility relation that is reflexive and symmetric: Mary does not recognize John's view as legitimate, unless he recognizes *her* view as legitimate. Plausible enough? Perhaps. The real argument for *KTB*, derives from the behavior of iterated modalities.

To simplify, consider first a sentence ψ. The three mutually exclusive sentences:

$$\Delta\psi \qquad \mathbb{B}\psi \qquad \Delta\neg\psi$$

constitute a rough scale: the truth of $\Delta\psi$ marks a higher "truth degree" for ψ, than the truth of $\mathbb{B}\psi$, which, in its turn, marks a higher degree than $\Delta\neg\psi$. This order is reversed for $\neg\psi$ (note that $\mathbb{B}\psi$ is equivalent to $\mathbb{B}\neg\psi$). We can, if we wish, subdivide $\mathbb{B}\psi$ into $\mathbb{B}\psi\wedge\psi$ and $\mathbb{B}\psi\wedge\neg\psi$ (for certain purposes it is preferable not to include that division). For any k>1, this rough scale can be refined by subdividing $\Delta\psi$ according to additional "markings": $\Delta^k\psi$, $\Delta^{k-1}\psi$,..., $\Delta^2\psi$, where the exponents denote iterations of the operator; we can do the same for $\neg\psi$. The mutually exclusive sentences, representing the resulting "intervals", give us a scale of *degree sentences*:

$$\Delta^k\psi,\quad \neg\Delta^k\psi\wedge\Delta^{k-1}\psi,\quad \ldots, \neg\Delta^2\psi\wedge\Delta\psi\,,\ \ \mathbb{B}\psi\wedge\psi,\quad \mathbb{B}\neg\psi\wedge\neg\psi, \ldots, \neg\Delta^k\psi\wedge\Delta^{k-1}\neg\psi,\ \ \Delta^k\neg\psi$$

Call a scale obtained by using all or some of these markings, a *standard scale*. In the illustration below, all the steps are carried out, for k = 3. In general, the full scale contains 2(k+1) degree sentences, each of which uses a ≤ k nesting of modal operators.

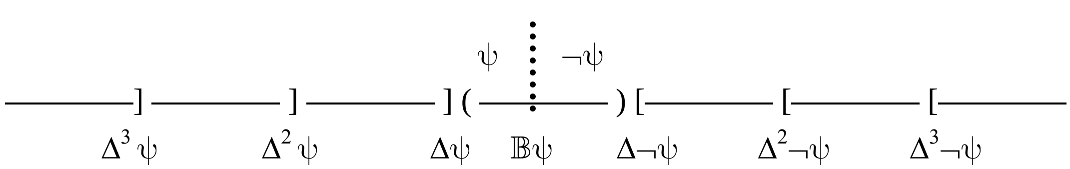

Assuming *KTB*, we can partition each degree of a standard scale as follows. Add, inside the $\mathbb{B}\psi$ region, two scale marks, for ŏ$\Delta\psi$ and for ŏ$\Delta\neg\psi$; it is not difficult to see that the remaining part of $\mathbb{B}\psi$ is $\Delta\mathbb{B}\psi$:

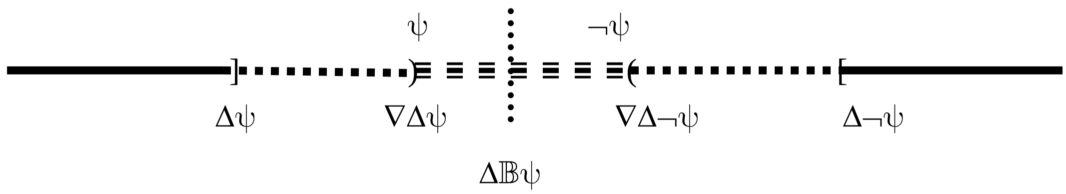

And in each of $\neg\Delta^j\psi\wedge\Delta^{j-1}\psi$, add a scale mark for $\nabla\Delta^j\psi$:

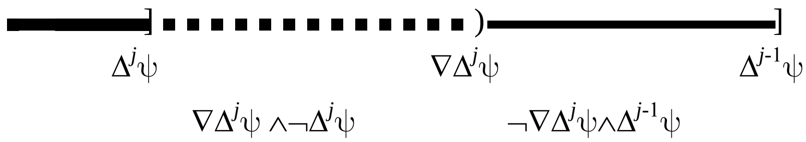

Do the same for $\neg\psi$ (producing the mirror image of the above diagram). Call a scale obtained by some or all of these subdivisions of standard degree sentences, a *one-refinement scale.* The one-refinement scale is based on the provability of certain sentences in the modal system. The provability of $\nabla\Delta\psi\rightarrow\psi$, and of $\nabla\Delta\neg\psi\rightarrow\neg\psi$, which are instances of axiom B, implies that the added markings in the center fall as indicated; other sentences are split as indicated, because $\nabla\Delta^j\psi \rightarrow \Delta^{j-1}\psi$ is provable. Without axiom B, we are not guaranteed that the new markings are even comparable to the old ones. Other refinements are possible, but they require specific theorems in the modal theory governing the vocabulary of ψ, which, it appears, are not instances of modal axioms of "nicely behaved" modal logics. In general, we can get refinements by using alternating blocks of Δ's and ∇'s. Their behavior—even in *KTB*, if we go beyond the one-refinement scale—can be quite complicated. The subject is beyond this paper's scope and I hope to address it elsewhere. A degree sentence, ϕ, can "collapse" (be empty); this happens when $\neg\phi$ is a theorem of our theory; e.g., if $\Delta^2\psi \rightarrow \Delta^3\psi$ is a theorem, then the degree sentences $\neg\Delta^2\psi\wedge\Delta^3\psi$, and all of its sub degrees are empty.

Not every sentence can serve as a degree sentence; for example, $\Delta\psi \vee \Delta\neg\psi$ cannot. Intuitively, a degree sentence should correspond to an "interval" in the above picture. This notion can be rigorously established, but I shall not pursue the topic here. Further indications will be given shortly when we consider predicates.

In the degree sentences just considered, 'ψ' plays a schematic role. Abstracting away from ψ, we can identify the degrees themselves as operators obtained by iterating modal operators, in combination with sentential connectives. In modal logic a *modality* is any sequence of monadic operators, including negation (e.g., $\Delta\nabla\nabla\neg\Delta$). For our purpose we have to generalize this by allowing also binary sentential connectives. First let the "empty modality", $\mathbb{E}$, be defined by:

$$\mathbb{E}\psi =_{Df} \psi.$$

Let a *generalized modality* be any operator obtained via the following recursive definition.:

- $\mathbb{E}$ is a generalized modality.
- If M and M' are generalized modalities, so are $\neg M$, ΔM and $M \wedge M'$, defined by:

$(\neg M)\psi =_{\mathrm{Df}} \neg(M\psi)$ $\quad$ $(\Delta M)\psi =_{\mathrm{Df}} \Delta(M\psi)$ $\quad$ $(M\wedge M')\psi =_{\mathrm{Df}} (M\psi) \wedge (M'\psi)$.

This determines in the obvious way other modalities, such as ∇M, or $M \vee M'$. As noted, only certain generalized modalities can serve as degrees. A set, $\boldsymbol{D}$, of *potential degrees*—consisting of generalized modalities that, conceivably, may serve as degrees—can be defined as follows: Let $\boldsymbol{D}^+$ be the smallest class containing $\mathbb{E}$, such that if $D, D' \in \boldsymbol{D}^+$, then, ΔD, ∇D, $D\wedge D'$, $D\vee D' \in \boldsymbol{D}^+$. Let $\boldsymbol{D}^-$ be the smallest class containing $\neg\mathbb{E}$, satisfying the same closure conditions. (The members of each class are logically equivalent to the negations of the members of the other). Then $\boldsymbol{D}$ consists of all members of $\boldsymbol{D}^+$, all members of $\boldsymbol{D}^-$, and all conjunctions $D^+ \wedge D^-$, where $D^+ \in \boldsymbol{D}^+$ and $D^- \in \boldsymbol{D}^-$. For a given $D \in \boldsymbol{D}$, the degree (or truth degree) of a sentence ψ is D, if $D\psi$ is true. It is easily seen that all degrees of the standard and one-refinement scales are among the potential degrees.

Now, a scale is not very significant, if the degree sentences have large "margins of error", that is, if they have broad borderline regions. Assume, for example, that exams are graded on a scale of 1 to 10, where *a*'s grade is the degree of '*a* is a successful exam'. Getting grade 9 means little if '*a* is a grade 9 exam' is included in the borderlines of grades 8,7,6,5; because grade 5 would also count as a legitimate grade. In general, the borderlines of degree sentences can get out of hand. The virtue of *KTB*, and the main argument for its adoption for modeling vagueness, is that it provides a nice measure of control over the borderlines—at least for standard and one-refinement scales. Here is the relevant theorem:

Let $\phi_1, \phi_2, \ldots, \phi_{i-1}, \phi_i, \phi_{i+1}, \ldots \phi_n$ be either a standard or a one-refinement scale of degree sentences, for the sentence ψ. Then, assuming *KTB*, we have for all i:

(i) For a standard scale, $\phi_i \to \neg\mathbb{B}\phi_j$ is provable, for all $j < i-1$, or $j > i+1$.

(ii) For a one-refinement scale, $\phi_i \to \neg\mathbb{B}\phi_j$ is provable, for all $j < i-2$, or $j > i+2$, except, possibly, for the case where $\Delta\mathbb{B}\psi$ is split into $\Delta\mathbb{B}\psi\wedge\psi$ and $\Delta\mathbb{B}\psi\wedge\neg\psi$ and where $\phi_i = \psi\wedge\nabla\Delta\psi$, $\phi_j = \phi_{i+3} = \neg\psi\wedge\nabla\Delta\neg\psi$, or where $\phi_i = \neg\psi\wedge\nabla\Delta\neg\psi$ and $\phi_j = \phi_{i-3} = \psi\wedge\nabla\Delta\psi$.

(In these claims empty degrees count as well; e.g., if, in case (i), ϕ_{i+1} is empty, then $\phi_i \to \neg\mathbb{B}\phi_j$ is provable for all $j > i$.)

This means that, in a standard scale, the margin of error is no more than one degree (because a degree is disjoint from the borderlines of degrees at a remove than more than one degree). For a one-refinement scale, the margin of error is no more than two degrees, except that in the case indicated above it can be three. (The exception disappears if we lump the two central degrees $\Delta\mathbb{B}\psi\wedge\psi$ and $\Delta\mathbb{B}\psi\wedge\neg\psi$ into $\Delta\mathbb{B}\psi$).

Standard and one-refinement scales are special types that behave nicely under *KTB*. This behavior is what recommends both these scales and the adoption of *KTB*. Apparently, none of the other customary modal systems provides for systematic well-behaved scales. I therefore conclude with a cautious endorsement of *KTB*.

The application of degrees to predicates is straightforward. Given a predicate *P* and a generalized modality *M*, *MP* is the predicate defined by:

$$(MP)(x) =_{\mathrm{Df}} M(P(x)).$$

The *degree predicates* for *P* are the predicates *DP*, where *D* ranges over some scale of degrees. For some fixed scale, the extensions of these predicates form a partition of the domain over which *P* is defined. The diagrams above become illustrations of partitions, ordered according to the "*P*- degree" of the objects contained in the part.

