

Lewis on Williamson: Evidence, Knowledge and Vagueness

Daniel Nolan, 27 November 2018, dnolan2@nd.edu

This paper does not quite have the form of a typical philosophy paper. Instead, it is an extended discussion of a letter that David Lewis sent to Timothy Williamson, dated May 21 1999. (Lewis 1999a) There is a lot of philosophical interest in this letter, so I will not be able to be thorough: I am going to focus on a few parts of the discussion, particularly those on vagueness. The point of my discussion is to provide a little context, and to highlight some of the interesting and puzzling parts of Lewis's discussion, especially as they bear on Lewis's own opinions about vagueness as of 1999.

What is most exciting for me in this letter is what Lewis says about his view of vagueness, and the view he defends here that he describes as a form of supervenience. I am not even sure it should be classified with the views more commonly called supervenience—at any rate, if it is, it is an unorthodox supervenience. The view has some very surprising features: no sentence with vague vocabulary in it is true, or false, and truth is not part of what we aim for when making assertions, at least not in general. Lewis's responses to influential criticisms of supervenience from Williamson reveal several other surprising features of his view: an apparent pluralism about logical consequence that countenances paraconsistent logic as one correct system of logical consequence; that jointly inconsistent claims are all assertable about stages in a sorites sequence; that no case of e.g. tallness is, or could be, D*tall; and some very surprising constraints on when claims about what is true are even assertable. Or so I will try to bring out below.

The letter divides into three main sections. In the first (pp1-2), Lewis sets out his own views about the connection between evidence (or one thing we could call 'evidence') and knowledge (or things 'know' could mean in various contexts). The second section (pp 3-8) is a sustained discussion of Williamson's own arguments about evidence, knowledge and luminosity. (E.g. the question of whether you always know what your immediate sensory evidence is.) The third section, from p 8-11 is a defence Lewis offers of his own "supervenience" views about vagueness from influential criticisms that appeared in Williamson's 1994 book *Vagueness*.

The first two sections are very interesting for epistemologists, especially those following the debates on luminosity of mental states (roughly, the doctrine that there is a range of mental states that you automatically in a position to know that you are in them when you are in them: luminosity about pain, for example, is the doctrine that when you are in pain you are automatically in a position to know that you are). Despite this interest, I will focus my discussion on the third part of the letter, on supervenience about vagueness.

This third section, together with some things Lewis says in a slightly later letter to Williamson in June 1999, have the status of an underground legend among some philosophers working on vagueness. I first heard rumours of this letter in 2001 or 2002. Before I heard about this letter, I had assumed Lewis was a "supervenientist" about vagueness in the sense of holding a view like that expressed in Fine 1975 or Keefe 2000, and the view criticised under that name by Williamson 1994. Lewis's remarks about the "linguistic theory of vagueness" in *On the Plurality of Worlds* (Lewis 1986) as involving semantic indecision, for example, are often read as expressing standard supervenientist sentiments. It becomes clear in this letter that Lewis's view by the late 1990s was *not* that view. As Lewis himself notes in his letter, there are some remarks in his "Many, But Almost One" (1993) paper that also suggest a very non-standard supervenientist attitude to vagueness.

I will make a few remarks about the first section, and despite the richness of the second section of Lewis's letter, I will skip Lewis's responses to Williamson's anti-luminosity arguments.

Section 1 of the Letter: Lewis, Evidence and Knowledge (pp 1-2)

Section 1 is illuminating about a number of Lewis's opinions, particularly about evidence.

Williamson famously argues that all and only our knowledge is our evidence ('E=K', as the slogan goes). Lewis's stipulative use in *Elusive Knowledge* (Lewis 1996) is in part in agreement with this, and part in disagreement with this. For Lewis, one's "evidence" is fixed by "the subject's entire perceptual experience and memory": what this is does not change as we raise or lower our standards for use of the word "know", and Lewis's framework ensures that we automatically know anything that is entailed by the occurrence of our perceptual experience and

(perceptual) memory. (Lewis 1999 p 1, see Lewis 1996 p 553) Lewis's usage requires him to disagree, however, that everything we know is part of our evidence: or more carefully, there will be many contexts where "I know something beyond my evidence" will be true to utter, and indeed in most contexts what counts as known will outstrip what counts as evidence, in that context.

