

Vagueness, tolerance and contextual logic

Haim Gaifman

Received: 13 April 2009 / Accepted: 12 May 2009 / Published online: 9 February 2010
© Springer Science+Business Media B.V. 2010

Abstract 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. One is tolerance: the insensitivity of predicates to small changes in the objects of predication (a one-foot increment of a walking distance is a walking distance). The other is the existence of borderline cases. The paper shows why these should be treated as different, though related phenomena. Tolerance is formally reconstructed within a proposed framework of contextual logic, leading to a solution of the Sorites paradox. Borderline-vagueness is reconstructed using certain modality operators; the set-up provides an analysis of higher order vagueness and a derivation of scales of degrees for the property in question.

Keywords Vagueness · Tolerance · Contextual logic · Semantic indeterminacy · Sorites paradox · Higher order vagueness · Degrees

1 Overview

This paper represents work done mostly in the period 1996–2002, which has been on my website since the beginning of 2002.¹ The idea of using syntactically represented context operators, which originated in this work, has been applied for other purposes

¹ 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

H. Gaifman (✉)
Philosophy Department, Columbia University, New York, NY, 10027, USA
e-mail: hg17@columbia.edu

in a paper on contextuality, Gaifman (2008). There is no overlap between that paper and the present one.

The present work contains an analysis of basic reasoning patterns that are characteristic of vague predicates. The analysis leads to rigorous reconstructions of the phenomena within formal systems. The first main thesis of the work is that one should sharply distinguish two different kinds of phenomena associated with vagueness: *tolerance* and *borderline-vagueness*. Tolerance is the insensitivity of predicates to sufficiently small changes in the objects of which they are predicated. A walking distance is still a walking distance if we increment it by one foot (but not by 5 miles); a child is still a child one hour later (but not 5 years later); and so on. Since any big quantitative change can be produced by accumulating sufficiently many small ones, we get the Sorites paradox. Borderline-vagueness is the existence of borderline cases of the given predicate—objects of which it is not clear whether they should be classified under the predicate or under its negation. This is a phenomenon of semantic indeterminateness. Vagueness as it is usually conceived is manifested in the existence of borderlines. I shall refer to it as *vagueness per se*. Tolerance and, in particular, the Sorites paradox is a more specialized “technical” notion.

As a rule, philosophers who wrote on these subjects bundled together tolerance, and in particular the Sorites, with semantic indeterminateness. This is understandable, given that many of the standard examples display both tolerance and semantic indeterminateness. Indeed, there is a connection. Yet, I will argue, these are distinct aspects that require different treatments. It is not difficult to establish by direct analysis, using easily available examples from natural language, that vagueness per se does not imply tolerance. The non-implication from tolerance to borderline-vagueness is more difficult since, in natural language, tolerance is manifested by predicates with borderlines. Nonetheless, an analysis, aided by a specially constructed example, will show that, in principle, tolerance need not imply vagueness per se.

Tolerance turns out to be a contextual phenomenon, whose rigorous analysis is best done within a framework of *contextual logic*. This involves a formalism, which provides for explicit representation of contexts by means of operators that combine with wffs (well-formed formulas). The idea and the way it resolves the Sorites paradox can however be explained in a non-formal way, and this will be the first stage of the exposition. The more technical part contains a formal system, a semantics and a sound and complete deductive system. It extends first-order logic and fully preserves

Footnote 1 continued

Symbolic Logic, and its ideas are sketched in Gaifman (2001). The results on borderlines, higher order vagueness, and *KTB* originated in the summer of 2001 and were 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. I would also like to thank Robert van Rooij, for his patience, without which this work would not have been published in Synthese. I also thank him for calling my attention to Klein (1980), cf. footnote 24.

the classical rules. One can apply it to vocabularies that include tolerant predicates, such as ‘poor’, ‘walking distance’, ‘noonish’, etc.^{2,3}

The more common phenomena of borderline-vagueness is treated in the second part of the paper. Also here there is an essential contextual factor. But this factor “mixes” with other factors and I found it more useful to use the framework of modal logic to clarify what goes on, in particular, the phenomenon of higher-order vagueness.⁴

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*.

To prevent misunderstandings, I should point out that, while I find the formal apparatus interesting by itself, it is used here as an instrument of analysis. It models, in idealized form, some basic aspects of the way language functions. The acknowledged distance between actual practice and the modeling is no more problematic than the distance between actual practice and formal first-order logic. In the second part of the paper I shall say something more of the modeling of vagueness by means of precise formalisms.

Another point I should briefly make here concerns the distinction between vagueness as a linguistic (or semantic) phenomenon and the so-called ontological vagueness. The approach pursued here is, of course, linguistic. Yet an analysis of such a basic aspect of language reveals something important about the relation between language and what it describes. The temptation to project crudely from the language/world to the world can lead to dubious problems, which arise out of a clash of metaphors without sufficient backup. It is not that there are no vague objects, but that the very question

² 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 Sect. 3.2 for further comments on the difference.

³ 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 led to the present system. I think it may succeed on the sentential level but would fail when it comes to quantifiers.

⁴ It goes without saying that my approach is logical, rather than linguistic. While there is considerable overlap between the two, the orientation is different. For a linguistic approach to degrees see Klein (1980) and footnote 24. No account of linguistic-oriented works on vagueness has been attempted here.

is bad, given the way in which it has been viewed. A reining in of some pictorial metaphysical tendencies 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.

In the next section I start with a general discussion of vagueness and tolerance; it includes the argument that vagueness per se does not imply tolerance. I shall then proceed to part I. This part, which consists of Sects. 3.1–3.5 deals with tolerance, the Sorites and the framework of contextual logic where it is formally represented. Part II consists of Sects. 4.1 and 4.2. It deals with borderline vagueness and its treatment in modal logic. Each part starts with a non-technical philosophical analysis, followed by the formal modeling. I hope that the presentation will give an adequate picture of the setup to readers who prefer to skip the more technical details. I shall conclude by outlining a combined system within which both tolerance and borderline-vagueness are represented.

2 Borderline–vagueness and tolerance

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.⁵ But nowadays, the term is used in philosophy to denote indeterminateness, exemplified by cases where the semantic rules seem to leave it open whether some predicate $P()$ is true of some particular object, a ; or, in general, whether a given description applies to some given case. Moreover, 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 to the question: $P(a)$? 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 not 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.⁶ This gives rise to the Sorites paradox, which in its original form was based on the tolerance of the concept *heap*: a heap of sand remains a heap after the removal of one grain.⁷ It 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

⁵ E.g., Russell (1923) conflate vagueness with ambiguity. This, however, is not a confusion stemming from loose usage, but a result of Russell’s attempt to define vagueness within his metaphysical framework.

⁶ The term ‘tolerance’ has been used in engineering to mark the amount of permissible deviation from sharply specific values. It was also used by Zeeman (1961) in a somewhat related sense. (I thank Peter 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).

⁷ ‘Sorites’ derives from ‘soros’, which in Greek means a heap. The paradox is probably due to Eubulides who lived in the middle of the fourth century BCE.

falls under P , so does the other. Applying instances of (iii) a sufficientl large number of times, we can deduce from (i) that the last object falls under P (e.g., a 20-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 \(1979a,b\)](#) has voiced even a more extreme position. Let me f rst take a look at semantic indeterminateness.

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;⁸ (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 the concrete expressions of semantic indeterminacy and they can serve us as a guide in the analysis. (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 current term).

