



OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

The British Society for the Philosophy of Science

Review: Vagueness, Ignorance, and Margin for Error

Reviewed Work(s): Vagueness by Timothy Williamson

Review by: R. M. Sainsbury

Source: *The British Journal for the Philosophy of Science*, Vol. 46, No. 4 (Dec., 1995), pp. 589-601

Published by: Oxford University Press on behalf of The British Society for the Philosophy of Science

Stable URL: <http://www.jstor.org/stable/687900>

Accessed: 15-08-2017 22:02 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://about.jstor.org/terms>



JSTOR

Oxford University Press, The British Society for the Philosophy of Science are collaborating with JSTOR to digitize, preserve and extend access to *The British Journal for the Philosophy of Science*

REVIEW ARTICLE

Vagueness, Ignorance, and Margin for Error¹

R.M. Sainsbury

This is a beautiful book. It traces the history of philosophical work on vagueness, evaluates the main contemporary treatments of the subject, and advances the view that vagueness is ignorance, supporting this position by a detailed account of the notion of inexactness in knowledge.

According to Williamson, the history of the subject is jerky: an early start, a pause, and then two fertile periods, separated from one another by a millennium and a half. The story begins in classical Greece with the discovery of the paradox of the Heap, attributed to Eubulides, a contemporary of Aristotle: 'If you admit that one grain does not make a heap, and are unwilling to make a fuss about the addition of any single grain, you are eventually forced to admit that ten thousand grains do not make a heap' (p. 8). The Stoics seem to have been the first to give systematic attention to the paradox and its relation to vagueness. Williamson endorses, and later develops, a view he finds in Chrysippus, that vague terms have sharp cut-offs, though we cannot tell where they fall. The second period in which interest in vagueness has flourished is the last hundred or so years. Williamson tentatively suggests that the explanation of this jerky development lies in the centrality of formal logic in both the fertile periods: logic makes it hard to fudge the paradoxical implications of vagueness.

Williamson traces the recent history from Frege, through Peirce, Russell, Wittgenstein, Black, and Hempel. As in his account of the Stoics, Williamson's exegesis constantly illuminates the philosophical issues. By Chapter 4, the history is complete, and Williamson turns to his presentation and evaluation of the main current positions: many-valued theories, supervaluations, and nihilism (the view that vague expressions are empty or meaningless). The remainder of the book (chapters 7 to 9) is devoted to the elaboration and defence of the epistemic theory: vague expressions draw sharp boundaries; their vagueness is their

¹ Review of Timothy Williamson [1994]: *Vagueness*, cloth £35.00, London and New York, Routledge.

connection with a particular kind of ignorance. Objections are considered, and consequences explored, including consequences for whether or not the world is vague.

Williamson's criticisms of theories other than the epistemic theory are very searching; many exploit the requirement that theories do justice to higher-order vagueness. He goes to great pains to explore how the positions he attacks could be developed to counter his criticisms. The reader is aware, however, that trouble is brewing: the more convincing the criticisms of the other theories, the harder it will be to avoid the epistemic theory.

The theory may be strange (but then 'the truth about vagueness must be strange', p. 166); however, it is not simply a last resort. For one thing, there is a general argument for bivalence, based on Tarski's schema, and bivalence arguably entails sharp boundaries (cf. Williamson [1992b], & [1994], pp. 185–98). Sharp boundaries do not entail Williamson's version of the epistemic theory; for example, they are, in a different way, accepted by theorists of semantic indecision. However, it is certain that if there are sharp boundaries, we do not know where they fall, and this goes a long way towards Williamson's position. Other theorists must either dispute the arguments for bivalence (for example, they might see no merit in retaining Tarski's schema), or dispute the connection between bivalence and sharp boundaries, or that between sharp boundaries and the epistemic theory. A semantic theorist might suggest that it is not failure of truth or falsehood that characterizes vagueness, but failure of *definite* truth or *definite* falsehood. This theorist is then faced with the task of giving a content to 'definite' which distinguishes it from 'knowable', a task which may be subsumed under the disconcerting general question: 'what more could it take for an utterance to be definitely true than just for it to be true?' (p. 194).