But as Lewis intends this view, it is less opinionated about evidence than one might have thought: since Lewis intends his use of the expression "evidence" here to be stipulative, and he quickly concedes that we could use the expression differently and still be "within our linguistic rights" to do so. Even if Lewis's use of "ruled out by our evidence" went beyond what we could normally mean by that expression, adapting pieces of ordinary language as terms of art is not uncommon in philosophy, or in virtually any other systematic inquiry for that matter.

Lewis also offers an interesting suggestion about what "evidence" might mean in more ordinary uses. He suggests that "evidence" might be a context dependent expression, picking out different bodies of our knowledge on different occasions. This would give us many of the advantages of the view that our evidence is our knowledge: the fact that quite esoteric and theoretical bits of knowledge can be cited as "evidence" for even more esoteric and theoretical hypotheses, for example. But it would also help with many of the qualms critics have had about the E=K theory. Lewis mentions one qualm: that we often do not take the conclusions of investigations to themselves be evidence relevant to the topic at hand, even when we think we have gained knowledge through an investigation. (If I ask you for the evidence you have that is relevant to the question of whether James is guilty, you should not include the claim that James is guilty, even if you take yourself to know that.)

Lewis also stresses a distinction between psychological luminosity about experience and epistemic luminosity about experience. Epistemic luminosity is the thesis that when you have an experience (of a suitable kind), you know, or are disposed to know, that you have that experience (and something informative about the experience, presumably). Psychological luminosity is the doctrine that when you have an experience of a suitable kind you are disposed to *believe* that you do. The latter is a more psychological claim, at least in the sense that it does not settle your

capacity to know about the experience. On some views of knowledge, epistemic luminosity will entail psychological luminosity: if knowledge requires belief, and luminous states go along with *knowledge* itself rather than just being in a "position to know", epistemic luminosity requires psychological luminosity. Even if it did, the two might come apart in the other direction. This distinction plays some complicating roles in Lewis's response to Williamson in section 2: but despite its interest I will pass over that response here.

For those wishing to track down the Williamson papers Lewis is responding to in section 2, the " $\underline{J}\Phi$ paper" is Williamson 1996, and the "Rutgers paper" Lewis is responding to became Williamson 2000. I should point out that the published version of the latter paper may not exactly match the paper as it was presented when Lewis heard it.

Section 3 of the Letter: Lewis on Supervaluationism (pp 8-10)

The view Lewis articulates as his supervaluationist view about vagueness is surprising to those who expect him to be a typical supervaluationist. Let me begin by contrasting what Lewis says with a more orthodox approach. What I think of as "orthodox" supervaluationism involves the following claims. First, that sentences with vague expressions in them can be associated, one-to-many, with "precisifications": assignments of precise semantic values to a sentence (or other expression). The usual semantic value associated by each precisification with an entire sentence will be a *proposition* with an entirely determinate truth-value. Some of these precisifications are "admissible" precisifications: assigning that precise semantic value to an expression would not get the wrong truth-value for any of the sentences that the expression appears in which are determinate in truth value. (There is not complete agreement about what makes a precisification an admissible one: you might think that an admissible assignment respects "all the rules of use for the expressions in the sentence", for example, but only if there is nothing about the rules of use for the expression that guarantees that sometimes it fails to apply completely precisely!)

Next, these sentences are *supertrue* if and only if all of their admissible precisifications are true; *superfalse* if and only if all of their admissible precisifications are false. A sentence in a context is plain *true* if and only if it is supertrue; plain *false* if and only if it is superfalse, and neither true

nor false otherwise. Finally, supertruth plays the role in assertion that you would normally take truth to play: plausibly (but controversially), a claim is assertable only if it is supertrue.

