

Reply to Yli-Vakkuri

Timothy Williamson

In 'Epistemicism and modality', Juhani Yli-Vakkuri rigorously treats the problem of developing a model-theoretic semantics for a formal language with the operators 'necessarily' and 'actually', interpreted in terms of metaphysical modality, and an operator 'definitely', interpreted in the spirit of the epistemicist theory of vagueness developed in my book (Williamson 1994). The challenge is especially pertinent because my epistemicist account of definiteness has a modal dimension, whose interaction with other forms of modal variation raises delicate issues for the overall theory.

1. *Worldly differences and semantic differences*

According to epistemicism, vagueness in a language induces no deviation from classical logic or from bivalent semantics. It is a purely epistemic phenomenon: borderline cases are just cases of unavoidable ignorance for a special sort of reason. The actual meaning of a vague

expression differs from some of its potential alternative meanings in ways indiscriminable by native speakers of the language. Following John Hawthorne, Yli-Vakkuri calls this idea 'Semantic Plasticity'. The upshot is a special obstacle to knowing some truths expressed using the vague expression. Very roughly, the definite truths are those knowledge of which is not blocked by an obstacle of that sort. I developed the underlying epistemology in safety-theoretic terms. Knowledge requires safety from error: one knows in a given case only if it is not relevantly close to an error possibility. Actuality is never close in that sense to metaphysical impossibilities, because they present no real danger.

Compare a case where one assents to a true proposition expressed by a vague sentence with a similar case where one assents to a false proposition expressed by the same sentence.¹ The difference in truth-value may result from either one of two kinds of difference between the cases (perhaps from both). The first kind of difference is *worldly*. The vague sentence may express the same proposition in the two cases, while the objects under discussion vary in their properties or relations. For instance, in both cases the sentence 'Jack is bald' expresses the proposition that Jack is bald, but Jack has less hair in one case than in the other. Thus the sentence 'Jack is bald' may express a true proposition in the former case and a false proposition in the latter. The second kind of difference between cases is *semantic*. Even if the objects under discussion have the same properties and relations in the two cases, the sentence 'Jack is bald' may express different propositions, because it is used in very slightly different ways. For instance, in one case the sentence 'Jack is bald' still expresses the proposition that Jack is bald, while in the other it expresses the distinct proposition that Jack is bald*, where the standard for being bald* is slightly more demanding than the standard for being bald. Thus, even though the state of

Jack's scalp is exactly the same between the two cases, the sentence 'Jack is bald' may again express a true proposition in the former case and a false proposition in the latter. Of course, two cases may also differ in both worldly and semantic respects simultaneously, and in other respects too.

Concerning our ignorance in paradigmatic borderline cases, I proposed that the error possibilities by which vagueness blocks knowledge are the close ones differing from the given case in semantic but not in worldly respects (1994, p. 231). By contrast, most discussions of safety in epistemology concern worldly but not semantic differences; the proposition at issue is held fixed across the cases under discussion. However, if one wants a non-trivial safety condition for knowledge of necessary truths, as in mathematics, one must look at semantic rather than worldly differences: if a sentence expresses a necessary truth in a given case, then it expresses a truth in all cases differing from the former only in worldly respects (assuming that all cases, or at least all close cases, are metaphysically possible). Thus there are reasons independent of vagueness for using a safety condition that takes account of cases that differ in semantic respects, as well as of those that differ only in worldly ones (Williamson 2000, pp. 181-2). However, the semantic differences relevant to mathematical knowledge will typically be much larger than those relevant to vagueness; the cases will not be semantically indiscriminable to native speakers.

For straightforward examples, we can treat the worldly and semantic dimensions as mutually orthogonal: we can vary a case on either dimension while holding it fixed on the other. For more complex examples, however, the independence assumption breaks down: varying a case on one dimension entails varying it on the other too. For instance, a vague word may be both used and mentioned in the same sentence, as in disquotational

biconditionals for truth such as “‘Jack is bald’ is true if and only if Jack is bald’, a type of example much discussed in the recent literature on vagueness and featured in Yli-Vakkuri’s paper. A difference between cases in the use of the word ‘bald’ is semantic in respect of the right-hand side of the biconditional (‘Jack is bald’) but worldly in respect of the left-hand side (‘“Jack is bald” is true’), since the latter is about a broadly semantic feature of the former.