Cases like 'tall' or 'walking distance', that come with a given ordering, or pre-ordering, ≤, of the domain,[16] yield insights into degree predicates. Very often, the predicate *P* is either monotone or anti-monotone with respect to ≤: monotonicity (exemplified by 'tall') means that '$x \leq y \rightarrow (P(x) \rightarrow P(y))$' is an axiom (or theorem) of the background theory; anti-monotonicity (exemplified by 'walking distance') is monotonicity with respect to ≥. Say that a predicate *S* (over the given domain) has the *interval property*, if it is provable, in the background theory, that *S*'s extension is an interval under ≤ :

$$x \leq y \leq z \rightarrow (S(x) \wedge S(z) \rightarrow S(y))$$

It should be obvious that, for a monotone or anti-monotone *P*, the interval property is a necessary condition for being a degree predicate. It can be shown that if *D* is a potential degree, then for *P* that is either monotone or anti-monotone, *DP* has the interval property. Note that $\mathbb{B}$ is a potential degree, but $\mathbb{B}^k$, for $k > 1$, is not; for, in general, the extension of $\mathbb{B}^k P$, for $k > 1$, is a union of separated intervals. Following these clues, we can construct scales, from appropriately chosen potential degrees, without presupposing an ordering of the domain. There is more to the story, but, as I said, it has to be told elsewhere.

The upshot of the above is a derivation of degrees, within classical two-valued logic, as a byproduct of the modal system. So far, we have considered degrees of atomic sentences. We can, however, apply generalized modalities that are potential degrees to sentences in general. The downside of the approach is that, due to the complexity of the system, finding the resulting patterns can be quite difficult. In actual cases, we appeal of course to some given ordering (or preordering) of the domain; we do not derive degrees via generalized modalities. But it is philosophically significant that, in principle, degrees can be established solely on the basis of the definiteness operator.

[16] A pre-ordering is a reflexive and transitive relation over the given domain, which is total (any two objects are comparable). We do not require, as we do in the case of an ordering that $a \leq b$ and $b \leq a$ imply $a = b$. For example 'less than or equally tall' defines a pre-ordering of the class of people (different people can have the same height).

The *order of vagueness* can be defined as follows. A sentence ψ has *vagueness of order (at least) k*, if, for some $n \geq k$, $\mathbb{B}^n\psi$ is true. (We cannot simplify this to the condition that $\mathbb{B}^k\psi$ is true, because it is possible that $\mathbb{B}^k\psi$ is false, but $\mathbb{B}^n\psi$ is true, for some $n>k$.) Vagueness of *exact* order k is defined as vagueness of order k, which is not of order k+1. When it comes to predicates the condition for vagueness of order (at least) k is: For some $n>$k, $\mathbb{B}^nP$ has non-empty extension. Exact orders are derived from this as above. A sentence, or predicate, has vagueness of infinite order if it is vague of order k, for all finite k. As observed, $\mathbb{B}^n$, for $n>1$, is not a potential degree; but there is a scale that, roughly speaking, is obtained from the partition induced by $\mathbb{B}^n$ (i.e., from the mutually exclusive degrees whose disjunction is equivalent to $\mathbb{B}^n$). This scale has some relation to the standard and one-refinement scales, but does not appear to have their nice properties. .

The present system handles both higher order vagueness and degrees, in terms of Δ only, paying the price of technical complexity. If we are prepared to ignore higher order vagueness, we can opt for *S5* and we can introduce degrees as an additional component, through a probability measure on sets possible worlds. The degree of a sentence is then the measure of the set of possible worlds in which it is true. This approach is compatible with the general framework of semantic modality, where indeterminacy is interpreted in terms of legitimate disagreement.[17]

3.1 Sharp Models for Higher Order Vagueness

Consider again the case where '$P(x)$' reads: 'x is a small fraction of 8', with 'x' ranging over 0,…,8. For any predicate, S, let $|S|$ be the extension of S. Let w_1, w_2, w_3, be the possible worlds defined earlier (in which $|P|$ is, respectively, {0,1} {0,1,2} and {0,1,2,3}). If every world is accessible from every world, we get the *S5* model discussed earlier. Let the model $\mathcal{M}$ be obtained from it by a slight modification of the accessibility relation: every world is accessible from itself, w_1 and w_2 are accessible from each other, w_2 and w_3 are accessible from each, and there are no other accessibilities. Then, in w_1, we have: $|\mathbb{B}P| = \{2\}$, $|\mathbb{B}^2P| = \{3\}$, $|\mathbb{B}^3P| = \{2\}$, $|\mathbb{B}^4P| = \{3\}$, and so on, flip flopping *ad infinitum*. In w_3 we have a similar flip flopping starting with {3}: $|\mathbb{B}P| = \{3\}$, $|\mathbb{B}^2P| = \{2\}$, $|\mathbb{B}^3P| = \{3\}$, etc. And in w_2 we have: $|\mathbb{B}^nP| = \{2,3\}$, for all $n> 0$. Each of these patterns, as well as the *S5* model, can be fitted with some plausible story. Take for example w_3 in $\mathcal{M}$. Assume the following scenario. I hesitate about 3, decide at the end that it is a small fraction of 8, but acknowledge as legitimate the view that it is not. On any view 4 is not a small fraction of 8; also, I do not acknowledge a view on which 2 is not a small fraction. Hence, on my view, $|\mathbb{B}P| = \{3\}$. When $\mathbb{B}P$ is given as an additional option, I classify 3 under it

[17] Again, full classical two-valued logic is retained. The possible worlds correspond of course to the (admissible) sharpenings; but there is no identification of truth with supertruth.

without hesitation, but hesitate whether 2 should belong there as well; I decide it does not, but acknowledge as legitimate a view on which it does: $|\mathbb{B}^2P| = \{2\}$. When the options are further extended by adding $\mathbb{B}^2P$, I classify 2 under it without hesitation, but hesitate about 3. And so on. The full account requires going through the whole model, but the above story can help in making sense of the pattern.

Small changes in the model can produce radical changes in the pattern: Let $\mathcal{M}'$ be obtained from $\mathcal{M}$ by adding a possible world w_0, in which $|P| = \{0\}$, and linking it, via symmetric accessibility, to w_1 (and to itself). Then in every world of $\mathcal{M}'$ we get exact third order vagueness. In w_2 the extensions of various predicates are:

$|P| = \{0,1,2\}$ $|\neg P| = \{3,4\ldots 8\}$ $|\Delta P| = \{0,1\}$ $|\Delta\neg P| = \{4,\ldots 8\}$
$|\Delta^n P| =\{0\}$, for all $n>1$, $|\Delta^n\neg P| =\{4,\ldots 8\}$, for all $n> 0$, $|\mathbb{B}P| = \{2,3\}$

$|\nabla\Delta P| = \{0,1,2\}$ $|\nabla\Delta\neg P| = \{3,4,\ldots 8\}$ $|\Delta\mathbb{B}P| = \varnothing$ $|\mathbb{B}^2 P| =\{1,2,3\}$

$|\nabla\Delta^2 P| = \{0,1\}$ $|\nabla\Delta^2\neg P| = \{3,4,\ldots,8\}$ $|\Delta\mathbb{B}^2 P| = \{2\}$ $|\mathbb{B}^3P| = \{1,3\}$

$|\nabla\Delta^3 P| = \{0\}$ $|\nabla\Delta^3\neg P| = \{4,\ldots,8\}$ $|\Delta\mathbb{B}^3P| = \{1,3\}$ $|\mathbb{B}^nP| = \varnothing$, for all $n > 3$.

My student Jonathan Simon has shown that, in the case of a monadic vague P defined over some set of numbers, we can get, for each n, vagueness of exact order n as follows. Let the model consist of 2^n possible worlds, $w_1,\ldots, w_{2^n}$, with a reflexive symmetric accessibility relation that links them in chain: w_k and w_{k+1} are accessible from each other for all $k < 2^n$; let the first half of that chain consist of copies of the same world, and the second half—of copies of a different world.

There is a philosophical perception that a faithful rendering of vague terms must be given in a vague language.[18] Yet, a vague picture can be revealed, under powerful magnification, to consist of a sharply defined collection of sharp pixels; conceivably, one can point out the features of this collection that are responsible for the picture's vagueness, and even suggest quantitative parameters for measuring it. The role of a sharp modeling of the borderline phenomena is analogous: the goal is not to produce "equivalent vagueness" but to reveal how the mechanism of vagueness works.

As a rule, precise models do not match concrete situations precisely. There is always some slack between the idealized drawing and the rough actuality. This is true in general; very much so in the case of vagueness, where the slack can be large indeed. The mismatch between the model and the linguistic phenomena can be perceived as

[18] Williamson (1994, p.191) observes that if vagueness is not ignorance, then the semantics of a vague language must correlate with vague statements vague propositions; hence the metalanguage must be vague. This argument overlooks the fact that our modeling need not lead to translations that produce "the same propositions". Its primary aim is to give us insight by providing a systematic general way of assigning truth-values. In the possible world semantic, 'necessarily Φ' is true just when Φ is true in all the worlds that are accessible from the actual world. But it takes a lot of stretching to say that 'necessarily Φ' and 'Φ, in every world accessible from *this* world' express the same proposition. (for one thing, it would require modal realism à la Lewis). If anything, Williamson's argument shows that "sameness of propositions" is not a good notion for describing what a semantic account does.

vagueness on a higher level. But we should separate the vagueness that is being modeled from the imprecision of the modeling. It is well to recall here that vagueness is an *inside* phenomena, acknowledged by the speakers, which enters into the meaning of the vague terms. Let us not confuse it with the slack that arises when the theoretician applies precise tools to imprecise phenomena. Also various parameters of the model—such as the exact order of vagueness—are underdetermined by ordinary usage. Each of the three worlds w_1, w_2, w_3, in the above model, $\mathcal{M}$, as well as the worlds in the *S5* model, yields a plausible semantic account of ‘a small fraction of a committee of eight’; but I cannot tell which of them represents better my own usage. For my own idiolect is not that clear-cut.

The most we can aspire is to produce the right sort of picture, one that captures the essential mechanism and yields plausible patterns. Questions about the “right” model, or the “correct” order of vagueness become meaningless if pushed too far. Such questions pertain to phenomena arising in repeated extensions of our language, through iterations of ‘the borderline of …’, a process we do not engage in, except in some sort of philosophical make believe; they cannot be decided by appealing either to actual practice, or to “the world”. It is therefore quite misleading to speak, as philosophers sometimes tend to, as if certain questions about higher order vagueness have matter of fact, or philosophically determined answers. A philosopher might extrapolate from familiar intuitions. For example, there is no sharp cutoff in the number of minutes that qualifies as ‘young age’, and the same seems to hold with respect to ‘borderline of ‘young age’ ’; it might appear that this should persist all the way, for ‘borderlines of borderlines of borderlines of …’. There is nothing inherently wrong with such a picture as long as we realize its source and its arbitrary aspect (and can you really be sure that this is how one should decide after 100 iterations of ‘borderline’? Or after 10.000?).

It would be a misconception to push the vagueness phenomenon into the sharp modeling. Confronted with vagueness of all finite orders, philosophers have considered the predicate ‘an *n*-order borderline case, for all finite *n*’; should not such a predicate be vague? And should not some object, say *a*, be in the borderline of *that* predicate? Sure, if you want, you can continue the game; you can go on using additional operators and a more complex modeling. It is an illusion however to think that some discoveries—linguistic, metaphysical, philosophical—are being made in this way. The only discovery is the technique, by which models that deliver certain effects can be constructed.

Here an analogy might help. We can define, in a purely extensional language, a possible-world semantics for a modal system. It would be a mistake to apply the notions of possibility and necessity, which are being modeled to the statements of this extensional language (“Is it *necessary* that the number of possible worlds is such and such? *Could* the accessibility relation be different from what it is?”) Similarly, we can define a semantics for an inherently vague language in a sharp metalanguage. And it would be a mistake to use the machinery of that language in order to generate additional vagueness.

A similar point applies to the modeling of tolerance, carried in the coming sections. Like the modeling of vagueness, many details of the model are underdetermined by the modeled phenomena. I shall remark on this, and on the fact that concepts used in the modeling appear themselves to be tolerant, at the end of section 9.