The way sentences interact with a sentential "determinacy" operator ("it is determinate that...") for orthodox supervaluationists is as follows: a sentence is determinately true on an admissible precisification provided all of its admissible precisifications are true. Or to put it directly in terms of the operator, $D\Phi$ is true on a given admissible precisification when Φ is true on all all admissible precisifications. An inference is deductively *valid* when it is necessarily supertruth-preserving: the supertruth of all of the premises always guarantees the supertruth of the conclusion. (This is, near enough, "global" validity in the sense of Williamson 1994 p 147-148.

Lewis's view is different. It shares some features with orthodox supervaluationism. It is true, for Lewis, that sentences with vague expressions in them are associated with a range of precisifications. The precisification of a complete sentence will typically determine a unique proposition with that sentence on an occasion of use, and that proposition will have a precise truth value. Some of the precisifications will be the admissible ones. Those sentences are *supertrue* if and only if their admissible precisifications are true, and *superfalse* when false on all precisifications. But Lewis does not identify supertruth with truth. Instead, supertruth is a feature of sentences that plays many of the usual roles truth is supposed to play. For Lewis, sentences containing vague expressions are not *true* at all (or false for that matter). Lewis sometimes points out an analogy with ambiguity: one plausible view of an ambiguous sentence is that it is not true or false, only true or false *relative to a disambiguation*. One difference between vagueness and ordinary ambiguity is that we would not want to insist that an ambiguous sentence is only assertable when it is true on all disambiguations (or near enough): it is fine for me to say "I went to the bank to withdraw some cash", even if I went nowhere near a river bank, provided it is clear enough in the conversational context which disambiguation I intended.

Lewis's propositions, as sets of possible worlds, have entirely precise truth-conditions: they are true when they contain the actual world, false when they don't, and there are no borderline cases of set membership. (Pedants may object that propositions *are* truth conditions, and so do not *have* truth-conditions, but since there's a perfectly good sense in which we can state the

conditions under which one is true I'll talk this way, trusting pedants can translate.) Given a context, a precisification of a sentence will be a proposition, so in the background for Lewis there is plenty of objects with precise truth-conditions: the phenomenon of vagueness comes from the fact that the relationship between pieces of vague language and the various precise semantic values is many-to-one.

A second way in which Lewis's view departs from what I identified as supervaluational orthodoxy arises in his account of assertability. Lewis does not require supertruth as a necessary condition for assertability: rather, the necessary condition is "*supertruth, or near enough to supertruth*". (Lewis puts this in terms of what the "goal of a cooperative speaker attempting to impart information" is, but I take it that Lewis is still happy for conversational pragmatics to play a role: just as we typically think not everything true is useful to say to communicate, Lewis would agree that not everything supertrue-or-near-enough would be useful to say in communicating.)

Incidentally, I think there is more work for a Lewisian to do here. The picture of communication we get from Lewis's writings in *Convention* (Lewis 1969) and later works is a picture where the goal of communication is, in part, getting our listeners to form beliefs: in the simplest case, I want you to form true beliefs through my making true claims. The additional work is to spell out how my uttering supertrue-or-near-enough claims is conducive to your forming *true* beliefs; or alternatively, a story should be given about what status the beliefs formed should have, if it is not truth. My hunch is that Lewis would have wanted to hold on to the standard being a matter of the listener forming true beliefs even when thinking about vague topics, and to handle the apparent vagueness in our attitudes to the world in a different way to the way he handles vagueness in language. But putting all of the pieces of a view like this together is not a trivial task.

In one respect, this is a radical view indeed. Almost all of our claims contain some vague language: perhaps all of them do. But on many occasions, the vagueness seems not to matter. "Rover knocked over his water bowl again" is usually uttered where the fuzziness of Rover's boundaries, which artifacts exactly count as water bowls, and how similar an event must be to a previous occurrence to count as happening "again" do not matter: however those boundary cases

were resolved, within reason, it would not have affected the truth of the sentence. But on Lewis's view, a sentence like the above is *never* true, and not even the kind of claim that *could* be true, given what it means. Instead, it is associated with multiple propositions (or functions from contexts to propositions), which might on a typical occasion of utterance be all true, and this is what we exploit when we use the sentence to convey information. Part of what is radical about this is that it makes the official theory of truth rather at odds with what we would ordinarily say and think about truth: for surely we would ordinarily count that sentence as a true one, and sincerely say it was true (unless we suspected the utterer of wanting to frame poor Rover). We might report that Jane is a truth-teller even if every sentence she has uttered in her life had some expression or other in it susceptible to vagueness: that is, on Lewis's view, even if she had never uttered a true sentence in all her days.