If Williamson is right, the epistemic view states just how one would expect things to be, given some under-appreciated implications of the nature of knowledge (p. 186, 215 and Chapter 8). This, I believe, is the most interesting and original part of this excellent book, and so will be almost the sole topic of this review. The main idea is to use margin for error principles to explain the ignorance postulated by the epistemic theory. It was presented as early as Williamson's *Identity and Discrimination* ([1990]; see also Williamson [1992a, 1992b]); but in *Vagueness* it is fully elaborated, and connections with related issues are explored in detail.

Vagueness allows borderline cases, cases in which we cannot answer questions involving the vague expression or concept. You may know how tall someone is to the millimetre, yet be unable to say whether or not he is tall. You may see a shade under perfect conditions for assessing its colour, yet be unable to say whether or not it is red. Many theorists take it that the

inability to decide springs from there being no fact of the matter, nothing to make a decision correct or incorrect. The inability is not a manifestation of ignorance, for in borderline cases there is nothing to know or to fail to know. Epistemic theorists, by contrast, deny this: they say that there is a fact of the matter, there is something to know, but we cannot know it. The distinctive kind of irremediable ignorance is, for these theorists, the essence of vagueness.

What is this kind of ignorance, and why can it not be remedied? To appreciate Williamson's answer, we must first consider his notion of inexact knowledge. Take a case in which there is a precise fact, but our knowledge of it is not exact. Williamson's example is a crowd in a stadium. For some number, n , there are exactly n people in the stadium (at a given moment), but I cannot tell, just by looking, exactly what this number is. I have some numerical knowledge: I may know that there are more than 500 people and fewer than 200,000; but it is not exact. Even if I happened to have the true belief that there are exactly n people in the stadium, this belief is not reliable enough to count as knowledge: I can't tell the difference between the actual situation and one in which there are exactly $n + 1$ or $n - 1$ people; a person more or less would make no difference to how things would seem, but would make my belief false. This ensures that knowledge in such cases is subject to a 'margin for error principle'. An example is: if I know that there are n people in the stadium, then there are n people in all similar cases. 'The degree and kind of the required similarity depend on the circumstances' (p. 226). In the stadium example, cases with a person more or less than n are certainly similar, for I could not distinguish them, and this precludes my having exact knowledge of how many people are present.

The margin for error principle just mentioned has the form: 'A' is true in all cases similar to cases in which 'it is known that A' is true (p. 227). The example can also suggest related principles of a different form, for example: if I know that there are not exactly i people in the stadium, then there are not exactly $i - 1$. If I know that there are other than exactly i people, then I am reliably right, that is, am right in all similar cases, so in particular I would be right even if there were other than exactly $i - 1$ people, since this case is similar to the actual case in which there are other than exactly i people.

Appeal to margin for error principles can explain why knowledge of the boundaries of vague concepts is impossible. Consider a soritical series of colour tiles, red at the left and slowly shifting to orange towards the right. Suppose the following margin for error principle is true: if I know that the n th tile is red, then the $n + 1$ th tile is also red. Then it would be impossible for me to know of the last red tile that it is red. It would likewise be

impossible for me to know anything of the form: the i th tile is red and the $i + 1$ th tile is not. It would remain, however, to justify the principle.

Alternatively, suppose the following margin for error principle were true: if I know the i th tile is red, then the i th tile is red in all similar cases. What are similar cases? Ones in which the tile is a little different in colour? Or ones in which a different but similar tile is the i th? Or ones in which the tile which was the source of that belief was other than the actual i th? Whichever answer we give, we can explain the ignorance which the epistemic theory postulates. Suppose the tile which is actually i th is close to the boundary. If it had been a little less red it would not have been red, so my actual belief, though in fact true, would have been false. Or, if the boundary had come a little earlier, that is, if the word or concept *red* had been different, so as to rule the i th tile non-red, then my belief would have had a content fixed by this different ruling, and it would have been false. Or, if a similar but non-red tile had given rise to a belief, it again would have differed in content from my actual belief and would have been false. It would remain, however, to justify the principle, and the selection of the relevant dimensions of similarity.