4. A Closer Look at Tolerance

A predicate is tolerant, in as much as sufficiently near objects are classified alike with respect to it. What counts as "sufficiently near" depends on the predicate. Let us associate with the tolerant predicate P a nearness relation, $N_P(\, , \,)$, and express P's tolerance by

(TC) $N_P(x,y) \rightarrow (P(x) \rightarrow P(y))$

By a *Sorites conditional* I shall mean a conditional of the form $P(a) \rightarrow P(a')$, where $N_P(a, a')$ holds; the conditional can be derived from an instance of (TC).

Since the nearness-relation is symmetrical, we can replace, in (TC), $P(x) \rightarrow P(y)$ by $P(x) \leftrightarrow P(y)$. Very often $N_P(x,y)$ has the form: $|m(x) - m(y)| < \varepsilon$, where $m(x)$ is some numerical quantity associated with x (e.g., x's height); sometimes the predicate can be construed as taking a numerical argument (e.g., the number of inches); in this case the condition is simply $|x - y| < \varepsilon$. In principle, ε can depend on x, but in all the paradigmatic examples it is some fixed small enough number. We assume that the nearness relation is such that there are sequences, $a_0 , a_1 , \ldots, a_n$, such that $N_P(a_i,a_{i+1})$ for all $i< n$, and such that $P(a_0)$ and $\neg P(a_n)$ hold on any plausible account. The Sorites conditionals, $P(a_i) \rightarrow P(a_{i+1})$, and $P(a_0)$ imply, via n applications of modus ponens, the unacceptable conclusion $P(a_n)$. Instead of using the Sorites conditionals as premises, we can combine them, via conjunction, into one premise $\bigwedge_{i<n} (P(a_i) \rightarrow P(a_{i+1}))$; in many cases we can also use as a premise a universal generalization of the form: $\forall x[P(x) \rightarrow P(x+\delta)]$, such that $a_n = a_0 + n\cdot\delta$. Any acceptable analysis of the Sorites should handle equally well all these variants.

Attempts based on current theories of vagueness, which have by far and large lumped together tolerance and borderlines, have tried to resolve the Sorites by impugning—in different ways—the validity of the Sorites conditionals. On such a theory, our acceptance of the conditionals derives from some sort of illusion. My common argument against these accounts is that they fail do justice to tolerance, as a semantic phenomenon; they fail to account for the fact that conditionals of the following kind are utterly compelling:

(SC) If John was young one second ago, then John is young now.

I will argue that, far from being illusions, these conditionals are normative rules that contribute to the meaning of the predicates in question. What is wrong in the Sorites reasoning are not the conditionals but the way in which they have been put together to produce the contradiction. The unsophisticated intuition that bans the stringing of "too many" Sorites conditionals is essentially correct; the trick is to find a smooth, non ad-hoc way of building such constraints into the logic.

Since this work does not intend, in any way, to survey existing approaches, I will limit myself to short comments. The supervaluation theory construes a vague predicate, P, as one that is partially defined. It bases the predicate's semantics on the class of all admissible completions of it—referred to as 'sharpenings', or 'precisifications'—where a sharpening is admissible if it satisfies certain obvious semantic axioms; e.g., a distance shorter than a walking distance is a walking distance (henceforth, let 'sharpening' mean one that is admissible). Since, for borderline cases, each Sorites conditional is falsified by some sharpening, no conditional that involves borderline objects is true; but neither is it false. The conjunction of these conditionals is false however (since in every sharpening some conjunct is false). Hence, we can hold the conjunction false, without being committed to the falsity of a particular conjunct (which would have amounted to a sharp cutoff). A similar effect is achieved by using intuitionistic logic, as suggested in [Putnam, 1983]; from $P(a_0)$ and $\neg P(a_n)$ we can derive in that logic $\neg \bigwedge_{i<n} (P(a_i) \rightarrow P(a_{i+1}))$, but this does imply any particular $\neg(P(a_i) \rightarrow P(a_{i+1}))$. Neither of these two accounts gives good reason for the compelling intuition that backs (SC) and its ilk. One may invoke the fact that, for each particular time point, (SC) is true in the great majority of sharpenings: all those that do not place the cutoff for ‘young’ within the given one-second interval. This explanation, appealing as it does to the “amount” of sharpenings that satisfy a given sentence, belongs to the degree-theory framework, which I will discuss shortly. Let me first comment on another, different approach, which treats vagueness as a kind of ignorance.

The epistemic thesis posits sharp unknown cutoffs. There is, on this view, a particular unknown walking distance whose one-foot (or even one-inch) increment is not a walking distance; and there was a particular heartbeat at which I ceased to be a child. The main argument for the view is that it accords best with our practice of applying classical logic across the board, irrespective of vagueness. A considerable part of the argument rests on shortcomings of the supervaluation account, in particular, with regard to the truth concept and the identification of truth with supertruth (truth in all sharpenings). The semantic modality theory of section 3 provides an interpretation of ‘definitely’, without problematic side effects, which fully preserves classical logic. To the extent that this approach is successful, much of the motivation for the epistemic account is gone. But even without semantic modality, the price of epistemicism is simply too high. If I say to my friend ‘Let us meet around 3 o’clock’, then, on the epistemic account, my proposal is that we meet between 3–*a* and 3+*b*, where *a* and *b* are sharp times, up to a tenth of a second, which are unknowable to human beings. Most people would find such an account absurd and here I agree with most people. But to say that the account is incredible is not to say that it is refutable. *Au fond* the epistemic account is of the hidden-parameters variety, which can be always maintained by endowing the hidden parameters with the required features, fantastic as they may be. Perhaps epistemicism is best read as a directive to *posit* unknown, sharp boundaries for vague terms, as the best way of systematizing our practices.

Williamson [1992] [1994, pp. 216-234], suggested a weakened version of the Sorites conditionals, acceptable to the epistemicist. In cases of imprecise information, claims of knowledge must concord with margins of error. I cannot claim to know that the number of people in some audience is at least 200, if the true number is 200 and my knowledge is based on an observation that has a margin of error of, say 15%, because the truth of my guess is accidental. I can only know in that case that the number of people is at least 170. He therefore suggested that we replace the original Sorites conditional, $P(a_i) \rightarrow P(a_{i+1})$, with the weaker conditional, acceptable to the epistemicist:

(EC$_i$) $K(P(a_i)) \rightarrow P(a_{i+1})$, where ‘*K*’ is the knowledge operator: ‘X knows that …’, X being some hypothetical, agent. This blocks the modus ponens applications that lead to contradiction. I suspect however that it does no more than to shift the paradox to another place, where it is obscured by complications arising out of the knowledge modality.[19] In any case, it is very doubtful that the weakened conditional, ‘If I know that John was young a second ago, then John is young now’ explains why we find (SC) compelling, or that it throws light on it. A more promising line is this: Since (on the epistemic view) all, but one, of the many Sorites conditionals are true, the chance of hitting on the false conditional is very small and we ignore it. This type of explanation, which appeals to considerations of quantitative reasoning, will be discussed shortly, when we come to degree theories.

The point of using predicates like ‘around 3 o’clock’ is that there *should not be* sharp cutoffs. I could have said, ‘Let us meet between 2:55 and 3:15’, or mentioned some other precise interval, but this would not have come to the same. ‘Nearly’, ‘approximately’, ‘roughly’ and similar modifiers are systematic fuzzifying devices, used for the purpose of avoiding commitments to sharp boundaries (which is not the same as commitment to boundaries that are unknowable). Not only does the fuzzifier eschew sharp boundaries, but it rules out any criteria by which sharp boundaries might be determined. If I view 3:14:45 as falling within the range of ‘around 3 o’clock’, then I can legitimately persist in my judgment after being informed that 55% percent of competent speakers (or any other reference group) disagree. The function of ‘around’ is to make place for such latitude.[20] Imagine a superhuman computer, who, taking as input all the English speech acts ever performed, derives from it a precise interval for ‘around 3 o’clock’; this, again, would make no difference, since the prescription is no more bending than any particular cutoff point. The guaranteed latitude is a built-in feature of ‘around 3 o’clock’.

We may deem certain problems beyond what humans can solve. There is good reason to think that there are mathematical questions that, for reasons of complexity, are of this nature. A physical theory can imply that certain facts are not knowable to us, in principle. But in all such cases, there is a rich system, and a well-developed methodology for establishing truth, that gives these statements meanings that are independent of their being knowable. One cannot require unknowability as part of what is being defined. If a mathematician claims to have found that CH (the continuum hypothesis) is true, then, barring a hoax or a mistake, I will suspect that he has some argument for the truth of CH, something that might motivate the adoption of a system in which CH is a theorem. But if

---

[19] From $K^nP(a_0)$ and $K^{n-1}(\mathrm{EC}i)$, for all $i<n$, we can deduce, via standard modal logic, $P(a_n)$; e.g., $K^{n-1}(KP(a_i)\rightarrow P(a_{i+1}))$ implies $K^{n-i-1}(KP(a_i)\rightarrow P(a_{i+1}))$, which implies $K^{n-i}P(a_i)\rightarrow K^{n-i-1}P(a_{i+1})$; putting all of these together, *n* application of modus ponens yields $P(a_n)$. To avoid the paradox we must assume that either $K^nP(a_0)$, or some $K^{n-1}(\mathrm{EC}_i)$ is false. But we know that $P(a_0)$, e.g., ‘A one-day old human is young’, is true and, having followed Williamson’s argument we know the truth of (EC_i), for all $i<n$. Reflecting on the way in which we have arrived at this knowledge, we know that we know these facts, i.e., we know $KP(a_0)$, and $K(\mathrm{EC}_i)$, for all $i<n$. This move can be repeated as needed, in order to get the premises that imply $P(a_n)$. My formal argument appeals to the so-called KK axiom: $K\Phi\rightarrow KK\Phi$, which Williamson may want to reject. But although we do not have to subscribe in general to KK, failures of KK should be accounted for; e.g., we might know something without being aware of our knowledge. In the present case there is nothing that stands in the way of KK (besides the goal of avoiding a version of the Sorites paradox). (EC_i) rests on an empirical known fact: a certain limit of our discerning ability. If repetitions of KK are to be rejected here, they should be rejected in any other case of empirical knowledge.

[20] If I find that 95% disagree, I will probably change my usage in order not to be radically out of tune with other speakers. The point however remains that I need not change my usage as long as it does not conflict too much with that of other speakers, and this guarantees sufficient latitude.

someone claims to have found that 5500' is walking distance and 5501' is not, then, barring a hoax, I do not know what to make of the claim; I will suspect that this person does not understand what ‘walking distance’ means.

Being a walking distance is not something that can terminate under a one-foot increment, just as childhood cannot cease at a particular heartbeat, and being rich cannot be lost by losing a penny, and so on and so forth. These are rules of usage, constitutive of the meanings of the predicates in question.

The last point counts also against the degree theorist’s analysis of the Sorites. Degree theory prescribes the assignment of more than two truth-degrees. In the case of the Sorites, $P(a_0)$ has a maximal degree (that corresponds to the classical “true”); the degree of $P(a_i)$ gradually decreases as i increases, hitting the lowest (which corresponds to “false”) at $P(a_n)$. There is also a logic, according to which the degrees of each conditional $P(a_i)\rightarrow P(a_{i+1})$ is high—reflecting the fact that $P(a_i)$ has a slightly higher degree than $P(a_{i+1})$—but not maximal. On degree theory, the paradox arises when we fail to take into account the accumulated effect of many small decreases in truth-degree. It is therefore a species of what might be called the small-effect fallacy, where the smallness of each item blinds us to the magnitude of the cumulative sum.