Note also that Lewis's presentation of the theory is not itself intended to consist of true claims, or even the kinds of claims that could be true. (Rather, he presumably intends that it be close enough to supertrue.) Lewis would no doubt have responses to pointing out these features that would seek to make them seem less objectionable. Braun and Sider 2007 defend a similar view to the one articulated here by Lewis, and they say a lot to motivate their view despite it having these odd features.

Further surprising features of Lewis's view come to light when he turns to discuss a range of the specific objections Williamson offers against supervaluationism in Williamson 1994, chapter 5. Lewis's references are to subsections of this chapter: so e.g. "V.3" is section 3 of Williamson's chapter 5.

Lewis on Section V.3

Back to Lewis's picture of vague language. Lewis's refusal to identify truth with super-truth leaves him extra flexibility when it comes to saying what it is for an argument to be valid. Surprisingly (surprisingly to me anyway) Lewis endorses a kind of pluralism about logical validity. At least this is how I interpret "I don't think I need to choose once and for all what shall be my official notion of logical validity, in other words of truth-preservation". A more

deflationary reading would be to take Lewis to merely be making a remark about what sorts of things he needs to include in his "official" theory, but it seems to me that he is indicating that there are a range of notions of something like "truth preservation" that would be useful to have available in theorising, and no harm will come from not deciding one is "the" notion of logical validity, provided of course that we do not slide between them "when it matters". In particular, Lewis also counts the system he discussed in "Logic for Equivocators" (Lewis 1982) as one of the "conceptions of truth-preservation" that we may wish to treat as logical validity, for some theoretical purposes. Even a system as radically non-classical as Priest's LP or Dunn's RM can count as a correct account of logical validity.

I am torn between three interpretive options about what is going on here. The first, which I lean more towards, is that Lewis is endorsing a pluralism about logical validity: that each of these systems, and perhaps more besides, count as defining a notion that is one of the ones that deserves the name "logical validity". Another option, however, is that "logical validity" or "truth preservation" requires some disambiguation or resolution of unspecificity: and different resolutions are appropriate in different contexts. This second option is very close to pluralism, but differs from it in having fewer resources to talk about the various "logical validities" at once: once, in a context we have settled on one, then in that context it is correct to say that it is the unique relation of logical consequence, and other characterisations of logical consequence are just mistaken. (This kind of contextualist sounds like an out-and-out monist, when speaking entirely within a single context.)

A third would be that *none* of these notions deserve to be called "logical validity" outright: but they all serve as potential replacements of that original commitment. This would be analogous to the idea that there is no sentential truth for sentences with vague expressions: no sentential truth, and so no sentential truth-preservation. Despite the parallel with what Lewis wants to say about truth for vague sentences, I think this option does not fit the text of Lewis's letter quite as well.

Lewis on Section V.4

On pp 153-154 of Williamson 1994, Williamson points out that the supervaluationist must say, in the typical sorites, that there is a number n , such that n grains of sand do not make a heap but $n+1$ does. And this just strikes us as wrong: the idea was that one grain of sand was too small to make the difference. ("Supervaluationism makes the very claim that they [many who have intuitions about vagueness] find incredible", p 153). In Williamson's view, this uncomfortable commitment is not helped very much by pointing out that for each n it is not supertrue that *this* n is the number such that n grains of sand do not make a heap but $n+1$ do. (The generalisation "there is some n ..." lacks a witness: or at least the true generalisation lacks a true sentence about a specific witness.) Indeed, Williamson seems to be suggesting that it makes things worse if the supervaluationist makes the "incredible" generalisation and then cannot, even in principle, exhibit the specific counterexample.