Williamson suggests that suitable margin for error principles can be justified in cases in vagueness without assuming the truth of the epistemic theory. Margins for error arise through imperfect discriminatory powers. In the stadium case, for some numbers i , I cannot, merely by looking, discriminate situations in which i people are present from situations in which $i + 1$ people are present. This prevents me from having exact knowledge, because it prevents true belief retaining truth in situations which are similar, as far as I can tell. One might expect that the relevant analogue for vagueness would dwell on such facts as that I cannot discriminate adjacent shades in a sorites series, or cannot discriminate collections of grains differing by just one. The approach need not be piecemeal for, arguably, a relevant failure of discrimination is present whenever there is a possibility of a sorites argument; and this possibility might in turn be held to be co-extensive with vagueness, or vagueness of the relevant kind.

However, that is not Williamson's approach. Rather, he suggests that 'what distinguishes vagueness as a source of inexactness is that the margin for error principles to which it gives rise advert to small differences in meaning, not to small differences in the objects under discussion' (p. 230–1). The claim that such margin for error principles apply does not depend upon the epistemic theory, and may be grounded as follows. Someone who asserts ' n grains make a heap' might well have used the same words assertively, even if our linguistic practices had been slightly different, conferring upon this sentence the semantic status actually possessed by ' $n - 1$ grains make a heap'. 'The actual assertion is the outcome of a

disposition to be reliably right only if the counterfactual assertion would have been right' (p. 232). So if one knows that n grains make a heap, then $n - 1$ grains make a heap.

Why does Williamson select this explanation for margin for error principles in the case of vagueness, when others appear to be available, like principles relating to the limitations of perceptual discrimination? (For example, one might try to explain the principle at the end of the previous paragraph on the basis of our inability to discriminate between n -grained and $n - 1$ -grained heaps.) Williamson's epistemic theory is designed to be consistent with the claim that vague facts supervene on precise ones (p. 204).² If such supervenience conditions were knowable, then vagueness would not give rise to any inescapable ignorance. So the theory must allow for, and explain, unknowable truths of the form: n grains are the fewest that make a heap. Limitations of perceptual discrimination do not appear promising in this connection; limitations on our sensitivity to small shifts in usage are more promising.

Williamson considers two objections to his use of margin for error principles: that they lead to sorites paradoxes, and that his account of their nature in the case of vagueness is inconsistent with our knowing the meanings of vague words.

If one knows that n grains make a heap, then $n - 1$ grains make a heap. We do know that 10,000 grains make a heap. So we can infer that 9,999 grains make a heap. We know what we correctly infer from what we know. So we know that 9,999 grains make a heap; we seem set on a paradoxical course. However, each stage of the argument requires that we know the premises at that stage. At each stage, one of our premises has the form: we know that n grains make a heap; so at each stage we need to know something of that form. Williamson argues that this knowledge is not in general available. It would be available were the KK principle correct (if we know that A, then we know that we know that A); but previous arguments have made plausible, quite independently of vagueness, that the principle is not correct (see 8.2). Without it, there is no reason to think we have what is needed to generate a paradox.

Williamson suggests that inexact knowledge can give rise to failures of the KK principle, independently of vagueness. I do not know that there are not 20,000 people in the stadium. Assume that this is not a vague matter (or that any vagueness in the only potential source, knowledge, has been

² Williamson's official stand is as I say: he does not explicitly endorse the supervenience thesis. However, as Michael Martin pointed out to me, it looks as if the way in which Williamson describes the relevant margin for error mechanisms presupposes this supervenience: worlds differing in a vague respect must differ in a precise one.

refined away).³ Then we have conditions for applying the least number principle (that every non-empty set of natural numbers has a least member) to generate the conclusion that there is a least number, n , such that

(2) I do not know that there are not exactly n people (p. 218).

By definition of n ,

(1) I know that there are not exactly $n - 1$ people (p. 218).

Reflecting on my discriminatory limitations, for all relevant numbers, including n :

(3) I know that if there are exactly n people, then I do not know that there are not exactly $n - 1$ people (p. 219).