The very terminology—'borderline', 'penumbra'—is based on geometrical metaphors that suggest distance. The problem of classifying borderline cases can be seen as the problem of determining whether an object is sufficientl near, or similar (in the relevant respects) to some other object; sufficientl near to warrant the classificatio of both in the same class: if one falls under the predicate so does the other. 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 Sects. 3.1 and 3.2) the crucial contextual element that underlies tolerance. For the moment let me make a more obvious point: indeterminateness does not imply tolerance. It should be clear that conceptually (that is, without the geometrical metaphor) indeterminateness per se need not involve tolerance; moreover, there need not be an implication even when the indeterminateness involves some notion of "distance". Here are some of the many examples that support the claim.

First, many cases of indeterminacy lack the kind of scaling that tolerance, in any appreciable sense, requires. The classificatio of a newly discovered object, or a new situation, can be indeterminate, because the item lacks sufficien 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](#),⁹ 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

⁸ By this I mean the usual kinds of facts that a competent speaker will fin 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.

⁹ [Hart \(1961\)](#). I am indebted to my former student Robert T. Miller for this example.

skateboards share relevant 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 justify 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.¹⁰ (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 baffl its classifie , by possessing combinations of “incompatible” features. The platypus is a celebrated case; like a bird it has a bill and, like birds and 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 classificatio 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 *monotremes*, which it shares with two other living, similarly “strange”, species of *echidna*.¹¹ 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-define 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 finger 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; one qualifie as a small fraction, four obviously does not, but three 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 3–15), ‘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

¹⁰ A physical configuratio 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.

¹¹ I thank Laura Franklin for calling my attention to these biological details.

and the difference between 40 and 41 chimpanzees is as clear-cut as the difference between 5 and 6 fingers— *ut because in the contexts in which the predicate is used the difference can be practically ignored*. I shall have more to say on this in the next section when I take a closer look at tolerance.

3 Part I: tolerance, the Sorites and contextual logic

3.1 A closer look at tolerance

A predicate is tolerant, to the extent that sufficientl near objects are classified alike with respect to it. What counts as “sufficientl near” depends on the predicate. Let us associate with the tolerant predicate P a nearness relation, $N_P(\cdot)$, and express P ’s tolerance by

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

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 $P(x) \longleftrightarrow P(y)$. 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); if we fix the basic unit we can construe the predicate as taking numbers as arguments (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. Assume that the nearness relation is such that there are sequences, a_0, a_1, \dots, a_n , for which $N_P(a_i, a_{i+1})$ for all $i < n$, 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 non-contextual 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 to do justice to tolerance, as a semantic phenomenon; they fail to account for the fact that conditionals of the following kind are utterly compelling:

$$(SC) \quad \text{If John was young one second ago, then John is young now.}$$

I shall argue that, far from being illusions, these conditionals are normative rules that are part of 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. For simplicity, let ‘sharpening’ mean henceforth a sharpening that is admissible. In the supervaluation semantics a sentence is true (or false) if it is true (or false) under all sharpening of the interpretation of P . Sentences that are true for some sharpening and false for some other do not have a truth-value.

Obviously, a Sorites conditional that does not involve borderline objects is true for every sharpening. Every borderline object is classified under the predicate in some sharpening, and under its negation—in another. Hence each Sorites conditional that involves a borderline object is falsified by some sharpening which puts a cutoff between the object and its neighbor. It is also true for every other sharpening that puts the cutoff elsewhere. Hence each conditional that involves a borderline object is neither true nor false. The conjunction of all conditionals is false however (since in every sharpening some conjunct is false). Hence, we can hold the conjunction false, which avoids the paradox. On the other hand we need not hold any particular conditional false. 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 not imply a disjunction of the wffs $\neg(P(a_i) \rightarrow P(a_{i+1}))$. Neither of these two accounts gives good reason for the compelling intuition that (SC) and its ilk are true. If, as I shall later argue, the truth of the Sorites conditionals is part of the meaning of tolerant predicates, the supervaluation approach fails.

One may invoke the fact that, for each particular time point, (SC) is true in every sharpening except one. This explanation, appealing as it does to the proportion of sharpenings that satisfy a given sentence, belongs to the degree-theory framework, which I shall discuss shortly. Let me first comment on a 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 distance, which is the maximal walking distance; if you add to it one foot, or even one inch, it is no longer 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 of borderline-vagueness, in particular, with regard to the truth concept and the identification of truth with *supertruth* (truth in every sharpening). The semantic modality theory of the second part of this paper provides an interpretation of *definite truth*, without problematic side effects, which fully preserves classical logic. To the extent that this approach is successful, the above 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 time points determined sharply up to a second, which are unknown, or even unknowable, to human beings. Most people will find such an account absurd (and I must agree). Perhaps epistemicism is best read as a directive to behave *as if* there were sharp unknown boundaries for vague terms, because in various cases this seems to fit our practices. Indeed, the sharpening of vague terms is not an arbitrary matter. Zoologists were looking for the right classification of platypus. Astronomers have been arguing whether Pluto falls under the concept *planet*. These arguments are sometimes phrased as argument about facts. There is an element of discovery in finding the best fit between our language and the world. The epistemicist error is to treat in this way the sharp determination of ‘around 3 o’ clock’, or the heart beat at which one’s childhoods ends. To do so is to ignore a basic aspect of linguistic usage.

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 15%; because the truth of my guess is accidental. In that case, I only know that the number of people is at least 170. Williamson therefore suggested that we replace the original Sorites conditional, $P(a_i) \rightarrow P(a_{i+1})$, with the weaker epistemically acceptable variant:

$$(EC_i) \quad K(P(a_i)) \rightarrow P(a_{i+1}),$$

where ‘ K ’ is the knowledge modality (‘ $K \dots$ ’ means that the presupposed agent knows that ...). 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.¹² 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’ clarifies why we find the unmodified conditional (SC) compelling. 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 belongs, as I mentioned above, to the degree-theory brand.

¹² From $K^n P(a_0)$ and $K^{n-1}(EC_i)$, for all $i < n$, we can deduce, via standard modal logic, $P(a_n)$; e.g., $K^{n-1}(K P(a_i) \rightarrow P(a_{i+1}))$ implies $K^{n-i-1}(K P(a_i) \rightarrow P(a_{i+1}))$, which implies $K^{n-i} P(a_i) \rightarrow K^{n-i-1} P(a_{i+1})$; putting all of these together, n applications of modus ponens yield $P(a_n)$. To avoid the paradox we must assume that either $K^n P(a_0)$, or some $K^{n-1}(EC_i)$ 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 $K P(a_0)$, and $K(EC_i)$, 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.

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 chosen some other precise interval, but this would not have come to the same. ‘Nearly’, ‘approximately’, ‘roughly’ and similar modifier are systematic *blurring* devices, used for the purpose of ruling out commitments to sharp boundaries. Let me coin the technical term ‘fuzzify’. The point of fuzzifying is not only to eschew sharp boundaries, but to rule out any criterion 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% of competent speakers (or any other reference group) disagree. 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, where “too much” is itself vague. This guarantees sufficient latitude. The function of the modifier ‘around’ is to introduce such latitude. Imagine a superhuman computer, who, taking as input all English speech acts ever performed, derives sharp interval (say up to a second) for ‘around 3 o’ clock’; this, again, would make no difference, since the computer’s prescription is not binding. And if some authority were to make it binding, it would amount to a legislative act that changes the meaning of ‘around 3 o’ clock’. We would have in that case sharpened the vague predicate, as we occasionally do for scientific, legal or political reasons. This of course would make the modifier ‘around’ completely pointless.