It is often unclear how to apply the rough dichotomy between the world and our descriptions of it to examples where we cannot ignore the fact that our descriptions are part of the world. Similarly, it is often unclear how to apply the rough dichotomy between the world and our knowledge of it to examples where we cannot ignore the fact that our knowledge is part of the world, as in the so-called paradox of knowability (we cannot know a conjunction of the form: P and it will never be known that P ; see Williamson 2000, pp. 270-301). Yet any theory of vagueness must be applicable to such complex examples. Problems of vagueness in the law remind us that the distinction between clear cases and borderline cases arises in practice for arbitrarily complex sentences. One benefit of trying to construct an adequate model theory for a formal language with a ‘definitely’ operator is that it forces us to think systematically and quite generally about how to apply vagueness-theoretic distinctions to such sentences.

At this point, some philosophers will be inclined to scrap the account of definiteness based on the worldly-semantic distinction as insufficiently general, and start again from scratch. But that is just the sort of methodological impatience that so much impedes progress in philosophy. In particular, the fact that the worldly and semantic dimensions are not fully independent of each other does not entail that an initial assumption of their

independence makes a bad starting-point. After all, if we are trying to understand a phenomenon in natural science, we may initially make the simplifying assumption that two aspects of it are mutually independent, even though we know that the assumption is not strictly correct. For instance, we may treat propositions about the shape of an object as probabilistically independent of propositions about its colour, so that we can calculate the probability of the conjunction of a shape proposition and a colour proposition as simply the product of the probabilities of the conjuncts. We can then use the results obtained under that simplifying assumption as a starting point for a more general account that allows for the interdependence of the two dimensions. The same goes for the treatment of the worldly and semantic dimensions as mutually independent in the initial explanation of ignorance in borderline cases. Rather than scrapping the first pass explanation, we may do better by refining it. A specific reason for optimism is that vagueness in complex expressions derives somehow from vagueness in their atomic constituents and modes of composition. The worldly-semantic dichotomy works best for atomic expressions, although doubtless not perfectly even there.

A further consideration comes from the nature of the 'definitely' operator. Its treatment as a sentence operator in a formal language with a compositional semantics turns out to involve a significant degree of idealization. Within limits, the theorist has discretion to decide what idealizations of its meaning will conduce to the most illuminating logic of vagueness.

'Definitely' is not just a philosopher's invention. To describe someone as 'neither definitely tall nor definitely not tall' is quite intelligible in ordinary English, even if it sounds a bit pedantic. Although the complement of 'definitely' in English is normally an adjectival

phrase rather than a sentence, that goes for modal words such as ‘necessarily’ and ‘possibly’ too; it does not prevent them from being most usefully formalized by operators on sentences. Many ordinary uses of ‘definitely’ are epistemic in a way not specifically related to vagueness. For instance, a detective asking ‘Who was definitely in the building last night?’ probably means something like ‘Who is known to have been in the building last night?’; borderline cases of being in the building last night are probably not at issue. Still, theorists of vagueness can legitimately requisition the word ‘definitely’ for their more specific purposes. So far, so good.

2. *Use and mention*

A more pressing concern is that the formalization of ‘definitely’ as a sentence operator Δ tailored to suit the theory of vagueness may not fit the structure of a compositional semantics. To start simple, consider a semantic theory that assigns a set of (metaphysically) possible worlds as the semantic value of a sentence. The compositional semantics makes the semantic value of the sentence operator Δ a function from sets of possible worlds to sets of possible worlds. Let ‘small’ be a vague predicate of natural numbers whose application to them is nevertheless non-contingent. Thus, we may suppose, 17 is a borderline case for ‘small’; it is vague whether 17 is small, but not contingent whether 17 is small. If 17 is small, it is necessary that 17 is small; if 17 is not small, it is necessary that 17 is not small. Consequently, the semantic value of ‘17 is small’ is either the set of all worlds or