Lewis admits to an even more odd consequence. Naturally, since it is supertrue that there is a number n , such that n grains of sand do not make a heap but $n+1$ does, that is assertable. But for each n in the border area, it is almost superfalse that *it* marks the boundary: for each n , in simple models, all precisifications except for one will either say that the n step and the $n+1$ step are both associated with heaps, or that they both fail to be associated heaps. Since the norm of assertion is to say things that are near-enough supertrue, it is assertable to claim, about each n , that *it* is not a number such that n grains of sand do not make a heap but $n+1$ does. Lewis's speaker is prepared to make a set of claims that are jointly inconsistent: the generalisation that there is a number in the sorites sequence such that... and then specific claims, about each number in the sequence, that *it* is not the number such that...

Despite this apparent inconsistency, at least at the level of sentences produced, Lewis appears unworried. "Maybe this got some getting used to, but I don't remember that it did." In the passage Lewis refers to here in Lewis 1993, in response to the charge that various of the things the supervaluationist says are "peculiar", he says "So it is. But once you know the reason why, you can learn to accept it." (p 173 of Lewis 1999b).

Incidentally, it is perhaps a little ironic that Williamson is pointing to the claim that there is a number n , such that n grains of sand do not make a heap but $n+1$ does, as pointing to something

objectionable in the supervaluationist account. Notoriously Williamson thinks that verdict is *correct*, and furthermore that there is a particular witness: in a given sorites sequence of sand-piling, there is a specific n that has that feature. We cannot know, even in principle, which particular one it is, by Williamson's lights. Williamson is careful to talk of *other* people who find the supervaluationist's verdict on the generalisation incredible.

Lewis on Section V.5

David Lewis and Hans Kamp have both suggested that something like supervaluations can be pressed into service to give a semantics for comparatives: see Lewis 1970 and Kamp 1975. I hedge a little because the discussion Williamson cites on this is Lewis 1970 (Williamson cites p 229 of the Lewis 1983 reprint of this paper), and despite the way that paper is sometimes read, I doubt the view of vagueness and valuations of a language given there is a kind of supervaluationism either, at least in any narrow sense of "supervaluationism".

The attempt to employ precisifications in the service of a semantics for comparatives (e.g. explaining the behaviour of "taller than" in terms of precisifications of "tall") would require a fair bit of discussion to characterise, and even more to go through Williamson's objections to it and Lewis's telegraphic responses, so despite the interest of the topic I will set it aside for this discussion. I will, however, point out two things that are relevant to Lewis's view of vagueness beyond the topic of comparatives:

Lewis does not demur to Williamson's assimilation of Lewis 1970's "delineations" to the supervaluationist's precisifications, which is by itself useful evidence about how Lewis saw the relationship between the treatment of vagueness offered in Lewis 1970 and his later writings on the topic.

At the end of Lewis's remarks on this section, he suggests that what the admissible precisifications are might be context dependent as well as vague. On p 3 of this letter he expresses some sympathy for the doctrine that some phenomena of vagueness are to be explained in terms of context, following Kamp (presumably Kamp 1975, but see also Kamp 1984). This

would complicate Lewis's account even further, since it would eventually need to include an account of how this context dependence worked.

Lewis on Section V.6

Most theories of vagueness face a "problem of higher order vagueness": if sentences receive some special treatment in "borderline" cases, there must be some demarcation between the borderline cases and the determinate cases: but it is on the face of it implausible that this borderline be entirely sharp. Just as (we might think) it should not be that some particular hair is the one that makes the difference between being bald and not, there shouldn't be a particular hair that is the one that makes a difference between being determinately bald and being in the grey area.

Which cases are definite cases of baldness, for example, seems to be itself a vague matter: eventually when considering a series of cases we pass from the definitely bald to the no-longer-definitely bald, but not at any precise point. Orthodox supervaluationists handle this by holding that expressions like "admissible precisification" and "borderline case" are themselves vague. While expressions like "there is a first admissible precisification that counts someone as bald but not determinately bald" are supertrue, for each precisification that counts someone as bald but not determinately bald it will not be supertrue that *that* one is the first one which counts someone as bald but not determinately bald. Or to put it another way, there will be a range of candidates to be the start of the grey area, and on different ways of precisifying "grey area", different ones of those candidates will count as the start of the grey area. So the boundary of the grey area will be indeterminate too.