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 beyond our ken. A physical theory can imply that certain facts are not knowable to us. But in all such cases, there is a rich system, and a well-developed methodology for establishing truth, that gives these statements meanings. 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, or something that might motivate the adoption of CH as an axiom. But if 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 change with 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 (which 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 Sorites considered so far, the degree theorist's is the most satisfying. There is no denying the graded nature of vague predicates: the aptness of applying them is usually 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 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) requires more in way of justification.¹³ 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, need not necessarily 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 theories, 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 the use 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, which he prefers to the slower and 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 preferable than taking any larger number. Some degree theorists compare the Sorites to the lottery paradox.¹⁴ 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.

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 . . . '. There is nothing semantic about it.

¹³ 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.

¹⁴ McGee and McLaughlin (1994, pp. 221, 222).

To recap, the non-contextual approaches to vagueness explain one's acceptance of the Sorites conditionals as a kind of illusion, understandable—but nonetheless an illusion. They fail to come to grips with the fact that the conditionals express norms of linguistic usage. The normative aspect becomes clear when we consider scenarios in which the norm is contravened, and how bizarre they are. 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.

Tolerance means that the Sorites conditionals are valid. In practice we avoid a contradiction by not stringing together too many of them in a single deduction. The conditionals are used, subject to such a restriction. Non-philosophers are amused but undismayed by the paradox; they know that each conditional is valid but stringing too many of them is illicit. This is the brute-force solution to the paradox. It is different from the solution I shall propose, but it indicates the right direction. Let me add that there is nothing incoherent in imposing some such restriction. 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 use. We have one-shot (disposable) gadgets, and, in principle, a device can come with a manual containing an instruction: “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 intended “meaning”. A theodolite is a portable instrument for measuring angles, used in land surveys. Each measurement has a certain margin of error. If we try to find the angle between l and l_n by adding n measured angles: between l and l_1 , l_1 and l_2 , and so on, we may get an accumulated error, large enough to cause contradiction in our data. Those who use theodolites avoid the adding up of too many measured angles, where “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”.¹⁵

Now consider the restriction: *a derivation should not use too many Sorites conditionals*, where “too many” means that the set of conditionals makes 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 is 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 this restriction is respected.

¹⁵ In constructing a structure from many pieces, knowledge of the tolerances allowed for the pieces is crucial. A blue print for a machine is practically useless without tolerance specifications

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, tolerant predicates 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 single occasion, if we have to decide whether 5280' is walking distance and also whether 5281' is walking distance, the two answers had better be the same. We are safe from wrong answers—such as “10 miles is walking distance”—since no single occasion demands answers regarding each and every item in a Sorites chain.

The brute force approach is unsatisfactory since it imposes a restriction on proofs, without making explicit the underlying contextual element of local usage. 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 my proposed logical system, is a set consisting of objects that, on a single occasion, have to be classified with respect to the predicate. Tolerance is the requirement that near enough objects be classified alike. A context for which there exists a classification satisfying the requirement, as well as 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 see that also 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 loses its tolerance: it has some sharp cutoff. But unfeasible contexts do not arise in practice, as long as the predicate is not explicitly sharpened.

3.2 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 nineteenth century. [Stumpf \(1883\)](#) notes cases where a subject who experiences two stimuli, a and b , judges them to be equal; also b and c are judged equal; but when experiencing a and c , the subject reports that $a < c$. This contradicted accepted theories at the time, on which stimuli are classified by absolute intensity. [Kofka \(1922\)](#) who reports this, 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 Stumpf’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: $[\dots]$, where ‘ \dots ’ is a list of the compared items. Then, the subject’s reports are written as follows, where ‘ $i(x)$ ’ denotes x ’s intensity:

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

These do not imply a contradiction; 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 $[a, b, c](\dots)$. 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 classification of single items, by humans, depends on the end of the chain from which the item is approached.¹⁶ [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.¹⁷

In the earlier philosophical literature much emphasis has been laid on psychological factors as the source of tolerance. The most standard example of a Sorites chain is a sequence of colors. In this and in similar perceptual cases, tolerance might appear an inevitable outcome of perceptual limitations: we judge near enough items equal because we do not perceive the difference. We deduce the difference between indistinguishable items, a , b , only by bringing into play additional items: 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 perceptual emphasis is misplaced, since vagueness, tolerance and the Sorites do not depend on perceptual factors. Tolerance in ‘walking distance’, ‘rich’, ‘old’, ‘noonish’, ‘a large community of chimpanzees’, and their like does not hinge on indistinguishability; 4017’ and 4018’, or 50 and 51 chimpanzees, are as distinguishable as are 1 and 2. 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 this: We can ignore the distinction for the purposes for which these predicates are employed and by ignoring it we achieve an enormous gain in efficiency. The appeal

¹⁶ 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 I could not locate a reference. I will be thankful to any reader who can direct me to it.

¹⁷ Raffman’s approach is psychologistic and informal: the context is constituted by the short history of the agent’s impressions. The approach does not involve any logic or a notion of a feasible contexts.

to perceptual limitations, or to any psychological brute facts, is unfortunate, because it construes justifiable norms in terms of brute psychology.

The logic that can handle contextual effects 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, which uses linguistic and logical devices that are suited to work efficiently under the 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. I shall start with some general remarks on the proposed representation of contexts.

3.2.1 Representing contexts

The system is based on the following notation:

$$(C) [C]\phi,$$

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. 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", that is, 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. Now expressive power is not always a virtue. Temporal and dynamic logic are meant to be used for program verification where simplicity and easy application are crucial; if this can be done within limited setups so much the better. 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, which can be amended by incorporating context operators. In [Gaifman \(2008\)](#) there is a system that incorporates, among other things, context operators for proper names and various indexicals.

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 another token of that sentence, on another line, is true. In that case, the place where the sentence-type is written plays the role of context. That 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, that kind of context dependency is not of the indexical type.

3.2.2 Contexts and feasible contexts

If P is a monadic tolerant predicate, then our contexts will be finite lists of objects that are to be classified as P ’s and non- P ’s. To be sure, many other contextual factors determine the use of tolerant predicates, as they do in general. Walking distances are much longer in the Olympic village than in old people’s resorts. But these do not concern us and, if needed, they can be handled by adding other context operators of the form [C].

From now on, unless indicated otherwise, C in the context operator [C] is a finite list of terms denoting objects that are to be classified with respect to P . For convenience I speak occasionally of the context C , meaning the set of these objects, and I shall use ‘ a ’ ‘ b ’ and other terms, as names in the formal language as well as in my metalanguage. The intended meaning should be obvious. The tolerance conditional of the previous section,

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

is now to be relativized to the context C as:

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

From

$$N_P(x, y) \rightarrow [C](P(x) \rightarrow P(y)) \quad \text{and} \quad N_P(y, z) \rightarrow [C](P(y) \rightarrow P(z))$$

we can derive the conclusion: $N_P(x, y) \wedge N_P(y, z) \rightarrow [C](P(x) \rightarrow P(z))$. This requires that the context in the two conditionals be the same. Generally, if C is different from C' , we cannot derive something of this kind from:

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

A context is *feasible* if its members can be partitioned into P 's and non- P 's without violating any instance of the tolerance conditional or any of the semantic axioms governing P (for example, a distance smaller than a walking distance should be a walking distance). The non-violation of the tolerance conditionals means that we do not introduce any sharp cutoff, that is, a cut between some a and b for which $N_P(a, b)$ holds. This means that there should be at least one sufficiently large gap that can serve as a divisor between P 's and non- P 's. For any given C we can express in the formal language the condition that C is feasible. We can add a formula that expresses the statement that C is feasible, as a conjunct to the antecedent of (TC*), this will give us the scheme: $C \text{ is feasible} \wedge N_P(x, y) \rightarrow [C](P(x) \rightarrow P(y))$. The satisfaction of this scheme guarantees that P is tolerant in all feasible contexts.