the set of all worlds. In the former case, '17 is small' has the same semantic value as ' $0 = 0$ ', so by the compositional semantics ' $\Delta(17 \text{ is small})$ ' has the same semantic value as ' $\Delta(0 = 0)$ ', and so is true (definitely 0 is self-identical). In the latter case, the semantic value of ' $\neg(17 \text{ is small})$ ' is the set of all worlds and so the same as the semantic value of ' $0 = 0$ ', so by the compositional semantics ' $\Delta\neg(17 \text{ is small})$ ' has the same semantic value as ' $\Delta(0 = 0)$ ', and so is true. Either way, the disjunction ' $\Delta(17 \text{ is small}) \vee \Delta\neg(17 \text{ is small})$ ' is true: it is definite whether 17 is small. But that contradicts the hypothesis that it is vague whether 17 is small.

One might simply draw the moral that sets of possible worlds are too coarse-grained to serve as the semantic values of sentences in the semantics for a logic of vagueness and modality. That is hardly surprising. The envisaged semantics is one-dimensional, whereas Yli-Vakkuri argues powerfully that it needs to be three-dimensional rather than two-dimensional, to accommodate both the usual two modal dimensions for the operators 'necessarily' and 'actually' and a further dimension for the operator 'definitely'. But further reflection suggests that the problem goes even deeper. *Any* compositional semantics for Δ makes the meaning of the sentence $\Delta\phi$ a function of the meaning of the sentence ϕ , in some appropriate sense of 'meaning'. Thus if ϕ is synonymous with ψ (they have the same meaning in the relevant sense), then $\Delta\phi$ is synonymous with $\Delta\psi$, and moreover $\neg\phi$ is synonymous with $\neg\psi$, so $\Delta\neg\phi$ is synonymous with $\Delta\neg\psi$; thus ϕ and ψ have the same borderline cases. But is that principle defensible?

To see the difficulty, let us make the standard assumption that the natural kind terms 'furze' and 'gorse' are synonymous. They are also presumably not perfectly precise: for instance, it is slightly vague when furze and gorse first evolved. Nevertheless, for purposes of meaning, the two words are vague in the same way. Let ϕ be the sentence 'All

and only furze is furze' and ψ the sentence 'All and only furze is gorse'. On a standard compositional semantics, ϕ and ψ have the same semantic structure and synonymous corresponding constituents, and so are synonymous. Thus, by hypothesis, $\Delta\phi$ and $\Delta\psi$ are synonymous too. Since ϕ is a trivial logical truth, $\Delta\phi$ is true. Hence $\Delta\psi$ is true too. But is it? We may assume that 'furze' and 'gorse' are understood independently of each other; a speaker may understand either word without understanding the other. No special rule coordinates them; they just happen to have the same meaning. Although the use of 'furze' happens exactly to match the use of 'gorse', there could very easily have been a slight mismatch, where the use of one word was slightly more liberal than the use of the other, so that the extension of the former included slightly earlier samples than the extension of the latter did. The vagueness of 'furze' and 'gorse' makes it unclear to native speakers whether there is indeed such a slight difference in extension. Thus it seems to be definite that all and only furze is furze, but not definite that all and only furze is gorse. On this view, ϕ and ψ are synonymous, but $\Delta\phi$ and $\Delta\psi$ differ in truth-value and so in meaning. If so, the operator Δ does not have a compositional semantics.

Could a theorist respond by individuating meanings in a more fine-grained way? We may assume that the only relevant difference between 'furze' and 'gorse' is the syntactic difference between the words themselves. A way of individuating meanings on which 'furze' and 'gorse' differ in meaning makes synonymy between distinct words virtually impossible. A supervaluationist might implement such a fine-grained standard by identifying meanings with functions from admissible sharpenings of the language to Kaplanian characters. Suppose that one admissible sharpening assigns the character C to both 'furze' and 'gorse', while another assigns a distinct character C^* to both words. Then if a third admissible

sharpening assigns C to 'furze' and C* to 'gorse', while a fourth assigns C* to 'furze' and C to 'gorse', the account makes 'furze' and 'gorse' differ in meaning, even if the overall range of admissible sharpenings associates each of them with exactly the same range of characters. The trivialization of synonymy is a high price to pay for such a view, especially since it threatens to trivialize the requirement of compositionality too.