Williamson describes a number of odd results that an orthodoxy supervaluationist may need to accept to be able to offer this theory of the indeterminacy of the boundary between clear cases and indeterminate cases. (I am not sure which ones Williamson intended as the basis for *objections* to the supervaluationist view, as opposed to just describing how such a theory might go.) It seems plausible that a supervaluationist might want to say the same sort of thing about the cases that are determinately determinate, and the cases about which it is indeterminate whether

they are in the grey area: that boundary is plausibly an indeterminate one too. And so on. But we can ask about which cases, if any are *determinate** or *D**: that is, they are determinate cases, and determinately determinate, and determinately determinately determinate... etc. (or to put it another way: they are not borderline, not borderline whether they are borderline, not borderline whether they are borderline whether they are borderline...).

Williamson suggests that the supervenient has a challenge in saying how the distinction between the *D** cases and the other cases works: if it is a sharp boundary, then we may have not done enough justice to higher-order vagueness; but if that distinction is itself a vague one, we will need an explanation of that vagueness that goes beyond the resources we have used to explain the levels of determinacy below that. (At least this seems to me the thrust of his discussion on p 160 of Williamson 1994.) Alternatively, Williamson notes, a supervenient could say that there are no cases that fall into the *D** category, for any possible sorites sequence. If no man is *D**tall, even if he were 100ft high, then the supervenient would not have to explain how the transition from *D** cases to the others should go. Williamson has little to say directly about this option: he points out that it would mean there would be no "perfectly straightforward application" of a term (in the sense that at some level of iteration we could make a claim about its application that is vague?), and that this means that a supervenient would have to conduct her discussion of vagueness in a vague meta-language (p 161). Reading between the lines, though, I suspect Williamson thinks that driving a supervenient to a position where she admits there are no *D** cases of any vague term is a reductio. Even Everest would fall into a borderline associated with "mountain" somewhere up the scale, and a cue-ball smooth hairless head would not be *D**bald. My guess, admittedly going beyond the text, is that Williamson would think those cases are "perfectly straightforward applications" of language, in his sense, and in the ordinary sense.

Lewis takes this option in dealing with a theory of what is *determinate**, suggesting that the *borderline** cases in a sorites sequence *will* include all the cases in the series. As Lewis notes, this removes the risk that there will be a sharp line between the *borderline** cases and the others. When Lewis suggests this is "not because *borderline** is itself vague" he seems to be suggesting that "*borderline**", as opposed to "*borderline*", is not a vague expression. It is possible that he

just thinks any vagueness present is not relevant to its behaviour in this case. That is a surprising commitment for him to make. Even more surprising is the claim that "Those who think otherwise are lazily thinking that borderline* is a special variety of borderline, which it isn't".

I am not sure why Lewis says this. I am sure it is not "laziness" that leads people to think that exceptionally clear cases are D* cases. Consider a sequence of colour strips from red to blue, arranged in the colours of the rainbow. We might well think that there is a vague boundary between red and orange: and if pressed, we might also think that which strips are in the boundary zone and which are the clear red and orange cases is also somewhat vague. On reflection, we might think the boundaries for "is in the boundary area between the red strips and the (definitely) boundary strips" is also not an entirely settled matter, and the boundary for this category is also somewhat vague... and so on. Why this should make us think that somewhere in the hierarchy of borderline-of-borderline etc. we will find *every* shade of the colour spectrum ("every case is borderline*")? On the contrary, even after we include the borderline cases, the cases that are on the borderline of the borderline, the cases that are on the border of the previous classification, and so on, I do not see a reason why we would *ever* include e.g. one of the blue strips in that series. (Once you have even got as far as yellow, you have moved too far from the borderline of the red cases to even be in the border of the border, or the border of the border, or...) For that matter, I do not see the pressure to include every red strip somewhere in this hierarchy: and if a red strip (e.g. the first) is not somewhere in the hierarchy, then it will not be among the borderline* cases.