It looks as if I can apply (1) to what I know in (3) to infer that there are not exactly n people. If I know the premises of this argument, and can attain knowledge by reasoning correctly from what I know, then I know that there are not exactly n people, which contradicts (2). One reasonable conclusion is that I do not know all the premises. Since by (3) I do know the conditional, one has reason to say that the premise I do not know is (1),⁴ and this is to say that the KK principle is false.

The other objection to be considered is that Williamson cannot consistently hold that we know the meanings of vague terms. Suppose that this knowledge is of something like: u says that P (where u is an utterance containing some vague expression). If I know this, then I must be reliably right; so I must be right for cases which, from the perspective of my discriminations, are similar; however, a case in which u meant something a little different from P is relevantly similar (I could not tell if I were to be placed in such a language-using group), and yet would be one in which I would be wrong about the meaning of u . So I do not know what u means.

Williamson recognizes that this is a 'tempting conclusion' (p. 235) but

³ It is not obvious that Williamson's confidence in the possibility of this refinement is consistent with his scepticism about the possibility of attaining complete precision, e.g. at pp. 170, 192, 248. True, the refinement requires only the elimination of relevant vagueness, not of all vagueness (p. 222). Here a distinction is important: some refinements may themselves involve vagueness (we might refine 'tall' in terms of heights, but heights themselves are vague); but if a refinement can fully eliminate vagueness in a certain respect, we are entitled to wonder why further refinements cannot eliminate vagueness in the other respects. The respects cannot be literally infinite.

⁴ It is tempting to think that if one properly stipulates whether or not p one ends up knowing whether or not p . Williamson envisages that borderline cases for 'know' are resolved by stipulation (see previous note); so one knows of these cases whether they are cases of knowledge or not. If they are cases of knowledge, one knows it. This appears to suggest an independent route to knowledge that (1). Williamson may wish to meet this objection by giving the stipulations in terms of general rules (p. 222, middle), which do not yield knowledge of singular statements like (1).

seeks to persuade us to resist it. He says that the epistemic theorist may concede that, for some P^* other than P , one may not know that u does not say that P^* , but he insists that it does not follow that one does not know that u says that P . An analogy with discriminating faces is advanced. We often know who people are by their face. There could have been a look-alike, but this mere possibility does not defeat our actual identificatory knowledge. ‘Similarly, why should our ability to recognize the meaning of utterances in our language be undermined by the mere possibility of indiscriminably different meanings?’ (p. 236). If we answer this question as Williamson hopes (saying ‘there’s no reason’), how can we retain the claim that the mere possibility of indiscriminably different meanings defeats putative knowledge in cases of vagueness? We are asked to believe that the mere possibility of indiscriminably different meanings may defeat knowledge that n grains make a heap, but does not defeat knowledge of what ‘ n grains make a heap’ means.

To clarify the question, we need to bring the underlying conception of reliability into the open. Given the supervenience of the vague on the precise, what is actually said by ‘ n grains are the fewest which make a heap’ is true at all worlds or none. Hence the reliability of a belief in such a proposition is not to be gauged in terms of whether it is true at all worlds similar to one in which it is believed. The very fact that reliability is as much an issue for necessary truths as for contingent ones shows that only a very preliminary conception of reliability would understand it in terms of truth at all similar worlds. Our proper target is the quality of the mechanisms involved in forming and sustaining a belief. We must try to discover whether an appropriate account of this quality would permit us to combine the following views: we know that ‘ n grains are the fewest to make a heap’ means, so this meaning is, presumably, reliably, available to other belief forming mechanisms; yet we do not know that n grains are the fewest to make a heap, on account of cross-world instabilities in the meanings used to express this belief.

Suppose I guess correctly that $193 + 245 = 438$. I do not know that $193 + 245 = 438$, because my belief does not issue from a reliable mechanism. This cannot be brought out by pointing to similar possible worlds at which $193 + 245$ sum to something other than 438, since there are no such worlds. Arguably, it can be brought out by pointing to possible worlds in which I exercise essentially the same capacities and yet arrive at a falsehood. A lucky guess is not a proposition which might easily not have been true, but a way of reaching a belief which might easily not have delivered a true one.