An alternative way of interpreting the semantics is to retain a conception of meaning that allows 'furze' and 'gorse' to be synonymous, but to treat the Δ operator as tacitly *quotational*: $\Delta\phi$ is really about the sentence ϕ itself, at least partly (a move Yli-Vakkuri may accept). In that case, the so-called compositional semantics for the language is not what it seems, since it does not determine the meanings of complex expressions just from the meanings of their constituents and the way those constituents are put together. For a genuinely compositional semantics, one might have to make the quotation explicit — and supply a semantics of quotation.

These problems are reminiscent of those facing attempts to give a compositional semantics of variable-binding, on which Yli-Vakkuri has written illuminatingly elsewhere (2013). In the simple case of an extensional, non-indexical language, if one identifies the meaning of a singular term with a function that maps each assignment of values to variables to the denotation of the term on that assignment, the variables 'x' and 'y' count as different in meaning, simply because they have distinct values on at least one assignment, even though that is just a corollary of their syntactic difference. In one respect, however, the problems raised by Δ are more acute, since they arise for closed as well as open expressions.²

Similar problems arise anyway for epistemic operators, and so pose no threat to the epistemicist view that Δ is an epistemic operator. For instance, even though 'furze' and 'gorse' are synonymous, it may be natural to describe someone as knowing that something is furze without knowing that it is gorse, or as knowing that all and only furze is furze without knowing that all and only furze is gorse. If one is doing the semantics of 'know' in terms of epistemically possible worlds, one may then be driven to assign 'furze' and 'gorse' distinct intensions (here, functions from epistemically possible worlds to extensions), even though they are synonymous.

Of course, one may take a more coarse-grained view of knowledge and simply tough it out, insisting that to know that all and only furze is furze just *is* to know that all and only furze is gorse, even if those who recognize that truth under the guise of the sentence 'All and only furze is furze' cannot recognize it under the guise of the sentence 'All and only furze is gorse'. In principle, someone might take the same line with vagueness and simply tough it out, insisting that for it to be definite that all and only furze is furze just *is* for it to be definite that all and only furze is gorse, even if those who recognize that truth under the guise of the former sentence cannot recognize it under the guise of the latter. However, the tough line is better motivated for 'know' than for 'definite'. When we speak of knowledge and ignorance, we often wish to prescind from the words in which a given piece of knowledge might be expressed. By contrast, when we speak of definiteness and vagueness, the expression in given words is typically central to our interest. To prescind from it would be to miss the very point at issue.

These questions also arise for Yli-Vakkuri's preferred epistemic interpretation of his three-dimensional semantics. To make all the desired distinctions, an intended model will

have to assign distinct 3D profiles to sentences that, by a natural compositional standard, are synonymous. That is very far from making the theory useless: it can still serve as a good guide to reasoning about the interaction between vagueness and metaphysical modality. But it will not be a compositional semantics in the sense one might originally have expected.

To the two non-epistemically modal dimensions needed for a standard semantic treatment of 'necessarily' and 'actually', Yli-Vakkuri adds a third modal dimension for the treatment of 'definitely', with promising results.³ He argues that the third dimension is best understood as epistemically modal. Since he is implementing an epistemicist account of vagueness, that is not altogether surprising. He requires the epistemic dimension to include metaphysically impossible epistemic possibilities as well as metaphysically possible ones. On his interpretation, 'when ϕ is borderline, ϕ expresses in the actual world, according to some close epistemic possibility, a proposition other than the one it actually expresses in the actual world', but 'any such epistemic possibility will be metaphysically impossible, because it is a noncontingent matter which proposition is actually expressed by a sentence'. Such a use of metaphysically impossible worlds in the semantics for 'definitely' poses an apparent threat to epistemicist theories that explain actual borderline cases in terms of Semantic Plasticity, for surely no metaphysically impossible world is close in the safety-relevant sense to the actual world.