Of the solutions to the paradox considered so far, the degree theorist’s is the most satisfying. There is no denying the graded nature of vague predicates—i.e., that the aptness of applying them can be a matter of degree—and there is no denying the gradual decrease in degree, as we move in the Sorites chain. More than other approaches, degree theory does justice to these facts. But from this to the institution of a many-valued logic, where connectives are interpreted as functions over truth-degree (or in terms of some measure on sets of sharpenings), there is a big jump.[21] The fact that assignments of numeric values to sentences of the form $P(a)$ may have plausible meaning, and may even serve in certain efficient algorithms, does not make such assignments a basis for a *logic*. This is a wide topic into which I cannot venture here. My present argument, which does not depend on a general evaluation of degree theory, is that by construing the Sorites as a small-effect fallacy, we miss the crucial element of tolerance: its being part of the semantic norms governing of tolerant predicates.

Small-effect fallacies are failures of quantitative thinking, which do not involve a semantic element. Some slippery slopes derive solely from such fallacies and should not be counted as instances of the Sorites. Here is one; call it “the occasional cab paradox”. John’s finances allow him to take a cab from time to time (rather than use the cheaper, less convenient public transportation). “One more cab will not make a difference,” John keeps telling himself, and then at the end of the month... John simply failed to appreciate the effect of many small expenses. This is not an instance of the Sorites, for there are no compelling conditionals of the form: ‘If n cabs are acceptable, then n+1 cabs are acceptable’. In fact, if we model the situation in terms of preferences, or utilities, then there will be sharp cutoffs: the first n, such that taking n cabs per month is strictly more

[21] One kind of degree theory interprets the sentential connectives as functions from degrees to degrees. Another kind presupposes some measure over sets of sharpenings and defines the degree of a sentence as the measure of the set of sharpenings that satisfy it. The first kind has served as a basis of efficient algorithms for certain specialized problems. The second kind has the advantage of preserving, on the syntactic level, classical logic.

preferable than taking any larger number. Some degree theorists compare the Sorites to the lottery paradox.[22] Again, a closer look will show the essential difference. The lottery paradox, recall, consists in the inference:

> For any particular person, it is unlikely that the person will win the lottery; therefore, it is unlikely that someone will win the lottery.

Say we consider conditionals of the form: $Unlikely(x) \rightarrow Unlikely(x+\varepsilon)$, where $Unlikely(x)$ stands for: 'An event of probability x is unlikely', and where ε is some fixed small number. Even if we agree to treat 'Unlikely' as a tolerant predicate over numeric arguments, it is clear that this is not the source of the lottery paradox; else, we could have constructed a Sorites chain without any mention of a lottery. The lottery paradox exemplifies confusion in probabilistic thinking, which is also expressed in terms of an illicit quantifier switch (from 'it is likely that for some *x* …' to 'for some *x*, it is likely that …'; or rather from the negation of the second to the negation of the first). There is nothing semantic about it. The lack of semantic element counts also against the "small chance" explanation offered above on behalf of the epistemic account.

To recap, the main approaches to vagueness explain one's subscribing to Sorites conditionals as a kind of illusion, understandable at best. Such approaches fail to come to grips with the fact that the conditionals express norms of linguistic usage. This becomes clear if we consider how bizarre are scenarios in which the norm is contravened. Say, John wants to know if Meg is rich. Being informed, by Meg's accountant, of all her assets, to the last penny, John concludes that she is not. But when the accountant, who happens to glimpse Meg passing by, adds "and, by the way, she has just picked up a penny from the pavement," John changes his verdict: "Now, she is rich."

People are quite aware of the possibility of sharp cutoffs in gradually changing situations. Threshold phenomena are known and respected. We have the story of the straw that broke the camel's back, but we do not have the story of the penny that made someone rich.

If tolerance means that the Sorites conditionals are valid, then how are we to avoid the resulting contradiction? In practice, we do this by refraining from carrying out all the modus ponens steps. The conditionals are thus subject to a restriction: one should not string too many of them in one chain of reasoning. Non-philosophers, who are amused but undismayed by the paradox, recognize implicitly that, although each conditional is valid, stringing too many of them is illicit. What I shall propose differs from this brute force solution. But let me pause to observe that, unsatisfying as it may appear to philosophers, there is nothing incoherent in imposing such a restriction, as part of the rules that determine the meaning of a tolerant predicate. A rule that limits the use of another rule can itself be a meaning-determining rule.

Non-linguistic analogies are easy to come by. Some instruments are intended for limited utilization. We have one-shot (disposable) gadgets, and some that must be changed after a certain amount of use. In principle, a device can come with a manual that contains an instruction of the form: "Do not use more than …times", or "do not use more than … times in succession". To transgress the restriction is to use the device contrary to its

[22] McGee and McLaughlin [1994, pp. 221,222].

intended "meaning". Such restrictions can arise also with respect to information gathering. A measuring procedure may involve a negligible error, which is not negligible when the procedure is repeated many times and the results are added. There is therefore a restriction against piecing together too many measurements. How much is "too many" depends on the margin of error of each measurement and the desired accuracy of the total. To the extent that the latter is vague, so is "too many".

Exploring this avenue, we may consider the following kind of restriction: *A derivation should not use too many Sorites conditionals*, where "too many" means that the set of conditionals make it possible to traverse a Sorites chain. Here, 'derivation' is defined in a standard way, using a standard system based on modus ponens and universal generalization; a wff is *used* in a derivation if it is either equal to, or a component of a wff occurring in the derivation. This restriction blocks all the versions of the Sorites paradox, including the version based on a universally quantified premise; that derivation is blocked, since it involves instantiations to all the conditionals of the chain. Universal sentences can be used and they can be instantiated, as long as the instantiations do not yield all the conditionals of a Sorites chain.

The restriction would certainly suffice for the purpose of everyday reasoning with vague predicates. We can tighten it further: *A derivation should not contain the names of all objects that make a Sorites chain.* Further technical details can be worked out. The set of theorems of the resulting system is not closed under modus ponens; e.g., we can prove $P(a_0)\rightarrow P(a_i)$, using the first half of the Sorites chain, and $P(a_i)\rightarrow P(a_n)$, using the second half, but we cannot prove $P(a_0)\rightarrow P(a_n)$. The system thus allows for some forgetfulness; we can employ, on separate occasions, different collections of Sorites conditionals, without being able to use jointly all the conclusions. Arguably, predicates that are vague and tolerant are intended for local use. I may judge on one occasion, 5280' (1 mile) to be walking distance, and on another—to be non-walking distance. But on every occasion, if *on the same occasion*, there are also questions about 5279' and about 5281', the answers had better be the same. We are safe from wrong answers (such as, "10 miles is walking distance") since no single occasion contains questions about each and every item in a Sorites chain.

The brute force approach is unsatisfactory as it imposes a proof theoretic restriction designed for the purpose of avoiding a contradiction. Yet the needed positive account is already indicated by the above considerations. The system should express the "local" aspect of the predicate's use: the fact that, on each particular occasion, the predicate has to be applied only to a restricted collection of objects, one that is not expected to contain Sorites chains. A context, in the proposed logical system, is a set of objects, consisting of objects that, on a single occasion, have to be classified with respect to the predicate. Tolerance means that near enough objects should be classified alike. A context for which there exists such a classification, which is moreover compatible with the semantic axioms governing the predicate, is called *feasible*. In all the standard examples of tolerant predicates, this is equivalent to the requirement that the context does not contain a Sorites chain. Now, ruling out non-feasible contexts leads to a cumbersome system. The better way is to allow all contexts but to restrict the tolerance requirement, so that tolerance is, by definition, tolerance in all feasible contexts. We shall furthermore see that sentences and proofs have associated contexts. Those whose contexts are feasible form the feasible portion of the language; and it is within this portion that a tolerant predicate is meant to be used. The proof of the Sorites contradiction fails, because it requires an unfeasible context and in unfeasible contexts a tolerant predicate looses its tolerance: it has some sharp cutoff. Unfeasible contexts do not arise in practice.

5. Tolerance, Context, and Feasible Contexts.

That perceptual judgments can depend on context is an old idea that goes back to psychological researches at the end of the 19th century. Stumpf [1883] notes cases where a subject who experiences two stimuli, *a* and *b*, judges them to be equal, and also judges equal *b* and *c*; but when experiencing *a* and *c*, reports that $a< c$. This contradicted accepted theories of the time, on which stimuli are classified by absolute intensity. Kurt Kofka [1922] who reports this, and on whom I draw for this history, notes attempts (by Ebbinghaus and Titchener) to explain the phenomenon by assuming a certain "friction": a lingering influence of a previous sensation on a later one. The true solution to Stump's puzzle has been proposed by Cornelius [1897]: The subject's judgments should be relativized to the experimental settings: (i) comparison of *a* and *b*, (ii) comparison of *b* and *c*, (iii) comparison of *a* and *c*. The subject's report should not be read as: "*a* and *b* have equal intensity," but as: "in the setting of comparing *a* and *b*, they have equal intensity". The paradox disappears, when the contextual parameter is introduced. (As Kofka tells the story, Stumpf admitted this possibility but considered it too complicated.) Let the context consist of the items under comparison, and let the relativization to context be effected by prefixing an operator: […], where '…' is a list of the compared items. Then, the subject's reports are written as follows, where '$i(x)$' denotes *x*'s intensity:

$$[a, b]\,(i(a) = i(b)) \qquad [b, c]\,(i(b) = i(c)) \qquad [a, c]\,(i(a) < i(c))$$

These do not imply a contradiction, but if we remove the contextual operator, they do. The context in which *a*, *b*, and *c* are simultaneously compared may not be psychologically feasible, but it can be part of the formal language and we can consider sentences of the form $[a,b,c](\ldots)$. Connections between sentences in different contexts can be introduced through axioms, e.g., $[x,y]\,(i(x) \neq i(y)) \rightarrow [x,y,z]((i(x) \neq i(y))$.

In Sorites scenarios, it appears that human's classification of single items depends on the end of the chain from which the item is approached.[23] D. Raffman [1994], noting that the Sorites contradiction disappears upon relativization to the history of previous perceptions, proposed this as a way of avoiding the contradiction.[24]

Much emphasis has been laid in the philosophical literature on psychological factors as the source of tolerance. The most standard example of a Sorites chain is a sequence of colors. In this case—as well as in cases of other predicates of perceptual nature—tolerance might appear an inevitable outcome of perceptual limitations: since we do not

---

[23] When drawings in a "cat-dog" sequence, in which a cat gradually becomes a dog, are successively presented to the subject, the verdict ("cat" or "dog") for borderline cases will depend, as a rule, on the end from which this sequence is run. People tend to stick with their initial classification. I have it on the authority of reliable people who have seen such a report, but we could not locate a reference. I will be thankful to any reader who can direct me to it.

[24] Raffman's approach is psychologistic and informal: the context is constituted by the short history of the agent's impressions; there is no logic and no concept of "feasible context" which plays center role in the present setup.

perceive the difference between near enough items we judge them as equal. We deduce the difference between indistinguishable patches, *a* and *b*, only by bringing into play additional items, e.g., *a* is distinguishable from *c*, but *b* is not. The so-called non-transitivity of perceptual sameness has been invoked to explain what goes on in the Sorites paradox.

The emphasis on perceptual examples is misplaced. Vagueness, tolerance and the Sorites do not depend on perceptual factors. In 'walking distance', 'rich', 'old', 'noonish', 'a large community of chimpanzees', and so on, tolerance is a conceptual phenomenon that does not hinge on indistinguishability (4017' and 4018', or 50 and 51 chimpanzees, are as distinguishable as 1 and 2). Yet, it is part of the predicates' meanings that if 4017' is walking distance, so is 4018'; and if 51 chimpanzees constitute a large community, so do 50. The crucial factor is that the distinction can be ignored for the purposes for which these predicates are employed, and that by ignoring it we achieve an enormous gain in efficiency. The emphasis on perceptual aspects is unfortunate, because it is impossible to argue with perceptual facts, but it is possible to "argue" with patterns of reasoning. Even if—to put crudely a subtle point—our usage is not subject to an external evaluation, we can analyze, criticize, and apply rationality standards to modes of reasoning, in ways that would be impossible, had these been brute psychological facts, like the indistinguishability of two sensations.