By analogy, suppose I guess correctly but luckily that n grains are the fewest that make a heap. The relevant mechanism presumably calls,

among other things, on my knowledge (as we are supposing) of the actual meaning or content of the word or concept *heap*. Williamson's theory requires us to identify a meaning-related unreliability in this mechanism. To do this, we must hold the mechanism, as it actually is, constant, and ask whether it would be truth-delivering in other situations. We can allow that the word 'heap' as actually used, or the actual concept *heap*, has its semantic features essentially. What is not essential is that the mechanism which constitutes my actual grasp of this content essentially constitutes grasp of that (semantically individuated) word or concept. There is a possible situation in which, thanks to some variation in the overall usage of the relevant linguistic community, the same mechanism, non-semantically individuated, would have constituted my grasping a word with a different meaning, or a different concept, one associated with a different cut-off point. In that situation, the mechanism which delivered the actual true judgement would have delivered a false judgement. For this story, we need to think of the relevant mechanism as individuated in part by the information on which it actually calls. It actually calls on some known information about the word or concept *heap*; but in an alternative situation, to call on that information would be to call on misinformation.

Arguably, this unreliability is consistent with my reliably identifying, and so knowing, the meaning of '*n* grains are the fewest that make a heap'. The belief which reflects this identification issues from a mechanism which aligns my linguistic beliefs with the overall usage of the relevant linguistic community, whatever that usage may be. In assessing its reliability by a modal feature, the mechanism is identified by its general capacity to deliver the right meaning, not by the actual meaning it delivers. If the overall usage of the relevant linguistic community had been different, my usage would also have been different, and my actual mechanism for aligning my linguistic beliefs appropriately with usage would have adjusted accordingly: it would have delivered a truth, albeit a different one from that which it actually delivered. The relevant mechanism could in principle track even small changes of meaning across words.⁵

On this approach, the counterfactual whose truth shows the unreliability of purported knowledge that *n* grains are the fewest which make a heap is this: if the mechanism which in fact delivered this judgement had been transplanted to a situation in which a 'heap'-affecting shift of usage had occurred, it would have delivered a similar-sounding, but in reality different, judgement, and it would have been false. The counterfactual

⁵ Williamson does not envisage a very finely tuned tracking: he allows that social determination of meaning may make same-sayers of speakers with different use-dispositions (p. 211). But if the consistency of the claims I am discussing at this point in the review can be made out for finely tuned tracking, it could be made out for coarsely tuned tracking.

whose truth shows the reliability of purported knowledge of the meaning of ‘*n* grains are the fewest which make a heap’ is this: if the mechanism which in fact delivered this judgement, a mechanism which serves in general to align judgements of meaning with ambient usage, had been transplanted to a situation in which a shift of usage affecting ‘heap’ had occurred, it would have delivered a different but still true judgement.

This attempt to defend Williamson is inadequate. If one cognitive mechanism consumes information produced by another, it cannot be right to test the reliability of the consuming mechanism by considering a situation in which the output of the producing mechanism is held constant, but other relevant matters varied, for example, matters relating to what information the producing mechanism should or would have produced in that situation. This would make many reliable mechanisms count as unreliable. In the following diagram, suppose that M_2 consumes information supplied by M_1 . We do not properly test the reliability of M_2 by asking whether, in a situation in which the input varies, and a different output is called for, but M_1 produces the information it actually produces, M_2 would produce the appropriately different output.

$$\text{input} \rightarrow M_1 \rightarrow M_2 \rightarrow \text{output}$$

M_1 is analogous to the mechanism which produces semantic information about some vague term, information which, along with other information, is consumed by M_2 , the mechanism immediately responsible for producing a judgement exploiting that term.