However, we may be able to reconcile Yli-Vakkuri's semantics with the safety-theoretic considerations by elucidating its explanatory ambitions. For just as a theory of the meaning of the word 'light' is not tasked with explaining the underlying nature of light, and a theory of the meaning of the word 'knowledge' is not tasked with explaining the underlying nature of knowledge, so a theory of the meaning of the word 'indefinite' is not

tasked with explaining the underlying nature of indefiniteness. The point is easier to miss in the latter case, for while semantics is obviously separate from physics, and only slightly less obviously separate from epistemology, it is not at all obviously separate from the theory of indefiniteness as a linguistic phenomenon. A comprehensive book on semantics *should* say something about vagueness and indefiniteness. Nevertheless, that task must be distinguished from the task of giving the semantics of the word 'indefinite', just as the task of explaining the underlying nature of meaning must be distinguished from the task of giving the semantics of the word 'meaning'. Thus, even if safety-theoretic considerations are central to the underlying nature of indefiniteness, they need not figure in a good semantic account of the word 'indefinite' or the Δ operator, just as, even if safety-theoretic considerations are central to the underlying nature of knowledge, they need not figure in a good semantic account of the word 'knowledge'. By itself, the appeal to merely epistemically possible worlds in Yli-Vakkuri's semantics for Δ poses no threat to a safety-theoretic account of the underlying nature of indefiniteness, since the semantics is not expected to describe that underlying level. Of course, if for some sentences in some contexts the semantics predicted truth-values incompatible with the safety-theoretic account, that would be a problem, but just appealing to merely epistemically possible worlds forces no such incompatibility. The point can be restated in terms of the distinction between semantics and metasemantics, where a metasemantic theory is supposed to explain why a given semantic theory holds for the language of a given community at a given time. The possible worlds of the safety-theoretic considerations are the concern of the metasemantic theory; Yli-Vakkuri's impossible worlds are the concern of his semantic theory.

Also to be taken into account is the extent to which the Δ operator and related English words such as ‘indefinite’, ‘borderline’, and ‘vague’ are all being used as theoretical terms, and rightly so. Although their introduction by examples is good enough for initial working purposes, some refinement is needed once they are used in formulating a positive theory of vagueness. For when we encounter more complex cases than the usual sorites series provide, especially cases where the semantic and worldly dimensions fail to be independent, the initial paradigms give no clear guidance as to how those terms should most fruitfully be applied for theoretical purposes. We see many apparent joints, and are unsure where to cut. For any theory that uses an operator like Δ to characterize vagueness-related phenomena, choosing a semantics for Δ is a prescriptive as well as a descriptive task. It is a matter of deciding what it is *best* to mean by Δ . In this respect, the model-theoretic nature of Yli-Vakkuri’s semantics strikes a helpful balance between informativeness and flexibility. It is primarily designed to generate a good *logic* for Δ , especially as it interacts with the modal operators. It validates nice structural principles for Δ , in particular those corresponding to the modal system T, thereby imposing constructive discipline on our use of Δ . At the same time, it does not specify in any detail an intended model (or even frame) for the model theory. It thereby leaves considerable scope for fine-tuning the specific application of Δ in line with future metasemantic developments.

Like most accounts of vagueness, the epistemicist account in my book focusses on simple borderline cases, and does not provide clear guidance on how to apply an operator like Δ to more complex cases, especially those where the semantic and worldly dimensions are not independent. As already noted, that does not make it irrelevant to such cases. Science is full of cases where we know how to apply a theory to simple examples of a

phenomenon but not to more complex examples. Typically, progress is then incremental, as models (in the scientific sense, not the semantic one) are developed to permit the theory to be applied to more complex examples, though never all at once, and always under simplifying assumptions. I see no reason why such progress should not also be made in the theory of vagueness, if philosophers are willing to do the hard work of model-building.