My hunch is that Lewis had the makings of an argument that every case is borderline* when we face a sorites sequence: and I assume it's not just a reworking of Williamson's argument that the supervaluationist is in trouble if not every case is borderline*. But I have to confess I do not know what that argument would have been. (I also think that this argument has an implausible enough conclusion that whatever it was it would not have been sound: but it might help illuminate Lewis's approach to know what he had in mind.) I also suspect that the specific accusation of laziness is one that Lewis would retract on reflection: many of the theorists who have written on higher-order vagueness have endorsed the view that not every step in a sorites sequence is borderline*, and even if people who have spent years grappling with the problem

disagree with Lewis, it is unlikely to be due to a lack of effort in trying to think through the problems.

Lewis on Section V.7

Williamson points out (p 162-164) that orthodox supervaluationists will have trouble keeping the T-scheme. Take a sentence S that expresses a true proposition on some, but not all, precisifications (e.g. "Bill is tall"), and consider an acceptable precisification on which Bill is tall. While "Bill is tall" is true on that precisification, "'Bill is tall' is supertrue" is not. For the latter sentence to be true on any acceptable precisification, it must be true on all acceptable precisifications: and by hypothesis S does not meet that condition.

If we were to identify truth with supertruth, $S \equiv T\langle S \rangle$ would not itself be true on all precisifications, and so would not be true. To the extent that the material conditional can stand in for "if" (or even if the material conditional is entailed by "if"), "'Bill is tall' is true if and only if Bill is tall" will not be supertrue (i.e. it will not be true). Orthodox supervaluationists do have some resources similar to the T-scheme they can employ. If they adopt an approach to validity that requires that valid inferences never take you from a set of true premises to a conclusion that fails to be true, they can hold on to the logical validity of the inference from "Bill is tall" to "'Bill is tall' is true" and vice versa. An "if" that signals entailment (i.e. if A then B is true if, and only if, 'B' is a valid consequence of 'A') can still link S and $T\langle S \rangle$. But Williamson complains this does not give us what we wanted: we wanted sentences of the form "S is true" to have the same truth-conditions as S.

Williamson also points out, as Fine 1975 does, that there is another notion of truth, truth_T available to a supervaluationist, which does ensure that "S is true_T " is true on a precisification just when S is. Williamson says that, given that truth_T has this feature and supertruth does not, it is hard to see why we should identify truth with supertruth rather than truth_T .

Lewis evades the letter of both charges by refusing to identify truth with supertruth. Furthermore, as he points out, he can treat *propositional* truth in a disquotational way: whenever a proposition

p obtains, the proposition that p is true will also obtain. (Indeed, given Lewis's intensionalism about propositions, these will literally be the same proposition: but let me set aside any potential oddness of this result here.) So one key notion of truth in his system is "deflationary", even if he has to say something less direct about the truth of sentences.

He faces a similar charge, however: sometimes a sentence S will be close to supertrue, while the claim that S is supertrue is superfalse, and so the claim that S is not supertrue will itself be supertrue. So both "Bill is tall" and "'Bill is tall' is not supertrue" will be assertable, when "Bill is tall" is true on almost all precisifications, but not true on all of them. This blocks any proposal that we should say "S is true" when and only when S is supertrue (or near enough to supertrue). Supertruth cannot guide *assertions* of the truth of sentences in one way a Lewisian might have been tempted by.

It is hard to see what Lewis *should* do with ordinary talk of sentences containing vague language being true. Strictly speaking, it appears, such sentences are not true, since that requires the absence of vague expressions. But it would make a hash of a lot of our talk of sentences being true to insist that this talk is not even acceptable, and there is some sort of widespread error that undermines our talk of sentential truth. And we should not, in general, assert that they are true only when they are supertrue, at least if we want to be in a position to assert a sentence's truth, roughly, when we are in a position to assert the sentence. (And perhaps in some other cases as well: I should be able to say you wrote down something true, even if I am not in a position to assert what you wrote. Perhaps I don't know exactly what you wrote, or I don't understand it, or for some other reason it was within your rights to say but not mine.)