A similar criticism can be made by exploiting a familiar phenomenon. Suppose that I judge that I do not have arthritis in my tibia (the inflammation is similar in kind; but the tibia is not a joint). The claim that this is knowledge ought not, it would seem, to be defeated by the following counterfactual: if the mechanism which in fact delivered this judgement had been transplanted to a situation in which an ‘arthritis’-affecting shift of usage had occurred, so that ‘arthritis’ properly applies to any inflammation similar to arthritis-in-joints, it would have delivered a similar-sounding, but in reality different and false, judgement. In the transplanted situation, the judgement-delivering mechanism would call on information which, in that context, was false, namely that the ambient concept of *arthritis* can apply only to inflamed joints.

Perhaps the ‘arthritis’ change is large, so that the changed world is not similar to the actual one, whereas the ‘heap’ change is small, so that the changed world is similar to the actual one. We should not expect a precise measure of similarity; but we should expect more than a fiat. Or perhaps, surprisingly, arthritis knowledge is genuinely defeated by my less than wholly reliable sensitivity to meanings. More interestingly, the precise way

in which the reliability criteria are developed can be refined so as to avoid the objection. The best option is not specified in the book, though it would need to be in a complete statement of the position.

It may be possible to hold the combination of views I have been discussing: that the meaning of vague terms is unstable in a way which can defeat purported knowledge expressed by their means; yet that we know the meaning of these terms. I claim merely that vindicating the combination requires some further work on very delicate matters concerning evaluating reliability.

Epistemic theories of vagueness are not as such committed to the combination, and one might read some passages in Williamson as envisaging a variant epistemic theory. He says that knowledge of meaning 'is not possession of a complete set of metaphysically necessary truths but complete induction into a practice' (p. 211). Suppose we see this induction as essentially producing practical rather than propositional knowledge.⁶ These different kinds of knowledge might well involve reliability in quite different ways, so that the tension I identified could be eased. He would have to admit that I do not have propositional knowledge of the meanings of vague terms, and so do not know such things as that '*n* grains make a heap' says that *n* grains make a heap. The apparent implausibility is mitigated by the reflection that on any account I do know that every such disquotational sentence is true, and it is easy to mistake this for genuine knowledge of what such sentences express.

Williamson also mentions a yet further fall-back position for the epistemic theorist: to allow that speakers only roughly know what their utterances mean (p. 236). It would be hard to prevent this position collapsing into that of the semantic indecision theorist, who holds that our usage of vague terms is indeterminate between various precise meanings, no one of which can we select; so we only roughly know what our vague expressions mean. To what data could Williamson appeal, were he to allow that we only roughly know the meaning of our vague terms, in order to establish that a given linguistic practice was best described his way rather than as manifesting semantic indecision?

Among various other aspects a critic will wish to explore, I will mention three. First, as Williamson says, 'knowledge is not the only cognitive relation one cannot have to cut-off points for vague terms' (p. 244). Given that non-factive attitudes are not constrained by margin for error principles, how can the epistemic theorist explain the impossibility of reasonably believing a truth of the form '*n* grains are the fewest which

⁶ Williamson tells me that he intended to allow that knowledge of meaning by induction into a practice should be propositional, as opposed to merely practical.

make a heap'? Williamson suggests that 'the epistemic theorist can apply the account of inexact knowledge to such attitudes by working with a hypothesis: that one's evidence is simply what one knows' (p. 245). Even those who are willing to accept the hypothesis may doubt whether it does all the required work. Nothing is explicitly said to block the suggestion that an oracle of known reliability could tell one that n grains are the fewest that make a heap.

Second, in considering semantic theories, Williamson attends mostly to those which posit a sharp metalanguage. He does not consider Tye's theory, which has what many, including Williamson (e.g. p. 191), would see as the merit of starting out with an unequivocally vague metalanguage.