Tolerance is, by definition, tolerance in feasible contexts. An unfeasible context is one in which we would be forced to sharpen the predicate to a degree that makes it no longer tolerant. For example, we might be required, on the same occasion, to decide for each of the 2,640 distances $d_n = n \times 10'$, $n = 1, \dots, 2,640$, which span the distance from 10' to 5 miles, whether it is a walking distance; assuming that 5 miles is not a walking distance, this would force us to specify with precision $\pm 5'$ the maximal walking distance. We do not expect to encounter unfeasible contexts; nonetheless context operators, $[C]$, for unfeasible C 's, are included in the language so as to round it off and avoid syntactic complications. Only the feasible part of the language—that which does not involve unfeasible contexts—is needed for actual usage and for actual reasoning.

I shall now consider an example that shows (as promised in the introduction) that tolerance need not imply vagueness. A certain educational institution awards, every 4 years, a highly prestigious prize. The finalists—those who pass very demanding preliminary tests—accumulate scores, from 1 to 20, by passing final examinations. The rules are the following:

- (i) Any finalist who scores 19 or 20 wins the prize.
- (ii) A minimal score of 16 is required to win the prize.
- (iii) A finalist whose score differs by 1 from a winning finalist also wins the prize.

The motivation for (iii) is the wish to avoid the possibility that an accidental small difference be construed as discrimination between the candidates. (iii) is obviously a tolerance constraint on the predicate 'winner of the prize'. There is a Sorites worry that the rules will lead to contradiction if there are five finalists with scores 15, 16, 17, 18, 19. This worry is ignored, since such a scenario is extremely improbable. The

preliminary screening examinations are so hard that in any given year the number of finalist is too small to lend the possibility of a Sorites chain appreciable chance.

Thus the predicate ‘winner of the price’ is tolerant, where two people are sufficiently near iff they are finalist at the same year whose scores differ by 1. Obviously the predicate is as non-vague as a predicate can be. The role of contexts and of feasible contexts is obvious here. The full context constitutes the finalist and their scores at each given year where the price is offered. Since only the scores matter, we can take as our contexts lists of scores. In the formalism proposed here any context containing the scores 15, 16, 17, 18, 19 is an unfeasible context.

Admittedly, this is an artificial example of a legalistic nature. Nonetheless it suffices to establish that, in principle, tolerance does not imply vagueness. In natural language the contexts are determined by the occasions on which a tolerant predicate is used, either in an assertion, an exchange, a thought, or in a single piece of actual (not philosophical-hypothetical) reasoning. The contextual element complicates matters to a high degree; in setting up a semantics, we try of course to avoid it. In natural language the role of context is manifest only in the case of indexicals and demonstratives, which carry their context-dependence on their sleeves. Now if $P(c)$ is true in some contexts, false in others, where the contextuality is not manifest, 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., started their life as tolerant predicates and became vague through our ignoring some precise contextual conditions. These predicates are both tolerant and vague, to start with. But the observation explains the connection between tolerance and vagueness in natural language.

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 Sect. 4.1.

3.3 TCL, tolerance contextual (first order) logic: the basic framework

We first provide a general formal setup within which tolerance can be treated. The way it is to be treated is described in Sect. 3.4.

TCL 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 $C = t_1, \dots, t_n$ and $C' = s_1, \dots, s_m$, then, $C, C' = t_1, \dots, t_n, s_1, \dots, s_m$; a single term is also regarded as a list of length 1, hence, $t, C = t, t_1, \dots, t_n$.

For notational convenience we include the empty list, where $[\]\alpha$ is, by definition, α .

3.3.1 The semantics: context dependency functions

The non-logical vocabulary includes also 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 indicate it as we go along.

A model for the language is a structure of the form (\mathcal{M}, f) , such that:

- (1) \mathcal{M} is a model for the context-independent vocabulary; the universe of \mathcal{M} is denoted as $|\mathcal{M}|$.
- (2) f is a function such that $f(X) \subseteq X$, for every finite subset, $X \subseteq |\mathcal{M}|$. The subset $f(X)$ 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 $X \cup Y$:

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

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:

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

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

Roughly, this means that the prefixin 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 and this is discussed in Sect. 3.4. Various important restrictions are expressible by sentences in the language of **TCL**. 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.

3.3.2 The deductive system

Any of the standard deductive systems for first-order logic can serve as a basis. A suitable enhancement of the basis can serve as a deductive system for contextual logic. All first-order axiom schemes and inference rules are retained, where the wffs range over the language of **TCL**, 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, we add the following axiom schemes:

- (I) $[x_1, \dots, x_n]\alpha \rightarrow [x'_1, \dots, x'_m]\alpha$, where $\{x_1, \dots, x_n\} = \{x'_1, \dots, x'_m\}$ (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) $R(x_1, \dots, x_n) \leftrightarrow [y_1, \dots, y_m]R(x_1, \dots, x_n)$, 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 $[t_1, \dots, t_n]\alpha$ as:

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

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

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

Here ‘ \vdash ’ and ‘ \models ’ denote, respectively, provability and logical implication in **TCL** (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).

3.3.3 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. **TCL** being a general framework, it leaves the contextual effect open. We can impose the desired connections between contextualization to different contexts through additional axioms. The enhanced system will be complete and sound, since the deductive system of **TCL** is. In Sect. 3.4, under ‘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.

3.3.4 The scopes of the context operators

Formalization in **TCL** may require decisions about contextual scope. The default reading of the wffs gives the context operators minimal scope, e.g., $P(t_1) \wedge P(t_2)$ 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, to give 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:

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

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.¹⁸ 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 a quantifier will be useless for the purpose of handling tolerance.

Quantification 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 $\forall x[x]\alpha$. A different reading takes into account the free occurrences of the quantified x in terms that occur under the context-dependent predicate. Let t_1, \dots, t_n be the list of terms that occur in α under P and contain an occurrence of x that is free in α . Then $\forall x\alpha$ is read as:

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

There is a variant, call it **TCL***, obtained by reading the quantifier in this way. It has a sound and complete deductive system, and it retains classical sentential logic. **TCL*** differs from **TCL** in the following items of quantification logic: The instantiation axiom (for legitimate substitutions) is changed to: $\forall x\alpha(x) \rightarrow [t_1, \dots, t_n]\alpha(t)$, where t_1, \dots, t_n is the above list of terms; 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 $\beta \rightarrow [t_1, \dots, t_n]\alpha(x)$, where x is not free in β and t_1, \dots, t_n is the above list, infer: $\beta \rightarrow \forall x\alpha$.¹⁹

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

¹⁸ 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**, i.e., the sentential fragment of **TCL**, is much more expressive than **CSL[#]** and its deductive system is an extension of classical logic, which **CSL[#]** is not.