3. *Semantic Plasticity*

Where do these general methodological considerations leave Semantic Plasticity? In section 5, Yli-Vakkuri argues directly against a version of it that appeals to metaphysical possibilities. He formulates the version thus:

- (*) Whenever ϕ is true and borderline, there is a close metaphysically possible world in which ϕ expresses a proposition other than the proposition ϕ actually expresses, and in which that proposition is false.

As a counterexample to (*), he discusses the case where ϕ is the material conditional $\# \rightarrow \chi$, χ is an actually true borderline sentence, and $\#$ is a relevantly precise actually non-borderline specification of the actual use facts that determine which proposition χ actually expresses (so $\#$ is actually true). Thus $\Delta\#$ and $\neg\Delta\chi$ are true, so $\neg\Delta(\# \rightarrow \chi)$ is also true by the normal logic of Δ . Therefore, since $\# \rightarrow \chi$ is true, it is also borderline. Thus, given (*), it follows that there is a close metaphysically possible world w^* in which $\# \rightarrow \chi$ expresses a proposition other than the proposition $\# \rightarrow \chi$ actually expresses, and in which that

proposition is false. We may assume that in all close metaphysically possible worlds $\#$ and \rightarrow express what they actually express. Thus if $\#$ is true in w^* , the actual use facts that determine which proposition χ expresses are the same in w^* as in the actual world, so $\# \rightarrow \chi$ expresses the same proposition as it actually expresses, contrary to hypothesis. On the other hand, if $\#$ is not true in w^* , it is false in w^* (given epistemicism, bivalence holds), so $\# \rightarrow \chi$ is true in w^* by the truth-table for \rightarrow , again contrary to hypothesis. Thus $\# \rightarrow \chi$ is a counterexample to (*). This argument against (*) does not depend on Yli-Vakkuri's formal semantics, although it does use the standard assumption that the logic of Δ is like that of \Box in a normal modal logic, which Yli-Vakkuri's semantics validates and I happily grant.

I will not defend (*). However, the falsity of (*) does not mean that the more general phenomenon of semantic plasticity as described in my book plays no role in explaining the borderline status of $\# \rightarrow \chi$, even when semantic plasticity is understood, as here, to be variation in what is expressed in close, so metaphysically possible, worlds (epistemically possible but metaphysically impossible worlds do not count). For the normal logic of Δ has these theorems, which Yli-Vakkuri's semantics also validates:

$$\Delta\# \rightarrow (\Delta(\# \rightarrow \chi) \leftrightarrow \Delta\chi)$$

$$\Delta\# \rightarrow (\Delta\neg(\# \rightarrow \chi) \leftrightarrow \Delta\neg\chi)$$

$$\Delta\# \rightarrow ((\neg\Delta(\# \rightarrow \chi) \& \neg\Delta\neg(\# \rightarrow \chi)) \leftrightarrow (\neg\Delta\chi \& \neg\Delta\neg\chi))$$

In brief, given that $\#$ is definite, the question whether $\# \rightarrow \chi$ is borderline reduces to the question whether χ is borderline. We can explain the definiteness of $\#$ in terms of its truth and lack of semantic plasticity, both of which Yli-Vakkuri stipulates. Thus, if we can explain the borderline status of χ in terms of semantic plasticity, then we can also explain the

borderline status of $\# \rightarrow \chi$ indirectly in terms of semantic plasticity, using our explanations of the definiteness of $\#$ and the borderline status of χ in those terms. For instance, if χ is a semantically simple sentence, we may be able to explain its borderline status directly in terms of its semantic plasticity, along the lines I originally proposed, for all that Yli-Vakkuri's arguments show. Of course, χ may itself be complex, and present similar difficulties to those $\# \rightarrow \chi$ did, but then we can apply the same strategy to χ as we did to $\# \rightarrow \chi$, and so on down in a recursive process terminating at the latest in semantically simple constituents.