The logic that can handle contextual effects is not, or need not be, a revision of classical logic (like intuitionism or quantum logic), but an extension of it (like modal logic, or dynamic logic). Contextual logic itself can be seen as part of the broader domain of *resource-bounded* reasoning. These are linguistic and logical organizations that are suited to work efficiently under the severe limitations of our deductive, and data processing, capacities. Take, for example, personal proper names. With thousands of people sharing the same name, disambiguation is achieved by context. It is, in principle, possible to do away with context, by instituting a universal system for tagging every past and future human with a unique number. On a smaller scale, consider using social security numbers as names of people. This may be practicable for computers but not for humans, who do incomparably better by letting the context determine the reference. For proper names, this is the end of the story; there is no further logical development, since the determination of the name's reference by context is not amenable to precise treatment. In the case of tolerance, there is an additional logical story, since "tolerance reasoning" can be formally systematized. Before going on, it might be worthwhile to put the proposed contextual notation in wider perspective.

Representing Contexts

The proposed system employs instances of the notational scheme:

(C) $[C]\phi$,

where ϕ is a wff, 'C' denotes some context and '$[C]\phi$' is to be read: 'ϕ in the context C.' The context can be specified unsystematically, by letting 'C' stand for some description of the occasion of use:

[Uttered in circumstances _ _ _] Edith Cohen is well to do.

Here there is no precise description of the way the context fixes the references; we cannot, in general, do better than: “the most plausible woman referred to by ‘Edith Cohen’ in circumstances _ _ _”. But there are contexts, and context dependencies, that are amenable to precise, systematic treatment. In cases of time and place indexicals ‘C’ stands for a description of a spatio-temporal region, which determines the reference of the indexicals within its scope, according to the indexicals’ types:

[Uttered on June 2, 2000 in Manhattan] Today is a beautiful day.

As far as I know, scheme (C) has not been used before. But quite a few existing systems incorporate context-dependency in other ways, usually through the semantics. In temporal logic, sentences are evaluated at time points; in dynamic logic, they are evaluated at “process points”, i.e., states in an execution of a given program. Kaplan’s [1989] system LD is of this kind, except that a sentence is evaluated at (i) a context consisting of a person, a time point and a place (the references of ‘I’, ‘now’, and ‘here’), and (ii) a non-indexical time point and a possible world. The context operator ‘[C]’ makes for richer expressive possibilities. We can have, for example, sentences in which ‘now’ is employed at different times, and we can combine them:

[Uttered at time t_1] Now it is day $\wedge$ [Uttered at time t_2] Now it is night.

This is impossible in the systems just mentioned. Expressive power however is not always a virtue. Since temporal and dynamic logic are used for verification and analysis of computer programs, where simplicity and easy application are crucial, the restriction of expressive power can be very useful. But in LD, whose aim is purely philosophical, the impossibility of treating statements by different speakers who employ the word ‘I’ is a serious limitation.

We can express the contextual effects of indexicals without appealing to (C), by basing the system on tokens, rather than types. We can recast ‘Today is a beautiful day’ as:

< t > is a beautiful day,

where ‘t’ denotes a token of ‘today’ and where ‘< t >’ denotes, by definition, the day in which t occurs. In general, if ‘i’ denotes a token of an indexical, then ‘<i>’ denotes the object referred to by ‘i’, according to the rules governing that type of indexical; <…> incorporates the rules of what Kaplan calls *character*. (To make this more explicit, we can display the type of the indexical within the notation: ‘<t, today>’.)

Finally, the Liar paradox provides an example of a different kind of context-dependency. ‘The sentence on line 1 is not true’ fails to express a proposition (or a true-or-false proposition) when written on line 1; but the proposition expressed by an additional token of that sentence, on another line, is true. The paradox, and the semantics of self-reference (direct and indirect) are treated in [Gaifman 1992, 2000], by using pointer systems, where pointers are generalizations of tokens. (As shown in those works, this kind of context dependency is not of the indexical type.) In contrast to the systems just mentioned, tolerance requires a logic that makes use of (C).

Feasible Contexts

The relativization of the tolerance conditional, (TC), of the previous section to context is:

(TC*) $N_P(x,y) \rightarrow [C](P(x) \rightarrow P(y))$.

where 'C' stands for a finite list of terms (including, possibly, variables) referring to the objects that have to be classified into P's and non-P's. For convenience, we speak occasionally of the context C, meaning by it the set of these objects. From:

$N_P(x,y) \rightarrow [C](P(x) \rightarrow P(y))$ and $N_P(y,z) \rightarrow [C](P(y) \rightarrow P(z))$,

we can derive the conclusion: $N_P(x,y) \wedge N_P(y,z) \rightarrow [C](P(x) \rightarrow P(z))$. But, generally, we cannot derive something of this kind, from:

$N_P(x,y) \rightarrow [C](P(x) \rightarrow P(y))$ and $N_P(y,z) \rightarrow [C'](P(y) \rightarrow P(z))$,

where C' is different from C. Tolerant predicates are meant for use in contexts that can be partitioned into P's and non-P's, in a way compatible with the semantic axioms governing the predicate (e.g., a distance smaller than a walking distance is walking distance) without violating tolerance—i.e., without introducing sharp cutoffs. The context can have many members, as long as there is at least one sufficiently large gap that can serve as a divisor between P's and non-P's. We shall call such contexts *feasible*. It will turn out that, for any given C, we can express in the formal language the condition that C is feasible. Adding the formula that expresses this, as a conjunct to the antecedent in (TC*), we get a scheme whose satisfaction guarantees the tolerance of P in all feasible contexts.

As stated in the previous section tolerance is, by definition, *tolerance in feasible contexts*. In an unfeasible context, a predicate looses its tolerance, since there must be a sharp cutoff. Unfeasible contexts are contexts we expect not to encounter; such a context, for example, will force us to define with 10' precision the maximal walking distance. There is however no harm in a system that has the operator [C] for all C's, which prescribes suitable sharp cutoffs for unfeasible contexts. To do otherwise is to impose an awkward syntactic restriction. Yet only the feasible part of the formal language corresponds to actual usage, and only derivations carried within this part represent actual reasoning.

Let us now consider an example showing that tolerance need not imply vagueness (this discharges a promise made in the introduction). A certain educational institution awards, every four years, a highly prestigious prize, according to the following regulations. The finalists—those who pass very demanding preliminary tests—accumulate scores by passing final examinations. The maximal total score is 20. A finalist who scores at least 19 wins the prize unconditionally (more than one prize can be awarded). There is, in addition, a tolerance condition: a finalist, whose score differs by one point from that of a winning finalist, wins the prize as well. (Such a policy can be plausibly motivated, e.g., by the wish to minimize the possibility that an accidental small difference be construed as discrimination.) There is a Sorites-type worry that someone with a very low score may get the prize. This worry is laid to rest by the fact that, due to the screening process, the possibility of there being four finalists, with scores 19,18, 17, 16, is extremely unlikely and has never happened. It is therefore ignored. For the sake of a formal system, we can add some a cutoff between, say 16 and 17, for unfeasible contexts: those containing a chain from 19 to scores smaller than 17.

The predicate ‘finalist who qualifies for the price’ is tolerant, where the nearness relation is such that finalists are near just when they are in the same year and their scores differ by one point at most. A context consists of all finalists in a given year. Although the predicate is tolerant, it is as sharp as a predicate can be. Admittedly, this is an artificial example of a legalistic nature. Nonetheless is suffices to establish that, in principle, tolerance does not imply vagueness.

While the non-implication from vagueness to tolerance is easily established by examples from natural language, we need an artifice for the non-implication in the other direction. Since tolerance is a contextual phenomenon, the role of context should be precisely specified in order for the example to work. In natural language, however, there is an overall tendency to leave the role of context unspecified and vague (indexicals are rather an exception). And if we have tolerance, but ignore the role of context, we are bound to get vagueness. If $P(c)$ is true in some contexts, false in others, where the contexts are left unspecified, then the inevitable implication is that there can be divergent opinions about c, i.e., that it is a borderline case. This is not to suggest that paradigmatic examples, such as ‘old’, ‘rich’ etc., became vague through our ignoring some precise contextual conditions. These predicates are both tolerant and vague, to start with.

Note that, although the predicate of the example is not vague, we have degrees, which, moreover, can be defined solely in terms of the predicate (without appeal to scores): Finalist a ranks higher than finalist b just when the set of contexts in which a is a winner properly includes the set of contexts in which b is a winner. Vagueness is absent, due to the sharp specification of the contextual winning conditions.

Contextual logic handles tolerance *per se*, i.e., without vagueness. The required dimension of vagueness is obtained by constructing on top of it a modal system with a possible-world semantics, along the lines of section 3, where each possible world is a model for contextual logic.

6. CL, Contextual First Order Logic: the Formal System.

CL is first order logic, augmented by context operators, which are of the form $[t_1,\dots, t_n]$, where $t_1,\dots, t_n$ is any list (finite sequence) of terms; terms are individual variables, individual constants, and—if the language has function symbols—expressions built from them in the usual way. Wffs are defined recursively, by the usual clauses and the following additional clause:

> If $t_1,\dots, t_n$ is a finite sequence of terms, and α is a wff, then $[t_1,\dots, t_n]\ \alpha$ is a wff. Occurrences of variables in the t_i's are free occurrences in $[t_1,\dots, t_n]\ \alpha$.

We refer to $t_1,\dots, t_n$ as a *context list*. The context corresponding to a list is the finite set consisting of the values of the terms. (For convenience, we sometimes use ‘context’ ambiguously: for the set of objects and for the context list.)

'x', 'y', 'z', 'x_1', 'x_2',…, 'y_1', 'y_2',…etc., stand for individual variables, 's', 't', 's_1', 's_2',…, 't_1', 't_2',…—for terms, and 'C_1', 'C_2', …, 'C'', 'C''', ….— for context lists. We use self-explanatory customary notations: if $C = t_1, \ldots, t_n$ and $C' = s_1, \ldots, s_m$, then, $C, C' = t_1, \ldots, t_n, s_1, \ldots, s_m$; a single term is also regarded as a list of length 1, hence, $t, C = t, t_1, \ldots, t_n$.

It is convenient to include the empty list, where $[]\alpha$ is, by definition, α.

The Semantics: Context Dependency Functions.

The non-logical vocabulary includes context-independent predicates and, possibly, function symbols. Their interpretation is given, in the usual way, as a model for a first-order language. Predicates that are tolerant are context-dependent; and every tolerant P has an associated nearness relation N_P, which is context-independent. The setup however is designed to treat context dependency *per se*, without reference to tolerance.

Whether a context-dependent predicate is true of a given object (or n-tuple) depends on a context: a finite set of objects. For the sake of simplicity, assume that we have one monadic context-dependent predicate, P. The generalization to the case of several context-dependent predicates, including predicates of higher arities, will be obvious; occasionally, we shall indicate it as we go along. A model for the language is a structure of the form:

$$(\mathcal{M}, f),$$

such that: (i) $\mathcal{M}$ is a model for the context-independent vocabulary; the universe of $\mathcal{M}$ is denoted as: $|\mathcal{M}|$, and (ii) f is a function that associates with every finite subset, $X \subseteq |\mathcal{M}|$, a subset of X. This subset is supposed to consist of all members of X that, *in the context* X, fall under P. Otherwise stated, f associates with each X an interpretation of P over X. We refer to f as a *context dependency function*, or for short: *cdf*.

(If P is an n-ary predicate, where n>1, then $f(X) \subseteq X^n$; it consists of all n-tuples in X^n that, in the context X, fall under P. If there are several context-dependent predicates, then f associates with each, an interpretation, $f(P, X)$, over X.)

For any cdf, f, and any context Y, let f_Y be the cdf that associates with every X the subset consisting of the members of X that fall under P, in the context $X \cup Y$:

$$f_Y(X) =_{\mathrm{Df}} f(X \cup Y) \cap X.$$

The idea is that Y serves here as a fixed contextual background; we add its members to any context on which the function operates. We refer to f_Y as the *conditionalization* of f on Y.