¹⁹ In **TCL***, provably equivalent wffs, need not have provably equivalent generalizations; e.g., α and $\alpha \wedge \beta$, where β is a tautology containing, under P , variables not occurring in α . By applying a quantifier to $\alpha \wedge \beta$ we “pull out” these variables into the global context. Therefore, mere mention of an item, even in a “vacuous” way, may have a substantial effect, since it changes the context. Indeed, in ordinary speech we would not add, for no specific reason, a vacuous conjunct, such as $t = t$. Regarded within this perspective, this feature of **TCL*** may reflect actual intuitions.

3.4 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. Here we can leave various details open. For example, 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.²⁰ They impose either monotonicity (every person of equal or higher height than a tall person is tall), or anti-monotonicity (every distance smaller or equal to a walking distance is a walking distance), or convexity (every time point between two noonish time points is noonish).

We shall assume that, in general, the conditions can be expressed by finitely many universal first-order sentences (including, possibly, quantifier-free sentences). Their conjunction is equivalent to a single universal sentence. Hence, the semantic constraint can be put in the form:

$$\forall u_1, u_2, \dots, u_k \phi(u_1, u_2, \dots, u_k), \tag{1}$$

where ϕ is quantifier-free (if $k = 0$, this is a quantifier-free sentence). Our analysis and constructions apply to the general case. As an illustration consider ‘walking distance’ whose anti-monotonicity is expressed by:

$$\forall x, y (x < y \rightarrow (P(y) \rightarrow P(x))). \tag{i}$$

In order to express it in **TCL**, 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:

$$\forall x, y [x, y] (x < y \rightarrow (P(y) \rightarrow P(x))). \tag{i'}$$

This is imposed in all contexts via the scheme:

$$\forall y_1, \dots, y_n [y_1, \dots, y_n] \forall x, y [x, y] (x < y \rightarrow (P(y) \rightarrow P(x))), \quad n = 1, 2, \dots \tag{i''}$$

²⁰ 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’ define a pre-ordering of the class of people (different people can have the same height).

The satisfaction of the scheme means that, for every context X , $f(X)$ is an initial segment of X under the ordering $<$. The requirement that some small distance, c , is a walking distance and some large distance, c' , is not is imposed by adding as conjunct within the scope of $[y_1, \dots, y_n]$ the wff: $P(c) \wedge \neg P(c')$; thus we get:

$$\forall y_1, \dots, y_n [y_1, \dots, y_n] \forall x, y [x, y] \{ (x < y \rightarrow (P(y) \rightarrow P(x))) \wedge P(c) \wedge \neg P(c') \},$$

$$n = 1, 2, \dots \tag{i'''}$$

It is not difficult to show that this scheme (the set of all these sentences) is equivalent in **TCL** to the scheme:

$$\forall y_1, \dots, y_n \forall x, y [y_1, \dots, y_n] [x, y] \{ (x < y \rightarrow (P(y) \rightarrow P(x))) \wedge P(c) \wedge \neg P(c') \},$$

$$n = 1, 2, \dots \tag{i^*}$$

(The two context operators can of course be amalgamated into one.)

In the more general case, where the semantic condition is expressed by (1), we have, instead of (i), the wff:

$$\forall u_1, \dots, u_k [t_1, t_2, \dots, t_m] \phi(u_1, \dots, u_k), \tag{1'}$$

where t_1, t_2, \dots, t_m are all the terms occurring in ϕ under P .²¹ The final general scheme analogous to (i*) is:

$$\forall y_1, \dots, y_n \forall u_1, \dots, u_k [y_1, \dots, y_n] [t_1, \dots, t_m] \phi(u_1, \dots, u_k), \quad n = 1, 2, \dots \tag{1^*}$$

Unrestricted tolerance can be stated as the universal generalization of (TC*) (see Subsect. 3.2.2). But this should be modified since we want to impose it only for 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.^{22,23}

Take again 'walking distance' as a representative case, where the scheme (i*) imposes the semantic conditions for P . For a context, X , let $b_1 < \dots < b_j < \dots < b_{m-1}$ be all the members of X in ascending order, which are strictly between c

²¹ (1') depends on the terms occurring in ϕ under P . 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 unimportant when we pass to the scheme (1*).

²² N_P can be vague, but this is an aspect we are not concerned with now; it can 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.

²³ 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 reflect better our ordinary intuitions. It is also preferable not to have too many free parameters to play with.

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 convexity condition) should be obvious: two sufficiently large gaps are required, to separate the noonish from 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 $X' \subseteq X$.

In these examples, the condition that characterizes the feasibility of a context of q members is easily expressible by a first-order wff in the context-independent vocabulary of **TCL**. Let $Fsble(x_1, x_2, \dots, x_q)$ be this wff (in fuller notation there is subscript 'q', since the wff depends on q). Tolerance is then expressible by the scheme:

$$(TOL) \quad \forall x_1, x_2, \dots, x_q \{Fsble(x_1, \dots, x_q) \rightarrow [x_1, \dots, x_q](N_P(x_1, x_2) \rightarrow (P(x_1) \rightarrow P(x_2)))\}, \quad q = 1, 2, \dots$$

In the general case, 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 generalization (1), so that every two N_P -related objects are in the same part. This requires that we characterize 'compatibility with (1)' in a way that is expressible in first-order logic. It can be done, using the fact that (1) is a universal sentence. The following is a short outline of the

construction, where we assume a single monadic tolerant predicate P . The interest of it is that the idea can be generalized to languages that have several tolerant predicates, of any arities, provided that the conditions governing the semantics of these predicates are expressible by finitely many universal sentences. The general construction is not needed for the treatment of the standard examples; readers not interested in these details should skip it.

3.4.1 Outline of the general case

For simplicity, we shall regard the objects of the relevant domain as names of themselves. Let us regard 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. Accordingly, for each $a \in X$, assign to $P(a)$ the truth-value **t**, if $a \in X'$; the value **f**—otherwise. If $a \notin X$, assign to $P(a)$ the value **u** (undetermined). The rest of the non-logical vocabulary is interpreted as in \mathcal{M} . First-order wffs in the language whose non-logical vocabulary includes also the predicate P can be now evaluated according to Kleene's three-valued logic.

Define (X, X') to be *compatible* with $\forall u_1, \dots, u_k \phi(u_1, \dots, u_k)$ if the value of this sentence under the partial interpretation of (X, X') is not **f**. This definition yields the right notion for the simple examples discussed above, as well as in general. To see this, consider the wff $[y_1, \dots, y_n][t_1, t_2, \dots, t_m]\phi(u_1, \dots, u_k)$. Assign to each u_i a

member $c_i \in |\mathcal{M}|, i = 1, \dots, k$, and to each y_j a member $b_j, j = 1, \dots, n$. Let X consist of all the b_j 's and all the values of the t_l 's, $l = 1, \dots, m$ under the assignment of the c_i 's to the u_i 's. Observe that $[y_1, \dots, y_n][t_1, t_2, \dots, t_m]\phi(u_1, \dots, u_k)$ is satisfied in (\mathcal{M}, f) , under the assignment of the c_i 's and b_j 's to its free variables iff $(X, f(X))$ is compatible with $\forall u_1, \dots, u_n \phi(u_1, u_2, \dots, u_n)$. Consequently, $(X, f(X))$ is compatible with $\forall u_1, \dots, u_k \phi(u_1, \dots, u_k)$ for every context X , iff (1^*) holds, i.e.,

$$(\mathcal{M}, f) \models \forall y_1, \dots, y_n \forall u_1, \dots, u_k [y_1, \dots, y_n][t_1, t_2, \dots, t_m] \phi(u_1, \dots, u_k).$$

Now 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 $(\{x_1, \dots, x_q\}, \{x_1, \dots, x_p\})$ is compatible with $\forall u_1, \dots, u_k \phi$. A context X is *feasible* if there exists $X' \subseteq X$ for which the following holds: (X, X') is compatible with $\forall u_1, \dots, u_k \phi$ and 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 only the context-independent vocabulary, that $\{x_1, \dots, x_q\}$ is feasible. Hence, we can formulate in **TCL** the 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 outline for the general case. \square

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 300’ to 6000’; contexts of less than 28 distances are feasible. On the other hand feasible contexts can contain thousands of distances, provided that they have large enough gaps.