An encouraging observation is that both the normal logic of Δ and Yli-Vakkuri's semantics yield the principle that any truth-function of non-borderline sentences is itself non-borderline. For instance:

$$(\Delta\phi \vee \Delta\neg\phi) \rightarrow (\Delta\neg\phi \vee \Delta\neg\neg\phi)$$

$$((\Delta\phi \vee \Delta\neg\phi) \& (\Delta\psi \vee \Delta\neg\psi)) \rightarrow (\Delta(\phi \& \psi) \vee \Delta\neg(\phi \& \psi))$$

The negation of a non-borderline sentence is itself non-borderline, a conjunction of non-borderline conjuncts is itself non-borderline, and so on. Contrapositively, if some truth-function θ of the simple sentences $\phi_1, \phi_2, \dots, \phi_n$ is borderline, then at least one of $\phi_1, \phi_2, \dots, \phi_n$ is borderline too. Thus if the explanation of borderline status in terms of semantic plasticity works for simple sentences, there is some hope of explaining θ 's borderline status indirectly in terms of semantic plasticity through the contributions of its simple constituents. One will not always be able to deduce θ 's borderline status from the borderline status of its constituents, for not every truth-function of borderline sentences is itself borderline: for instance, $\phi \& \neg\phi$ is a non-borderline truth-function of ϕ , even when ϕ is borderline.

Nevertheless, if θ is borderline, a reasonable goal is to explain its borderline status in terms

of the interrelated ways in which its constituents are borderline, where those ways essentially involve semantic plasticity.

More generally, suppose that *only* semantically plastic sentences are borderline; Yli-Vakkuri's arguments in no way undermine that principle. Let σ be a complex sentence with a compositionally determined meaning. Thus if all the simple constituents and modes of combination from which σ is composed are semantically *non*-plastic, so is σ itself; by hypothesis, then, σ is non-borderline. Contrapositively, if σ is borderline, then some of its simple constituents or modes of combination are semantically plastic. Thus we can hope to explain σ 's borderline status in terms of the semantic plasticity of those simple constituents or modes of combination. Although the present considerations do not guarantee success, they suggest that it is a reasonable strategy.

How does that strategy square with a safety conception of knowledge? For the truth actually expressed by $\# \rightarrow \chi$ is not known, yet in every close metaphysically possible world the sentence $\# \rightarrow \chi$ expresses a true proposition (by Yli-Vakkuri's argument). Moreover, it is quite compatible with the previous assumptions that the proposition actually expressed by $\# \rightarrow \chi$ is both actually believed and true in every close metaphysically possible world. Invoking epistemically but not metaphysically close worlds looks like abandoning the safety conception. If $\# \rightarrow \chi$ simply refutes that conception, then the original explanation of our ignorance in terms of semantic plasticity fails even for the simplest borderline cases, because it depends on a false epistemology.

At this point, it helps to recall how safety-theoretic considerations would explain the failure of irrational but true mathematical beliefs to constitute knowledge. Someone irrationally plumps for believing a proposition p expressed by a precise mathematical

sentence π . In fact, p is true, and so necessarily true. Thus in every close world, π expresses p (π is not semantically plastic) and p is true. However, in some close world w , he irrationally plumps in a similar way for believing a different proposition r expressed by a different precise mathematical sentence ρ . But r is false, and so necessarily false. The sentences π and ρ are not semantically indiscriminable to the subject; he understands both and is aware of their semantic difference. It is just that his plumping does not track their difference in truth value. Since his irrational false belief in the close world w is relevantly similar in method to his irrational true belief in the actual world, the latter fails to constitute knowledge.

Now consider Yli-Vakkuri's case. We may assume that the borderline true sentence χ is semantically plastic along the lines of (*), for otherwise his argument against (*) would assume more than it was entitled to by resting on stipulations hostile to (*). We may also assume that the subject actually believes the proposition expressed by the sentence $\# \rightarrow \chi$, for the challenge is to explain in appropriate terms why such a belief does not constitute knowledge. But then we should expect that in some close metaphysically possible world x , the subject believes in a similar way the proposition expressed in x by the sentence $\#^x \rightarrow \chi$, where $\#^x$ is a relevantly precise non-borderline specification of the use facts in x that determine which proposition χ expresses in x (so $\#^x$ is true in x), while by semantic plasticity the proposition χ expresses in x differs from the one it actually expresses and is false in x , even though x is just like the actual world in the worldly ways relevant to χ , though not in those relevant to $\#$ and $\#^x$. For we were not supposed to assume that the actual world was uniquely special in the subject's believing the proposition expressed by the material conditional with a true relevantly precise non-borderline specification of the use facts that