The truth-values of wffs are now defined recursively. As usual, the truth-value depends on assignments of members of $|\mathcal{M}|$ to the free variables of the wff. To avoid elaborate notations, we will assume an implicit assignment to the free variables. The value of a term under the assignment is determined in the usual way, via the interpretation of individual constants and existing function symbols.

Atomic Wffs: Satisfaction of atomic wffs with context-independent predicates is their usual satisfaction in the model $\mathcal{M}$. For other atomic wffs:

$(\mathcal{M}, f) \models P(t)$ iff $a \in f(\{a\})$, where a is the value of t.

This means that, by definition, $P(t)$ is evaluated in the context consisting of t, or more precisely, of t's value. (For an n-ary P, with $n>1$, the context list consists of all the terms that appear under the predicate.)

Context Operator: $(\mathcal{M}, f) \models [t_1,\ldots, t_n]\, \alpha$ iff $(\mathcal{M}, f_Y) \models \alpha$, where Y consists of the values of the t_i's.

Roughly, this means that the prefixing of $[C]$ amounts to augmenting any context by the values of the terms of C.

Sentential Connectives: The usual clauses of classical logic.

Quantifiers: The usual clauses of classical logic.

This system is quite general. We have not imposed any restriction, in particular, nothing relating to tolerance. Specific features are imposed through restrictions on the set of models. Various important restrictions are expressible by sentences in the language of **CL**. Since the forthcoming deductive system is sound and complete, we get also sound and complete systems for the specific cases, by adding these sentences, as axioms.

The Deductive System:

Any of the standard deductive systems for first-order logic can serve as a basis, a suitable enhancement of which constitutes a deductive system for contextual logic. All first-order axiom schemes and inference rules are retained (where the wffs range over the language of **CL,** and where the free variables of a wff are defined as above). Details of the enhancement depend on the first-order system we start with. Assuming that the first-order inference rules are modus ponens and universal generalization, add the following axiom schemes:

(I) $[x_1,\ldots x_n]\alpha \rightarrow [x'_1,\ldots x'_m]\alpha$, where $\{x_1,\ldots x_n\} = \{x'_1,\ldots x'_m\}$ (i.e., the same set of variables). The effect of the axiom is to make $[C]$ depend only on the set of terms occurring in C.

(II) (i) $P(x) \leftrightarrow [x]\, P(x)$
(ii) $R(x_1,\ldots, x_n) \leftrightarrow [y_1,\ldots,y_m]\, R(x_1,\ldots, x_n)$, for every context-independent R.

(III) $[C]\, [C']\, \alpha \leftrightarrow [C, C']\, \alpha$.

(IV) (i) $[C]\neg\, \alpha \leftrightarrow \neg[C]\, \alpha$,
(ii) $[C]\, (\alpha \rightarrow \beta) \leftrightarrow ([C]\, \alpha \rightarrow [C]\, \beta)$.

(V) $[C]\, \forall y\alpha \leftrightarrow \forall y[C]\, \alpha$, where y does not occur in C.

It is easily seen that $[C]\,(\alpha * \beta) \leftrightarrow ([C]\ \alpha\ *\ [C]\ \beta)$ is a theorem, for every binary sentential connective, $*$, and that we can similarly replace '$\forall$' by '$\exists$' in (V)

I have chosen (III), (IV) and (V) for clarity, rather than economy. Since we have universal generalization, it is sufficient to let *C* and *C′* be lists of variables. We can furthermore restrict them to the case where *C* consists of a single variable. We can also base the system on an operator of the form $[t]$, and define $[t_1,\ldots, t_n]\ \alpha$ as:

$$[t_1]\ [t_2]\ \ldots\ [t_n]\ \alpha.$$

This would make (III) redundant. Yet, (III) expresses the non-trivial semantics of conditionalizing on context; we do better not to disguise it as a syntactic convention. Alternatively, instead of using lists of terms, we can use sets of terms; this identifies *C* with any list containing the same terms and makes (I) redundant.

Note that, if α is any axiom of first-order logic, then $[C]\alpha$ is a theorem, for every context *C*: Take first the case where *C* consists of variables not occurring in α; using (IV) and (V), "push" *C* inside, all the way, getting a provably equivalent wff; this wff is a first-order axiom. From this the general case is obtained by universally generalizing over the variable of *C* and instantiating them to the desired terms.

Soundness and Completeness Theorem: If Γ is a set of sentences, then

$$\Gamma \vdash \alpha \quad \Leftrightarrow \quad \Gamma \models \alpha\ .$$

Here '$\vdash$' and '$\models$' denote, respectively, provability and logical implication in **CL** (a proof from a given set of wffs is defined in the usual way, and logical implication means that α is satisfied in every model $(\mathcal{M}, f)$ that satisfies all members of Γ).

Restricting the Effects of Contextual Change

Since, in general, the cdf can be arbitrary, the system does not impose connections between $[C]P(t)$ and $[C']P(t)$, where *C* and *C′* denote different contexts. **CL** being a general framework, it leaves the contextual effect open. We can impose the desired connections between different contextualizations through additional axioms. The enhanced system will be complete and sound, since the deductive system of **CL** is. In section 7, subsection `Conservativeness', we discuss a continuity principle that restricts the effect of context changes in the case of tolerant predicates, and we show how to implement it via suitable axioms. In a more general philosophical vein, there is a worry that contextualization might serve as an evasive technique, namely, one may excuse contradictory beliefs through relativization to different contexts. Such a worry should be addressed not by ignoring the possible effects of context, but by imposing explicitly the required continuities.

Contextual Scope

Formalization in **CL** may require decisions about contextual scope. The default reading of the wffs gives the context operators minimal scope, e.g., $P(t_1) \wedge P(t_1)$ is equivalent to $[t_1]P(t_1) \wedge [t_2]P(t_2)$, not to: $[t_1, t_2]\, (P(t_1) \wedge P(t_2))$. The latter is equivalent to: $[t_1, t_2]\, P(t_1) \wedge [t_1, t_2] P(t_2)$. Hence, giving the context operators maximal scope amounts to combining them and applying the combined context to each component. Depending on the case, this may seem the natural interpretation; e.g., '11:30 AM is noonish and 1:00 PM is noonish' is naturally evaluated in the context containing both time points. This problem does not arise with respect to negation, since $\neg$ commutes with $[C]$. We can consider a conjunction connective, $\wedge^{\#}$, which has an additional context-combining effect. It satisfies, among other axioms, the scheme:

$$[C_1]\alpha \wedge^{\#} [C_2]\, \beta \quad \leftrightarrow \quad [C_1, C_2]\, (\alpha \wedge^{\#} \beta).$$

Similar versions exist for all binary connectives. Since we can achieve the effect of the #-connectives by using the classical connectives and the context operator, we do not need them as primitives.[25] Note that when we interpret universal quantification as a (possibly infinite) conjunction, we should use conjunction in the sense of $\wedge$, not in the sense of $\wedge^{\#}$; because the $\wedge^{\#}$-based interpretation combines all the contexts of the instances, resulting, as a rule, in an unfeasible context. Such quantifiers will be useless for the purpose of handling tolerance.

Quantification, as well, may require contextual scope adjustment. Sometimes it is natural to give the quantified variable maximal contextual scope over the quantified wff; that is, we read $\forall x \alpha(x)$ as equivalent to $\forall x[\mathrm{x}]\alpha$. A more refined reading takes into account the free occurrences of the quantified x in terms that occur under the context-dependent predicate:

$$\forall x[t_1, \ldots t_n]\, \alpha,$$

where $t_1, \ldots t_n$ are all the terms which occur under P and contain x that is free in α. There is a variant, call it **CL***, obtained by reading the quantifiers in this way. It has a sound and complete deductive system, and it retains classical sentential logic. **CL*** differs from **CL** in the following items of quantification logic: The instantiation axiom (for legitimate substitutions) is changed to: $\forall x\, \alpha(x) \rightarrow [t_1, \ldots t_n\,]\alpha(t)$, where $t_1, \ldots t_n$ is the above list; the axiom $\forall x\, (\alpha \rightarrow \beta) \rightarrow (\forall x \alpha \rightarrow \forall x \beta)$ is deleted; universal generalization is changed to the following inference rule: from $\beta \rightarrow [t_1, \ldots t_n\,]\alpha(x)$, where x is not free in β and $t_1, \ldots t_n$ is as above, infer: $\beta \rightarrow \forall x \alpha$.[26]

---

[25] There is a system, call it **CSL#**, of contextual *sentential* logic, which is built only on the basis of the #-connectives. In sentences of **CSL#,** every sentence component contributes a context, and the sentence is evaluated in the context that is the union of all contributions; every sentence, α, of **CSL#** can be rewritten in equivalent form $[C]\alpha'$, where α' does not contain context operators. There is also a corresponding sound and complete deductive system. **CSL**, the sentential fragment of **CL,** is much more expressive than **CSL#** and its deductive system is an extension of classical logic, which **CSL#** is not.

[26] In **CL***, provably equivalent wffs, need not have provably equivalent generalizations; e.g., α and $\alpha \wedge \beta$, where β is a tautology containing , under P, constants not occurring in α. By applying a quantifier to $\alpha \wedge \beta$ we "pull out" these constants into the global context. Therefore, mere mention of an item, even in a "vacuous" way, may have a substantial effect, by changing the context. Indeed, in ordinary speech we would not add, for no specific reason, a vacuous conjunct, such as $c=c$. Regarded within this perspective, this feature of **CL*** may reflect actual intuitions.

There are other variants that I shall not discuss here. The big advantage of **CL** is that it retains full classical logic and has the greatest expressive power. The default scopes that underlie **CL** can be overridden by inserting context operators in the appropriate places; hence, whatever is expressed in some variant can be expressed in **CL**; but not vice versa. The price for this is a greater need for contextual scope adjustment. After Russell, we should not be deterred by scope adjustments.

7. Specific Systems, Tolerance, and Conservativeness.

When it comes to specific systems, the choice of $\mathcal{M}$ is obviously restricted by the semantics of the context-independent vocabulary. Regarding this, we can leave various details open. E.g., the first-order language can be either one-sorted, or many-sorted—with different sorts of variables ranging over different sub-domains of $|\mathcal{M}|$. The theory of $\mathcal{M}$ may or may not be characterized by a recursive set of first-order axioms; we can, for example, assume that $\mathcal{M}$ includes, among its parts, the standard model of natural numbers. The forthcoming analysis does not depend on these details.