Third, it has seemed to many that, however things might be with our actual vocabulary, we could introduce expressions which by stipulation lack sharp boundaries. If expressions without sharp boundaries are so much as possible, it would be hard to resist the claim that some actually exist. Suppose that we are an examining body, and we have received numerical marks out of 112 for the candidates. Our task is to group each candidate into just one 'qualitative' category: excellent, very good, good, satisfactory, and so on. We have no knowledge of the examined subject (the marks are fed to us by others), but we do know the distribution into the categories in previous years in both this and other subjects. (So, for example, we know that two years ago in this subject 4% of candidates were categorized as excellent, 12% as very good or better, and so on.) Categorization supervenes on marks: difference of category entails difference of mark. Ranking by category mirrors ranking by mark in obvious ways. One candidate scores 103, 3 score 99, 2 score 97, 5 score 91; and I will spare you the remaining data. We decide to classify the first 6 candidates as excellent, the next 10 as very good; and I will spare you the remaining decisions. However, in order not to set a precedent, we refuse to answer counterfactual questions about what we would have decided about candidates with scores between 96 and 92, had there been any. We deliberately aim to leave the lower boundary of excellence unsharp.

There are various ways in which one might criticize this as a purported model for the introduction of a term which is boundaryless by stipulation. Williamson's discussion of 'dommal' is relevant. The term is introduced by making only the following stipulations: (a) all dogs are dommals; (b) all dommals are mammals. In earlier work ([1990], p. 107), Williamson allowed that this could count as a case of semantic indeterminacy. It would seem that there is genuinely a question as to whether a cat is a dommal; a point that would fit ill with the general arguments he has given for bivalence, notably in the book under review (Williamson [1994], pp. 185–98). He now denies that a cat is a borderline case for 'dommal', even if

users refuse with equal firmness to answer the question ‘Is a cat a dommal?’ ‘Yes’ or ‘No’. The argument is based on his claim that the concepts of truth and of falsity are asymmetrical (pp. 208–9). Whereas truth involves the satisfaction of a specific condition, falsehood does not: it is merely failure of truth. One who would use the story about excellence as a model for how a term might be stipulatively boundaryless is now confronted with this dilemma: either ‘a candidate with a score of 95 is excellent’ says nothing, or else it says something. In the former case we have not mirrored the phenomenon of vagueness (the argument for this is given on pp. 195–7); in the latter case we must count it as false, through having failed to meet the only available condition for truth, namely a score of 97 or more.

Can this conclusion survive the explicit claim by the stipulators that they are not specifying a minimum? Of course, not just anything can be made so by stipulation; for example, nothing with existential import, nothing incoherent. One seemingly common failure of stipulations is to cover all cases, as lawyers know. It would be surprising if an abstract argument can rule out this possibility; it would be even more surprising if it can be shown that stipulations which explicitly claim non-comprehensiveness are incoherent. The argument relies on the principle that falsehood is failure of truth (among utterances which say something). But if this adds anything to the argument for bivalence, it adds something which lacks support.

Nothing should henceforth be written on vagueness which fails to learn from this book. It should not be read only for its contribution to vagueness, but also for what it says about knowledge; for the purity of its style, which achieves a very high level of rigour and clarity without the use of symbols;⁷ and as an example of philosophy at its very best.

*Department of Philosophy
King's College London
Strand
London WC2R 2LS
UK*

Acknowledgements

Williamson's book was read and discussed by a group of London colleagues in which I was included, the other members being Dorothy Edgington, James Hopkins, Keith Hossack, Michael Martin, and Scott Sturgeon. I found these discussions very useful, and I apologize for any inadvertent plagiarism. I would also like to thank the following for valuable written comments on a draft version of this review: Keith Hossack, Michael Martin, Scott Sturgeon, and Timothy Williamson.

⁷ Williamson is quite free of a tendency he attributes to Max Black: ‘a tendency to mistake pomposity for rigour’ (p. 286).

References

- Tye, Michael [1994]: 'Sorites Paradoxes and the Semantics of Vagueness', in J. Tomberlin (ed.), *Philosophical Perspectives VIII: Logic and Language*, Northridge, CA, Ridgeview Publishing, pp. 189–206.
- Williamson, Timothy [1990]: *Identity and Discrimination*, Oxford, Basil Blackwell.
- Williamson, Timothy [1992a]: 'Inexact Knowledge', *Mind*, **101**, pp. 218–42.
- Williamson, Timothy [1992b]: 'Vagueness and Ignorance', *Proceedings of the Aristotelian Society*, Supplementary Volume **66**, pp. 145–62.
- Williamson, Timothy [1994]: *Vagueness*, London, Routledge.