My guess is that Lewis should think the standard for assertion that a sentence is true, at least outside rarefied contexts like that of giving a theory of vague language, should track the assertion conditions for the sentence itself, at least in simple cases. Any way of giving a good answer to the question of the standard for assertion of sentential truth, however, runs a risk of undermining Lewis's theory. Suppose we find a feature of vague sentences, T , such that its having T explains why it is correct (or acceptable, or whatever) to ascribe sentential truth to those sentences. Then there would be a case, perhaps a strong case, to be made that T was sentential truth (or sentential

truth for sentences containing vague language) after all. But if *T* is sentential truth, then Lewis's claim that these sentences are all not true must be rejected after all.

This objection is probably not decisive: no doubt the Lewisian would have things to say in response. (Though, dear reader, it would be fine with me if you do regard it as decisive!) Lewis offers us in this letter a radical alternative to orthodox supervaluationism, which avoids some of the objections that can be raised to its orthodox cousin and has some other appealing features besides. It also faces serious challenges that merit further work. I recommend that we recognise it as a species of the linguistic theory of vagueness that deserves to be included among the salient options when theorising about the working of vague language.

Wrapping Up

Lewis's remarks on vagueness are intriguing and at a number of points tantalising. I suspect that had he lived, we would have seen more contributions from him on the topic of vagueness, since vagueness, and semantic indeterminacy more generally, were issues that lurked in the background of many topics he addressed. The early 2000s saw a mini-boom in the literature on vagueness as well, which may have made it especially likely that his attention would have been drawn again to these issues. There is certainly plenty of conceptual space still to explore for theories that treat vagueness as some sort of semantic indecision but do not treat it in the way orthodox supervaluationism does.

My thanks to all of those on the *Age of Metaphysical Revolution Project* at the University of Manchester, and particularly to Anthony Fisher for an invitation to be a guest poster on the Letter of the Month blog. Thanks also to Timothy Williamson for being willing to have me present this letter. I am sure that as more of David Lewis's philosophical letters become publicly available, a lot of additional exciting material will become available: exciting not just to those of us with a particular interest in Lewis's thought, but to those working on the many areas of philosophy in which he made valuable contributions.

References

- Braun, D. and Sider, T. 2007. "Vague, so Untrue". *Nous* 41.2: 133-146
- Fine, K. 1975. "Vagueness, Logic and Truth". *Synthese* 30: 265-300
- Kamp, R. 1975. "Two Theories of Adjectives" in Keenan, E.L. (ed.) *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press.
- Kamp, R. 1981. "The Paradox of the Heap" in Mönnich, U. (ed). *Aspects of Philosophical Logic, Some Logical Forays into Linguistics and Philosophy*. Dordrecht: D. Reidel
- Keefe, R. 2000. *Theories of Vagueness*. Cambridge: Cambridge University Press.
- Lewis, D. 1969. *Convention*. Blackwell, Oxford.
- Lewis, D. 1970. "General Semantics". *Synthese* 22.1-2, pp 18-67. Reprinted in Lewis 1983.
- Lewis, D. 1983. *Philosophical Papers, Volume 1*. Oxford: Oxford University Press.
- Lewis, D. 1982. "Logic For Equivocators". *Noûs* 16.3: 431-441.
- Lewis, D. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, D. 1993. "Many, But Almost One" in Campbell, K., Bacon, J. and Reinhardt, L. (eds) *Ontology, Causality and Mind: Essays on the Philosophy of D. M. Armstrong*. Cambridge: Cambridge University Press, pp 23-38. Reprinted in Lewis 1999b, pp 164-182.
- Lewis, D. 1996. "Elusive Knowledge". *Australasian Journal of Philosophy* 74.4: 549-567
- Lewis, D. 1999a. Letter to Timothy Williamson, 21 May 1999. Available on the *Age of Metaphysical Revolution* website, letter of the month, October 2018.
< <http://www.projects.socialsciences.manchester.ac.uk/lewis/letter-of-the-month-november-2018/> >
- Lewis, D. 1999b. *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge.
- Williamson, T. 1994. *Vagueness*. London: Routledge.
- Williamson, T. 1996. "Cognitive Homelessness". *Journal of Philosophy* 93.11: 554-573
- Williamson, T. 2000. "Scepticism and Evidence". *Philosophy and Phenomenological Research* 60.3: 613-628