The choice of f is constrained by requirements concerning the context-dependent vocabulary. In the usual examples of tolerant predicates, such requirements involve some ordering—or, more generally, pre-ordering—of the relevant domain (cf. section 3, footnote 16). They impose either monotonicity (e.g.: if a person is tall, any person of equal or greater height is tall), or anti-monotonicity (e.g.: if d is walking distance, every smaller or equal distance is walking distance), or the interval property (e.g.: any time point between two noonish time points is noonish). Sometimes more complex conditions are required, which involve multi-dimensional comparisons, or more than one tolerant predicate. But we assume here that they amount to a finite number of universal generalizations, which can be formalized in usual first-order logic as:

(1) $\forall u_1,u_2,\ldots, u_k\, \phi(u_1,u_2,\ldots, u_k)$,

where ϕ is quantifier-free (if $k = 0$, ϕ has no free variables). Our analysis applies to this general case, but we shall use, as a typical example, the anti-monotonicity of 'walking distance', formally expressed by:
$\forall x, y\, (x< y \rightarrow (P(y) \rightarrow P(x)))$. In order to express it in **CL**, we have to consider possible contexts. Minimally, we want to say that, *in the context of any two distances*, if the bigger distance is a walking distance, so is the smaller:

$$\forall x, y\, [x,y](x< y \rightarrow (P(y) \rightarrow P(x))).$$

Since $<$ is context-independent, it is equivalent to: $\forall x,y(x< y \rightarrow [x,y]\, (P(y) \rightarrow P(x)))$. Obviously, this should be imposed in any context whatsoever. Hence, the following scheme:

$$\forall y_1,\ldots y_n\, [y_1,\ldots y_n]\, \forall x,y\, [x,y](x< y \rightarrow (P(y) \rightarrow P(x))),$$

where $n = 0,1,\ldots$. The satisfaction of the scheme means that $f(X)$ is an initial segment of X under the ordering $<$, for every context, X. In general, the minimal **CL** version of (1) is:

(1*) $\forall u_1,\ldots, u_k\, [t_1,t_2\ldots,t_m]\, \phi(u_1,\ldots, u_k)$, where $t_1,t_2\ldots,t_m$ are all the terms occurring in ϕ under P (or under any context-dependent predicate, if there are several).[27] To ensure that this holds in any context we use the scheme:

$$(1^{**})\quad \forall y_1,\ldots y_n \forall u_1,,\ldots, u_k\, [y_1,\ldots y_n]\ [t_1,\ldots,t_m]\ \phi(u_1,\ldots, u_k),$$

where n=0, 1….. If k = 0, ϕ is quantifier free; e.g., '100' is walking distance', whose corresponding axiom scheme is: $\forall y_1,\ldots y_n\, [y_1,\ldots y_n]P(100')$.

From now on, let us assume that P is tolerant. Unrestricted tolerance can be stated as the universal generalization of (TC^*) (see section 5, "Feasible Contexts"). But this scheme has to be modified, since we want to impose it only in feasible contexts. Tolerance in the context X means that any two members of X related by N_P are classified both as P's, or both as non-P's. (N_P can be vague, but this is an aspect we are not concerned with now; it will be handled by using, instead of a single model ($\mathcal{M}$, f), a family of possible models. For the moment, N_P is a sharp binary context-independent predicate.[28])

Take 'walking distance' as a representative case. Besides the anti-monotonicity axiom, we assume an axiom specifying that certain distance, say c, is a walking distance, and another specifying that a certain distance, c', is not a walking distance. The axioms (formulated in **CL** as schemes of the form (1**)), hold in all contexts. For a context, X, let $b_1 < \ldots < b_j < \ldots < b_{m-1}$ be all the members of X in ascending order, which are strictly between c and c'; let $b_0 = c$, $b_m = c'$. Then X is feasible iff, for some $i < m$, b_i and b_{i+1} do not stand in the N_P relation. For if b_i and b_{i+1} do not stand in the N_P-relation, we can define $f(X)$ as the subset of X consisting of all members that are $\leq b_i$; that is, we use the gap between b_i and b_{i+1} as a separator between P's and non-P's. Tolerance is satisfied because there is no cutoff separating N_P- near members of X. Obviously, if the condition fails, then there is no way of defining $f(X)$ that does not violate tolerance.

The analogous requirement for 'noonish' (a predicate that satisfies the interval condition) should be obvious: two sufficiently large gaps are required, to separate the non-noonish times before 12:00 PM, and the noonish from the non-noonish times after 12:00 PM. It is easily seen that if X is feasible, so is every $X' \subseteq X$.

It is not difficult to see that, in these examples, the condition characterizing the feasibility of a context of q members is expressible by a first-order wff in the context-independent vocabulary of **CL**. Let $Fsble(x_1,x_2\ldots,x_q)$ be this wff (in fuller notation there is subscript 'q', since the wff depends on q). Tolerance is then expressible by the scheme:

$$(TOL)\quad \forall x_1,x_2\ldots,x_q\{Fsble(x_1,\ldots,x_q) \rightarrow [x_1,\ldots,x_q](N_P(x_1,x_2)\rightarrow(P(x_1)\rightarrow P\,(x_2)))\}$$

---

[27] Naturally, one should write ϕ in a way that avoids redundant terms under the predicate P (e.g., omit conjuncts of the form $P(t)\vee\neg P(t)$). The differences that are due to different ways of writing ϕ become, however, unimportant when we pass to the scheme (1^{**}).

[28] It is possible to construe N_P as context-dependent. We can fine tune it so that the nearness relation shrinks as the context approaches a Sorites chain, becoming at the end empty. In this way one can retain a nominal "tolerance" in all contexts. But a nearness relation that is context-independent reflects better our ordinary intuitions. I also prefer not to have too many free parameters to play with.

In general, the feasibility of X means that we can partition X into P's and non-P's in a way that is compatible with the universal generalizations, (1), that determine the semantics of P, so that every two N_P-related objects are in the same part. This requires that we define precisely 'compatibility with ψ', where ψ is a wff of the form (1). The rather simple universal wffs, expressing monotonicity, anti-monotonicity, or the interval property, do not capture all possibilities. For example, not all quantified variables of (1) have to appear under P and, moreover, they can appear under function symbols. The following outline indicates how this is done. The reader who is satisfied on this score can skip it.

The General Case

For simplicity, we shall regard the objects of the relevant domain as names of themselves. Consider any pair of sets (X,X'), such that $X'\subseteq X$, as a partial interpretation of the predicate P: Every $a\in X'$ falls under P, and every $a\in X-X'$ falls under $\neg P$. If $a\notin X$, its falling under P is left undetermined. Assign, accordingly, a truth-value to every $P(a)$, where the truth-value for the case $a\notin X$, is **u**, i.e. "undetermined". The rest of the non-logical vocabulary is interpreted as in $\mathcal{M}$. First-order wffs in the language of **CL** that do not contain context operators can be now evaluated according to Kleene's three-valued logic.

Say that (X,X') is *compatible* with (1), i.e, with $\forall u_1,\ldots, u_k\ \phi(u_1,\ldots, u_k)$, if the value of (1) under the partial interpretation (X, X') is not **f** (false). This definition yields the right notion in the examples discussed above, and is justified, in general, by the following observation. The wff $[y_1,\ldots y_n]\ [t_1,t_2\ldots,t_m]\ \phi(u_1,\ldots, u_k)$ is satisfied in $(\mathcal{M}, f)$, under the assignment of c_i to u_i, $i = 1,\ldots,k$, and b_j to y_j, $j = 1,\ldots,n$, iff, $(X, f(X))$ is compatible with $\forall u_1,\ldots, u_n\ \phi(u_1,u_2,\ldots, u_n)$, where X consists of the b_i's and of the values of the t_j's under the assignment of c_i's to u_i's. (Here, $t_1\ldots,t_m$ are as in (1^*)). Consequently, if, for every context, X, $(X, f(X))$ is compatible with $\forall u_1,\ldots, u_k\ \phi(u_1,.., u_k)$, then scheme (1^{**}) holds, i.e.,

$$(\mathcal{M}, f) \models \forall y_{1,}\ldots,y_n\ \forall u_1,\ldots, u_k\ [y_{1,}\ldots,y_n]\ [t_1,t_2\ldots,t_m]\ \phi(u_1,\ldots, u_k).$$

Taking the above as our definition of compatibility, it is not difficult to see that, for a given ϕ, for each q and each $p\leq q$, we can say in first-order logic that the pair $(\{x_1,\ldots,x_q\}, \{x_1,\ldots,x_p\})$ is compatible with $\forall u_1,\ldots, u_k\ \phi$. A context X is *feasible* if there exists $X'\subseteq X$ such that (i) (X,X') is compatible with each of the finite number of universal sentences governing P's semantics, and (ii) for all a, $b\in X$, if $N_P(a, b)$, and if $a\in X'$, then $b\in X'$. Clearly, for each q, we can say in first-order logic, using the context-independent vocabulary, that $\{x_1,\ldots,x_q\}$ is feasible. Hence, we can form in **CL** the above scheme (TOL) that expresses the tolerance condition for $(\mathcal{M}, f)$. It can be also shown that if a context X is feasible so is every sub-context, $X'\subseteq X$. This concludes the general outline.

Returning to the usual examples, any context that is not too large must be feasible, since it is bound to contain gaps. In the case of 'walking distance', even under a generous modulus of tolerance, by which the predicate is insensitive to a 200' difference, we need at least 28 distances to span the interval from, say, 300' to 6000'; contexts of less than 28 distances are feasible. On the other hand feasible contexts can contain thousands of distances, provides that they have large enough gaps.

Conservativeness

The principle of conservativeness is that context change should have minimal effect. The cdf should not be capricious. An object that is classified as a *P* in one context should not be classified as a non-*P* in another, unless there is some pressing reason. In our case, such a reason is the need to satisfy tolerance. If a feasible context is enlarged by filling the gap that separates *P*'s from non-*P*'s with a dense chain, then some changes in the classification will be necessary. Say, in the context {3000', 4500', 6000'}, 4500' is walking distance and 6000' is not; now augment the context by adding 4510', 4520', …, etc., up to 5990' ; in this context, tolerance dictates that 4500' and 6000' be classified alike, hence there must be a change in the status of one of them. We might even want to change the classification upon adding a smaller number of scattered distances inside the gap. Conservativeness means that changes should be justified by such considerations. Consequently, if *a* is classified as a *P* (or as a non-*P*) in the context *a*, then, except for certain large dense contexts, the classification will endure. It takes a specially designed predicate, like 'winning finalist' in the example of section 5, to create situations in which small changes in small contexts can cause reversals.

While conservativeness is a wide, general principle, some of its implications are quite precise. Suppose that $c \notin X$, but that the classification into $f(X)$ and $X - f(X)$ determines (because of the semantic axioms relating to *P*) the classification of *c*. Then, the addition of *c* to the context *X* should not affect the existing classification of the members of *X*. The addition of *c* should not matter, since, in the context *X*, we know already how *c* is to be classified. Take again 'walking distance'. If $a \in X$ and, in the context *X*, *a* is a walking distance, then by adding to *X* a new member $< a$, we should not affect the existing classification of the members of *X*. Only the addition of members inside a separating gap may, sometimes, force changes. Formally, this is expressed by the scheme:

$$[x_1,x_2\ldots,x_n]\, P(x_1) \wedge y < x_1 \;\rightarrow$$
$$\bigwedge_{i=1}^{n} \{[y, x_1,x_2\ldots,x_n](P(x_i) \leftrightarrow [x_1,x_2\ldots,x_n]\, P(x_i)\}$$

By the same reasoning, or by iterating this step, the addition of many distances $< a$, should not reverse the existing classification of members of *X*. Such an addition, however, can create a dense chain that makes it impossible to reverse the classification of *a*. The effect of adding new members, whose classification is already implied by the existing one, is to reinforce the existing classification, making it harder to change.

8. Feasible Formulas and Proofs

The recursive evaluation that determines the truth-value of a given sentence in ($\mathcal{M}$, *f*) uses a part of the cdf, *f*, but not all of it. Consider, for example, the sentence

$$\forall x\,[x]\,(P(x)\to \exists\, y, z(P(y)\wedge\neg P(z))),$$

which says that, for all x, if, in the context {x}: $P(x)$, then there exist y and z such that, in the context $\{x,y\}$: $P(y)$, and, in the context $\{x, z\}$: $\neg P(z)$. To compute its truth-value, we need the values of f for all contexts of one and two elements, but not others. If we change this to: $\forall x\,[x](P(x)\to \exists y, z(P(y)\wedge[y]\neg P(z)))$, then, in the new sentence, $\neg P(z)$ is in the scope of [y], hence we need the value of f also for all three-element contexts.