3.4.2 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 initially, that is, in context $\{a\}$, as a P (or as a non- P), then, the classification should endure except for certain large dense contexts. It takes a specially designed predicate, like ‘winning finalist’ in the example of Sect. 3.4, 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 classificatio 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. Consider ‘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 classificatio 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:

$$\begin{aligned}
 [x_1, x_2, \dots, x_n] P(x_1) \wedge y < x_1 &\rightarrow \bigwedge_{i=1}^n \{[y, x_1, x_2, \dots, x_n](P(x_i))\} \\
 &\leftrightarrow [x_1, x_2, \dots, x_n]P(x_i)
 \end{aligned}$$

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

3.5 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

$$\forall x[x] \{P(x) \rightarrow \exists y, z (P(y) \wedge \neg P(z))\},$$

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: $\forall x[x] (P(x) \rightarrow \exists y, z (P(y) \wedge [y] \neg P(z)))$, 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 , even when $f(X)$ has no effect on α 's truth-value; for example, the truth-value of a logical truth containing P is independent of the interpretation, but the sentence refers to some contexts; $\forall u(P(u) \rightarrow P(u))$ refers to all one-element contexts. The set of contexts referred to by a sentence is determined, syntactically, as follows. Correlate with every wff α a set, $cont(\alpha)$, of context lists (sequences of terms) associated with α , 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, $C \in \text{cont}(\alpha)$, 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 only to feasible contexts. Call a wff feasible if its universal generalization is a feasible sentence. Call a proof in **TCL** feasible if it consists of feasible wffs. In ordinary usage, tolerant predicates are meant to be used only in feasible sentences. It is obviously desirable that, restricted to the feasible portion of the language, the deductive part of the system should capture the semantics. And this indeed is the case:

The Completeness 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 considered semantic constraints on the tolerant predicate (or predicates, in the general case) that are expressible as a universal first order formula (1). Now the scheme (1*), which enforces the satisfaction of (1) in all contexts, has non-feasible instances—those in which the number n is too large. Also, the tolerance scheme (TOL) covers unfeasible instances. 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 cases of tolerant predicates 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 **TCL** yields, as a corollary, the completeness feasibility theorem, and plays a part in proving the strengthened version just mentioned. Say that two context lists are *equivalent* if the sets of terms occurring in them are the same. Call an *instantiation* of a context list any context list obtained by substituting variables by terms.

Localized completeness If $\Gamma \models \beta$, 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 .

3.5.1 Possible developments of TCL

A considerable enhancement of expressive power is obtained by adding to **TCL** quantifiable variables ranging over contexts, say $X, Y, X_1, \dots, Y_1, \dots$, which can occur in the context operator; along with them we add a membership symbol that can occur in wffs in the form: $x \in X$. 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 $Fsble_q(x_1, x_2, \dots, x_q)$ —each saying, for a particular q , that $\{x_1, x_2, \dots, x_q\}$ is feasible—we can say that a context is feasible by a single wff $Fsble(X)$. Schemes such as (1*), obtained by varying the number of variables in operators of the form

$[y_1, y_2, \dots, y_n]$, can be now replaced by single wffs containing operators of the form $[X]$.

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). Completeness is regained if, disallowing quantification over context variables, we treat them schematically; that is, the derivation rules allow to substitute X , in any theorem, by an arbitrary context list. Essentially, this version amounts to a convenient notational variant of **TCL**.

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.

4 Part II: borderlines and higher order vagueness

4.1 Semantic modalities, borderlines, higher order vagueness and degrees

I shall base my analysis on a definiteness operator, Δ , such that, for a vague predicate $P()$, $\Delta P(a)$ reads: “definitely $P(a)$ ”. It means that a definitely falls under P , that is, it falls under P and is not a borderline case.²⁴ Such an operator has been introduced by Fine in his classic account of vagueness in terms of supervaluations.²⁵ Borderline cases can be then characterized as those that neither fall definitely under P , nor definitely under $\neg P$.

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

In the supervaluation framework Δ can be interpreted in the obvious way. Outside this framework, it has been a disputed subject. It has been argued that the prefixin of ‘definitely’ does not change the content of a statement, or that being definitely true is just the same as being true.²⁶ Since in natural language ‘definitely’ is vague and

²⁴ The analysis leads to a definition of degrees in terms of iterated modalities. A linguistically oriented analysis of degrees is given in Klein (1980). That work outlines an extremely rich machinery for a semantic and syntactic analysis of English comparative adjectives and degree modifiers as exemplified in sentences such as ‘John is more happy than Mary is sad’, or ‘Sean is taller than Mary but not very tall’. It incorporates variants of generative grammars, elements of Motague’s intensional semantics, an abstraction operator and a notation for Kaplan’s *character*. The contexts are, roughly speaking, *comparison classes* (e.g., ‘tall’ is interpreted as ‘tall as a human’ in the context of humans, and as ‘tall as a mountain’—in the context of mountains.) Within each class the semantics is based on a partial truth-function, with an option for supervaluations. Sequences of shifts in comparison classes produce degrees. A common basic idea concerning degrees is shared by the present proposal. The special feature of my proposal is the use of iterated modalities as the only tool for producing degrees, where this is done within the modal system KTB (or possibly some offshoot of it).

²⁵ Fine (1975, p. 40).

²⁶ Williamson (1994, p. 194) argues that this is so, unless ‘definitely’ is interpreted epistemically.

ambiguous and can have many roles, there is ample room for arguments and counterarguments. The appeal to the truth predicate does not make things clearer. For one thing, the difference should not be construed as a difference between the truth of $P(a)$ and the definite truth of $P(a)$, but as the difference between the truth of $P(a)$ and the truth of $\Delta P(a)$. Invoking Tarski's biconditional is of no help either.²⁷

Conceivably, one can use 'definitely' merely to emphasize, or to express confidence. But 'definitely' adds true information, if the test for 'definitely P ' is 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 still white but not definitely white." Moreover, one's hesitation may be due to semantic indeterminacy, therefore an expression of confidence can mean that this case is *not* one of semantic indeterminacy, i.e., is not a borderline case. And this *does* add informative content. In any case, we are not concerned here with analyzing the natural-language meaning of 'definitely', but give it a narrow technical sense as an operator in our system. One way of doing this is to construe it in terms of borderlines: definitely $P(a)$, just when (i) $P(a)$, and (ii) a is *not* a borderline case of P . And I propose to characterize a as a borderline case of P just when the semantics does not provide grounds for ruling out either ' $P(a)$ ', or ' $\neg P(a)$ '. To spell it out, consider a competent judge who classifies a as a P , yet recognizes the legitimacy of classifying a as a non- P . Or imagine an idealized community of competent judges; some (to whom I belong) classify a as P , but others (whom I acknowledge) classify it as non- P . Moreover, the difference of opinions is due solely to lack of semantic determination, not to different factual knowledge. It is this possibility that 'definitely' is supposed to rule out. Wright (1994) took a similar line (if I understand him correctly), but seems to have misconceived an important aspect.²⁸

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, he or she 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 to be addressed shortly. Of course, in actual cases the speaker may also be at a loss; he or she may hesitate due to a combination of factors. I am not concerned here with a

²⁷ 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. 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'. We would do better to separate the treatment of 'true' from that of 'definitely'.