determine which proposition χ expresses as the antecedent and χ as the consequent. But then $\#^x \rightarrow \chi$ expresses a false proposition in x , for there $\#^x$ expresses a truth and χ expresses a falsehood. The false belief in the close world x in the proposition there expressed by $\#^x \rightarrow \chi$ is relevantly similar in method to the true belief in the actual world in the proposition actually expressed by $\# \rightarrow \chi$, so the latter fails to constitute knowledge. Thus we can explain in safety-theoretic terms the obstacle to knowing the truth expressed by $\# \rightarrow \chi$, and so the borderline status of $\# \rightarrow \chi$. The explanation crucially invokes semantic plasticity in relation to the vague consequent and a different safety-theoretic consideration in relation to the precise antecedent. That is a natural generalization of the original explanation by appeal to semantic plasticity of the obstacle to knowing the truth expressed by a simple vague sentence. The hybrid aspect of the explanans follows the hybrid aspect of the explanandum, where the sentence at issue combines the relevantly vague constituent χ with the relevantly precise constituents $\#$ and \rightarrow .

The previous paragraph made various simplifying assumptions. We can construct more complicated scenarios. For instance, we can imagine subjects hard-wired to believe propositions expressed by the sentence $\# \rightarrow \chi$ but not those expressed by similar sentences $\#^x \rightarrow \chi$ in any close world. On a suitable development of the story, the true beliefs expressed by $\# \rightarrow \chi$ of such bizarre creatures may bizarrely constitute knowledge.⁴ If so, and they also somehow know the truth expressed by the precise but massively detailed specification $\#$, may they not also know the truth expressed by χ inferentially, by modus ponens? But then χ is presumably not borderline for them. That does not mean that it is not borderline for us: our use of χ is different from theirs, even if we do in fact express the same proposition. If $\#$ correctly describes their use, then it does not correctly describe ours. If we enter into

communication with them, it may be unclear to us whether χ as they use it expresses the same proposition as χ as we use it.

The foregoing line of thought opens up many complex questions that I cannot hope to deal with adequately here. Rich explanatory resources are available to an epistemicist approach to vagueness articulated in terms of safety and semantic plasticity, given a willingness to apply those resources to complex cases in correspondingly complex ways, in the spirit of scientific model-building.

Notes

1 For the sake of brevity, when I speak of knowing or believing a proposition expressed by a sentence, I will mean that it is known or believed under the guise of that sentence.

Thus one may know or believe it expressed by one sentence (say, 'Hesperus is Hesperus') but not expressed by another sentence (say, 'Hesperus is Phosphorus').

2 The analogy may suggest a treatment along the lines of Fine (2007). This is no place to discuss such a radical departure from the usual form of compositional semantics; Yli-Vakkuri follows the latter.

3 The technical aspects of the logic were developed in collaboration with Jon Litland (Yli-Vakkuri and Litland 2016).

4 See Williamson 2016 for an analogous possibility concerning knowledge of mathematical axioms.

Fine, Kit 2007: *Semantic Relationism*. Oxford: Blackwell.

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

Williamson, Timothy 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, Timothy 2016: 'Absolute provability and safe knowledge of axioms', in Leon Horsten and Philip Welch (eds.), *Gödel's Disjunction*, Oxford: Oxford University Press.

Yli-Vakkuri, Juhani 2013: 'Propositions and compositionality', *Philosophical Perspectives* 27, pp. 526-563.

Yli-Vakkuri, Juhani 2016: 'Epistemicism and modality', THIS VOLUME.

Yli-Vakkuri, Juhani, and Litland, Jon 201X: 'Vagueness and modality', MS.