Say that α *refers to the context X*, if, in the recursive evaluation that determines α's truth-value, the value of f for X is needed. This, it is easy to see, depends on $\mathcal{M}$, but not on the particular f. Note that α may refer to X, though $f(X)$ has no effect on α's truth-value, e.g., the truth-value of a logical truth containing P is independent of the interpretation, but the sentences refers to some contexts; $\forall u(P(u)\to P(u))$ refers to all one-element contexts. The set of contexts referred to by a sentence is determined, syntactically, as follows. Associate with every wff, α, a set, $cont(\alpha)$, of *associated context lists*, by the recursion:

- $cont(P(t)) = \{t\}$ ('t' denotes also the one-element list)
- $cont(\alpha) = \emptyset$, if α is atomic and its predicate is context-independent,
- $cont([C]\alpha) = \{C,C' : C' \in cont(\alpha)\}$,
- $cont(\alpha*\beta) = cont(\alpha) \cup cont(\beta)$, for every binary connective $*$,
- $cont(\neg\alpha) = cont(\forall x\, \alpha) = cont(\exists x\, \alpha) = cont(\alpha)$.

Then a sentence α refers to the context X, iff, X is the set of values (in $|\mathcal{M}|$) of some context list, C, associated with α, under some assignment of values to the variables occurring in C.

The notion extends to wffs, provided that we add an assignment of values (in $|\mathcal{M}|$) to the free variables of the formula; the contexts referred to by the wff depend on such an assignment.

Call a sentence α *feasible* if it refers to feasible contexts only. Call a wff feasible if its universal generalization is a feasible sentence. Call a proof in **CL** feasible if it consists of feasible wffs. In ordinary usage, tolerant predicates are meant to be used only in feasible sentences. It is desirable, of course, that, restricted to this portion of the language, the deductive part of the system should capture the semantics. And this indeed is the case:

Feasibility Theorem: If $\Gamma \models \beta$, where Γ is a set of feasible wffs and β is feasible, then there is a feasible proof of β from Γ.

This is not the whole story. We have seen that, in specific cases, the semantics of P is governed by axioms, which are universal sentences of the form (1^*). These sentences are feasible; but the scheme (1^{**}), which enforces their satisfaction in all contexts, has non-feasible instances (those in which the number, n is too large). Also, the phrasing of the tolerance scheme (TOL) covers unfeasible instances. Now there is a strengthening of the feasibility theorem by which all feasible consequences of (1^{**}) can be derived from its feasible instances and the same holds for (TOL). This means that, as long as we are interested in feasible consequences, the additional axioms that characterize specific systems can be restricted to feasible sentences. We can therefore do all our reasoning within the feasible part of the system.

The following completeness result for **CL** yields, as a corollary, the feasibility theorem, and plays a part in proving the strengthened version just mentioned.

Localized Completeness: If $\Gamma \models \beta$, then, there is a proof of β from Γ consisting of wffs whose associated context lists are equivalent to instantiations of context lists associated with wffs of Γ or with β. (Two context lists are equivalent if they contain the same terms. An instantiation of a list is one obtained from it by substituting terms for variables.)

Possible Developments of CL

A considerable enhancement of expressive power is obtained by adding to **CL** quantifiable variables ranging over contexts, say $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{X}_1, \ldots, \boldsymbol{Y}_1, \ldots$, which can occur in the context operator; along with them we add a membership symbol that can occur in wffs in the form: $x \in \boldsymbol{X}$. The contextual variables range over the finite subsets of the domain in question. This obviates the need for using schemes that cover an infinite number of instances, such as (TOL); we can use instead single wffs that are straightforward translations from the English. Thus, instead of using the wffs $Fsble_q(x_1, x_2 \ldots, x_q)$—each saying, for a particular q, that $\{x_1, x_2 \ldots, x_q\}$ is feasible—we can say that a context is feasible by a single wff $Fsble(\boldsymbol{X})$. Schemes such as (1^{**}), obtained by varying the number of variables in operators of the form $[y_1, y_2 \ldots, y_n]$, can be now replaced by single wffs containing operators of the form $[\boldsymbol{X}]$.

The price for this convenience is that we no longer have a complete deductive system, since the system has the expressive power of weak second order logic (second order logic in which the second order variables range over finite sets only). Completeness is regained if, disallowing quantification over context variables, we treat them schematically; that is, the derivation rules allow to substitute $\boldsymbol{X}$, in any theorem, by an arbitrary context list. Essentially, this version amounts to a convenient notational variant of **CL**.

Further extensions come naturally to mind, such as the inclusion of infinitary contexts, (where this involves second order logic), or ordered contexts, in which the ordering of the members matters, or, more generally, in which there is some additional structure on the members of the context An investigation along these lines should be motivated however by specific examples.

9. CLV: Combining Tolerance with Vagueness

In order to model vagueness, we extend the language of **CL**, by adding a modal definiteness operator, Δ, according to the outline sketched in section 3. Call the resulting system **CLV**. Instead of one structure $(\mathcal{M}, f)$, we have a family $(\mathcal{M}_i, f_i)_{i \in I}$, each constituting a possible world, with an appropriate (reflexive and symmetric) accessibility relation. If only the tolerant predicates are vague, then the possible worlds differ only in the semantics of these predicates, i.e., in the f_i's. For simplicity, assume that this is the case. Accordingly, let the family of structures (each referred to as possible worlds) be

$(\mathcal{M}, f_i)_{i\in I}$. Assume also that P is the only tolerant predicate, and take 'walking distance' as an illustration.

The universal sentences governing P's semantics, and the tolerance condition (TOL), are true in every (M, f_i). Also conservativeness is imposed in all possible worlds. On the syntactic level, this means that various universal sentences are included as axioms, which are subject to the necessitation rule. Among these are axioms that fix certain objects (distances) as P's and others as non-P's, e.g., $P(200')$, $\neg P(20.000')$. These are true in every context, in every possible world. Other parameters can depend on the possible world; in particular, the maximal distance that—in all contexts—is a walking distance, and the minimal distance that—in all contexts—is a non-walking distance, can vary from one world to another. In this way we get vagueness of the end points of the interval from walking to non-walking distances. And, of course, for a given context, the cut that separates walking from non-walking distances can, in general, depend on the world; in particular, the truth-value of $P(c)$—that is, $P(c)$ in the context $\{c\}$—can vary. Higher-order vagueness is modeled as described in section 3. By thus combining the contextual dimension with the modality of definiteness, we produce the effects that characterize tolerant and vague predicates. Both components are required: the contextual machinery—for semantic tolerance, the modal apparatus—for semantic modality.

The observations made in 3.1 about sharp modeling of vagueness apply also to tolerance. Our usage determines only basic features of the corresponding formal systems, leaving finer details undetermined. And there is also the inevitable slack between the precise model and the rough actuality. The contextual dependency of P is transmitted to wffs in which P occurs. For example, assuming that P takes numeric magnitudes as arguments let $\phi(x)$ be:

$$\forall\, y\, [P(y) \rightarrow P(y+x)]$$

This wff says that P is insensitive to changes of x. Sorensen [1994] proposes a paradoxical Sorites argument for a predicate of this kind.[29] The context dependency of ϕ resolves the paradox. On the other hand, it can be argued that the nearness relation, N_P, is itself tolerant. This would be another matter since the nearness relation is not defined in terms of P, but is a presupposed sharp relation that determines the notion of tolerance. We can extend our formal system by treating N_P as tolerant, which would require the introduction of a second level nearness relation N_{N_P}, or say, N_P^2; if desired, we can go further and introduced a nearness relation N_P^3 for *that*; and so on ad infinitum. All of this is doable and gives rise to complex systems, in which contexts are more complicated entities than those of **CL**. But the move is unnecessary. The nearness relation is not one of the items to be modeled, but part of the theoretician's machinery. For the purpose of modeling tolerance, N_P can be chosen as a suitable sharp relation. It is true that, if N_P yields plausible results, so will a sufficiently small modification of it. This is simply one of the aspects of the slack phenomenon.

Bibliography

[29] Sorensen's example is somewhat different, but the difference does not matter for the present point.

In the following *Vagueness Reader* refers to: *Vagueness: A Reader*, ed. R. Keefe and P. Smith, MIT Press, 1996. Places in articles reprinted in it are referred by page numbers in the book.

Dummett, M. 1975, "Wang's Paradox" *Synthese* 30: 301-324. Reprinted in the *Vagueness Reader*.

Cornelius, H. 1897, *Psychology als Erfahrungswissenshaft*, Leipzig Teubner, , xv + 445 pp.

Evans, G. 1978, "Can there be vague objects?" *Analysis* 38:208. Reprinted in the *Vagueness Reader.*

Fine, K. 1975, "Vagueness, Truth and Logic" *Synthese*, 30: 265-300. Reprinted in the *Vagueness Reader.*

Gaifman, H. 1992, "Pointers to Truth" *The Journal of Philosophy*, pp. 223-262,

Gaifman, H. 2000, "Pointers to Propositions", in *Circularity, Definition, and Truth*, edited by Chapuis A. and Gupta A., Indian Council of Philosophical Research, New Delhi, pp. 79 --- 122.

Gaifman, H. 2001 "Contextual Logic and Its Applications to Vagueness" *The Bulletin of Symbolic Logic*, September 2001, vol. 7 No 3, p.241 (Abstracts of Invited Talks of the Annual 2001 Meeting of the ASL)

Kamp, H. 1981, "The paradox of the heap" in *Aspects of philosophical logic*, ed. U. Mönnich, Dorderecht: Reidel.

Kaplan, D. 1989, "Demonstratives'' in *Themes from Kaplan*, ed. J. Almog, J. Perry and H. Wettstein, Oxford Un. Press.

Kofka, K. 1922, "Perception : an introduction to the Gestalt-Theorie", *Psych. Bull*. vol. 19, 531-585; reprinted in *Classics in Psychology*, ed. Thorne Shiple , 1961, Philosophical Library.

V. McGee and B. McLaughlin, 1994, "Distinctions Without a Difference", in *Vagueness*, supplement to *The Southern Journal of Philosophy*, 1994, Vol. XXXIII, ed. T. Horgan, pp. 203-251.

Putnam, H. 1983, "Vagueness and alternative logic" *Erkenntnis* 19: 297-314.

Raffman, D. 1994, "Vagueness without paradox" *The Philosophical Review*, 103.

Russell, B. 1923, "Vagueness" *Australasian Journal of Philosophy and Psychology*, 1: 84-92. Reprinted in the *Vagueness Reader*.

Sainsbury, M. 1990, *Concepts Without Boundaries*, an Inaugural Lecture, King's College, London. Reprinted in part in the *Vagueness Reader*.

Sorensen, R. 1994, "A Thousand Clones", *Mind* 103: 47-54.
Stumpf, C. 1883, *Tonepsychology I*, Leipzig Hirzel, xiv + 427 pp.

Unger, P. 1979a, "I do not exist" in *Perception and Identity*, ed. G.E. Macdonald, 235-51, London, MacMillan.

Unger, P. 1979b, "There are no ordinary things" *Synthese*, 41: 117-54.

Williamson, T. "Vagueness and Ignorance", Proceedings of the Aristotelian Society, Supp. Vol. 66. Reprinted in the *Vagueness Reader*.

Williamson, T. 1994, *Vagueness*, Routledge.

Wright, C. 1975 "On the Coherence of Vague Predicates" *Synthese*, 30: 325-65. A shorter version appeared as "Language-mastery and the sorites paradox" in *Truth and Meaning*, ed. G. Evans and J. Mcdowell, Oxford 1976, reprinted in the *Vagueness Reader*.
Vaguenes : A Reader, edited R. Keefe and P. Smith, MIT press, 1996

Wright, C. 1987 "Further Reflections on the Sorites Paradox", *Philosophical Topics*, 15:227-90. Reprinted in the *Vagueness Reader*.

Wright, C. 1994 "Epistemic conception of Vagueness" in *Vagueness*, supplement to *The Southern Journal of Philosophy*, 1994, Vol. XXXIII, ed. T. Horgan, pp. 133-159.

Zeeman 1961 "The topology of the brain and visual perception" in *Topology of 3-manifolds and related topics*, Proceedings of The University of Georgia Institute, 1961, Prentice Hall, pp. 240-256.