²⁸ He proposes (p. 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 impugning in any way the vagueness of P . Sorensen's reply to Wright, in the same volume, has the same misconception.

psychological picture of someone’s “state of mind”, but with the objective aspect of semantic indeterminacy.

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). 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 this interpretation of ‘definitely’ is 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*. This corresponds to one reading of the inconsistent notation in Evans (1978)

My notation, which differs from that of various authors, is motivated by the need for a convenient natural dual of the necessity operator.

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:

$$(M) \quad \nabla\Phi \Leftrightarrow \neg\Delta\neg\Phi \quad \mathbb{B}\Phi \Leftrightarrow \nabla\Phi \wedge \nabla\neg\Phi \quad \Delta\Phi \Leftrightarrow \Phi \wedge \neg\mathbb{B}\Phi$$

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”.²⁹ 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 pile 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 three-fold classification. Will 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 considered the alternative as clearly 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-fold classification may give her pause also in the three-fold 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, “pass”. Nothing mandates that hesitations in the two-fold situation be always clear-cut “borderlines” in the three-fold one. ²⁹ I owe Achille Varzi the idea of using a grading situation in order to analyze vagueness.

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 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 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. But I find that hardly any 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 \(1975, p. 146\)](#).³⁰ 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

³⁰ The straightforward modal corollary of supervaluations is an *S5* 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, \dots, 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^{i+1}, \dots , is accessible from another, a^0, a^{i+1}, \dots , iff $b^i \in a^{i+1}$ for all $i < \infty$. 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 .

is to encompass, within one modal language, all the generated categories, expressing them in terms of iterated modalities.

Just as ‘definitely’ can narrow a concept by ruling out borderline cases, ‘definitely definitely’ can effect 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 define 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.

In terms of the possible world semantics, $\nabla\Phi$ is true in world w just when Φ is true in some world accessible from w . In terms of semantic modality, the truth of $\nabla\Phi$ 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 possible sharpenings. 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 $\nabla\Phi$ true, and Mary holds $\nabla\nabla\Phi$ true but not necessarily $\nabla\Phi$. 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 that it includes the views that achieve certain effects. Later we shall see some examples.

It is obvious and well known that $S5$ 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 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 Sect. 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, ∇P , $\mathbb{B}P$, and $\Delta\neg P$ have, respectively, the extensions $\{0,1\}$, $\{2,3\}$, and $\{4, \dots, 8\}$, and each is sharp (i.e., has an empty borderline). The outcome is plausible. Vagueness disappears, if instead of using the predicate P , and its negation, $\neg P$, we use ΔP , $\mathbb{B}P$, and $\Delta\neg P$. But we cannot translate statements involving ‘ P ’ into this language. The vagueness of P , which consists in the existence of semantically indeterminate cases in the two-fold classification is not diminished

by the fact that $\mathbb{B}P$ is sharp (see also footnote 28). There are others plausible ways of modeling of ‘a small fraction of 8’, in which, we shall see, the predicate has vagueness of infinit order; there are also, for each $n \geq 1$, models 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 \mathbb{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 infinit order in which all higher-order borderline regions of order ≥ 2 are the same. (A precise definitio of the order of vagueness will be given in Sect. 4.2.)

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.³¹ 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. But the real argument for KTB , derives from the behavior of iterated modalities.

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


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

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 any $k > 1$, this rough scale can be refined by subdividing $\Delta \psi$ according to additional markings: $\Delta^k \psi, \Delta^{k-1} \psi, \dots, \Delta^2 \psi$, 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*:


³¹ The use of KTB was also suggested in Williamson (1999), where iterated modalities are used in the analysis of higher-order vagueness. My paper is independent of his work.

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


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 $\leq k$ nesting of modal operators.



Assuming *KTB*, we can partition each degree of a standard scale as follows. Add, inside the central $\mathbb{B}\psi$ region, two scale marks, for $\nabla \Delta \psi$ and for $\nabla \Delta \neg \psi$; 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, in the modal system, of certain sentences. The provability of $\nabla \Delta \psi \rightarrow \psi$, and of $\Delta \nabla \neg \psi \rightarrow \neg \psi$, 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 refinement are possible, but they require specific theorems in the modal theory (that is based on the vocabulary of ψ) which, it appears, are not instances of modal axioms of nicely behaved modal logics. In general, we can get refinement 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. 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.

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 define by:

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

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 $M \wedge M'$, define by:

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

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, \mathbf{D} , of *potential degrees*—consisting of generalized modalities that may serve as degrees—can be define as follows: Let \mathbf{D}^+ be the smallest class containing \mathbb{E} , such that if $D, D' \in \mathbf{D}^+$, then, $\Delta D, \nabla D, D \wedge D', D \vee D' \in \mathbf{D}^+$. Let \mathbf{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 \mathbf{D} consists of all members of \mathbf{D}^+ , all members of \mathbf{D}^- , and all conjunctions $D^+ \wedge D^-$, where $D^+ \in \mathbf{D}^+$ and $D^- \in \mathbf{D}^-$. For a given $D \in \mathbf{D}$, 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 we grade exams on a scale of 1–10, 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 sufficient control over the borderlines—at least for standard and one-refinement scales. Here is the relevant theorem:

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

- (i) For a standard scale, $\phi_i \rightarrow \neg \mathbb{B}\phi_j$ is provable, for all $j < i - 1$, or $j > i + 1$.
- (ii) For a one-refinement scale, $\phi_i \rightarrow \neg \mathbb{B}\phi_j$ 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 $\phi_i \rightarrow \neg \mathbb{B}\phi_j$ 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 all degrees that are removed from it by 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*.

Let us now apply degrees to predicates, instead of sentences. The application is straightforward. Given a predicate P and a generalized modality M , MP is the predicate defined by:

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

The *degree predicates* for P are the predicates DP , where D ranges over some scale of degrees. For some fixed scale, the extensions of the degree predicates form a partition of the domain, over which the argument under the predicate ranges. The diagrams above become illustrations of partitions of the universe, according to the “ P -degree” of the objects.

Cases like ‘tall’ or ‘walking distance’, which come with a given ordering (or pre-ordering) of the domain, yield insights into degree predicates.³² Very often, the predicate P is either monotone or anti-monotone with respect to \leq . Recall that a predicate S (over the given domain) is *convex*, if it satisfies

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

It should be obvious that, for a monotone or anti-monotone P , convexity is a necessary condition for being a degree predicate (if two people have, on our scale, the same degree of tallness, then every person whose height is between their heights has also that degree). It can be shown that if D is a potential degree, then for P that is either monotone or anti-monotone, DP is convex. Note that, if P is monotone or anti-monotone, \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, which should be elaborated 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

³² For the definition of a pre-ordering, see footnote 20.

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.

4.2 More on higher order vagueness

The *order of vagueness* is defined as follows. A sentence ψ has *vagueness of order* $\geq k$ 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 $\geq k$ is: For some $n \geq k$, $\mathbb{B}^n P$ has non-empty extension. Exact orders are derived from this in the obvious way. A sentence, or predicate, has vagueness of infinite order if it has vagueness of order $\geq k$, for all finite k .

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*. We can also introduce degrees as an additional component, through a probability measure on sets of 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.³³

4.2.1 Sharp models for higher order vagueness

I shall now discuss the possibilities of sharp modeling of higher order vagueness. As an illustration consider again the case where ' $P(x)$ ' reads: ' x is a small fraction of 8', with ' x ' ranging over $0, \dots, 8$. 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 other 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: $|\mathbb{B}^n P| = \{2, 3\}$, for all $n > 0$. Each of these patterns, and the pattern determined by 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 without hesitation, but hesitate whether 2 should belong

³³ 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.

there as well; I decide it does not, but acknowledge as legitimate a view on which it does: $|\mathbb{B}^2 P| = \{2\}$. When the options are further extended by adding $\mathbb{B}^2 P$, 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 story is sufficient to show how this works.

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:

$$\begin{aligned}
 &|P| = \{0, 1, 2\} \quad |\neg P| = \{3, 4, \dots, 8\} \quad |\Delta P| = \{0, 1\} \quad |\Delta\neg P| = \{4, \dots, 8\} \\
 &|\Delta^n P| = \{0\}, \text{ for all } n > 1, \quad |\Delta^n\neg P| = \{4, \dots, 8\}, \text{ for all } n > 0, \quad |\mathbb{B}P| = \{2, 3\} \\
 &|\nabla\Delta P| = \{0, 1, 2\} \quad |\nabla\Delta\neg P| = \{3, 4, \dots, 8\} \quad |\Delta\mathbb{B}P| = \emptyset \quad |\mathbb{B}^2 P| = \{1, 2, 3\} \\
 &|\nabla\Delta^2 P| = \{0, 1\} \quad |\nabla\Delta^2\neg P| = \{3, 4, \dots, 8\} \quad |\Delta\mathbb{B}^2 P| = \{2\} \quad |\mathbb{B}^3 P| = \{1, 3\} \\
 &|\nabla\Delta^3 P| = \{0\} \quad |\nabla\Delta^3\neg P| = \{4, \dots, 8\} \quad |\Delta\mathbb{B}^3 P| = \{1, 3\} \quad |\mathbb{B}^n P| = \emptyset, \\
 &\text{for all } n > 3.
 \end{aligned}$$

My ex-student Jonathan Simon showed that, in the case of a monadic vague P define 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, \dots, w_{2^n} , with a reflexive symmetric accessibility relation that links them in a 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 another, different world. So much for the examples.

There is a philosophical perception that a faithful rendering of vague terms must be given in a vague language.³⁴ 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 quite large. The mismatch between the model and the linguistic phenomena can be perceived as 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

³⁴ 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.

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 qualify 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 one 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 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 vagueness in the metalanguage.

The observations about sharp modeling of vagueness apply also to tolerance. Our usage determines only basic features of the corresponding formal systems, leaving fine 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.³⁵ 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 sharp relation that determines the amount 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 introduce 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 highly complicated entities. 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 an instance of the slack phenomenon.

4.3 TCLV: Combining tolerance with general vagueness

A combined modeling of tolerance and general vagueness requires that we extend the language of **TCL**, by adding a modal definiteness operator, Δ , according to the outline just sketched. Call the resulting system **TCLV**. 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 possible worlds be $(\mathcal{M}, f_i)_{i \in I}$. Assume also that P is the only tolerant predicate; say it stands for 'walking distance'.

The universal sentence that determines P 's semantics, and the tolerance condition (TOL), are true in every (\mathcal{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. For example, the maximal distance that in all contexts is a walking distance can vary from world to world. As can the minimal distance that in all contexts is not walking distance. 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 $P(c)$ (i.e., $P(c)$ in the context $\{c\}$) can vary. Higher-order vagueness is modeled in the way described in Sects. 4.1 and 4.2.

References

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

Cornelius, H. (1897). *Psychology als Erfahrungswissenschaft*. Leipzig: Teubner, xv+445 pp.

³⁵ Sorensen's example is somewhat different, but the difference does not matter for the present point.

- Dummett, M. (1975). Wang's paradox. *Synthese*, 30, 301–324. (Reprinted in *Vagueness Reader*).
- Evans, G. (1978). Can there be vague objects? *Analysis*, 38, 208. (Reprinted in *Vagueness Reader*).
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30, 265–300. (Reprinted in *Vagueness Reader*).
- Gaifman, H. (1992). Pointers to truth. *The Journal of Philosophy*, 89, 223–262.
- Gaifman, H. (2000). Pointers to propositions. In A. Chapuis & A. Gupta (Eds.), *Circularity, definition, and truth* (pp. 79–122). New Delhi: Indian Council of Philosophical Research.
- Gaifman, H. (September, 2001). Contextual logic and its applications to vagueness. *The Bulletin of Symbolic Logic*, 7(3), 241 (Abstracts of Invited Talks of the Annual 2001 Meeting of the ASL).
- Gaifman, H. (2008). Contextual logic with modalities for time and space. *Review of Symbolic Logic*, 1(4), 433–458.
- Hart, H. L. A. (1961). *The concept of law*. London: Oxford University Press.
- Kamp, H. (1981). The paradox of the heap. In U. Mönnich (Ed.), *Aspects of philosophical logic*. Dordrecht: Reidel.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. New York: Oxford University Press.
- Klein, E. (1980). The semantics of positive and comparative adjectives. *Linguistics and Philosophy*, 4, 1–45.
- Kofka, K. (1922). Perception: An introduction to the Gestalt-Theorie. *Psychological Bulletin*, 19, 531–585 (Reprinted in *Classics in Psychology*, ed. Thorne Shiple, 1961, Philosophical Library).
- McGee, V., & McLaughlin, B. (1994). Distinctions without a difference. *The Southern Journal of Philosophy*, 33(Supplement), 203–251.
- Putnam, H. (1983). Vagueness and alternative logic. *Erkenntnis*, 19, 297–314.
- Raffman, D. (1994). Vagueness without paradox. *The Philosophical Review*, 103, 41–74
- Russell, B. (1923). Vagueness. *Australasian Journal of Philosophy and Psychology*, 1, 84–92. (Reprinted in *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 G. E. Macdonald (Ed.), *Perception and identity* (pp. 235–251). London: MacMillan.
- Unger, P. (1979b). There are no ordinary things. *Synthese*, 41, 117–154.
- Williamson, T. (1992). Vagueness and ignorance. In *Proceedings of the Aristotelian Society*, 66(Suppl), 145–162 (Reprinted in *Vagueness Reader*).
- Williamson, T. (1994). *Vagueness*. London: Routledge.
- Williamson, T. (1999). On the Structure of Higher-Ord Vagueness. *Mind*, 103, 127–143
- Wright, C. (1975). On the coherence of vague predicates. *Synthese*, 30, 325–365. A shorter version appeared as "Language-mastery and the sorites paradox". In G. Evans & J. McDowell (Eds.), *Truth and meaning*. Oxford: Oxford University Press, 1976. (Reprinted in *Vagueness Reader*).
- Wright, C. (1994). Epistemic conception of vagueness. *The Southern Journal of Philosophy*, 33(Supplement), 133–159.
- Zeeman, E. C. (1961). The topology of the brain and visual perception. In *Topology of 3-manifolds and related topics*, Proceedings of The University of Georgia Institute, Prentice Hall, pp. 240–256.