

Supervaluationism and Its Logics

Achille C. Varzi

Department of Philosophy, Columbia University (New York)

[Final version published in *Mind* 116 (2007), 633-676]

Abstract. If we adopt a supervaluational semantics for vagueness, what sort of logic results? As it turns out, the answer depends crucially on how the standard notion of validity as truth preservation is recasted. There are several ways of doing that within a supervaluational framework, the main alternative being between ‘global’ construals (e.g., an argument is valid if and only if it preserves truth-under-all-precisifications) and ‘local’ construals (an argument is valid if and only if, under all precisifications, it preserves truth). The former alternative is by far more popular, but I argue in favour of the latter, for (i) it does not suffer from a number of serious objections, and (ii) it makes it possible to restore global validity as a defined notion.

Supervaluationism is a mixed bag. It is sometimes described as the ‘standard’ theory of vagueness, at least insofar as vagueness is construed as a semantic phenomenon, but exactly what that standard theory amounts to is far from clear. In fact, it’s pretty clear that there isn’t just one supervaluational semantics out there—there are lots of such semantics; and although it is true that they all exploit the same insight, their relative differences are by no means immaterial. For one thing, a lot depends on how exactly supervaluations are constructed, that is, on how exactly we come to establish the truth-value of a given statement. (And when I say that a lot depends on this I mean to say that different explanations may give rise to different philosophical worries, or justify different reactions.) Secondly, and equally importantly, a lot depends on how a given supervaluational machinery is brought into play when it comes to explaining the logic of the language, that is, not the notion of truth, or ‘super-truth’, as it applies to individual statements, but the notion of validity, or ‘super-validity’, as it applies to whole arguments. (I am thinking for instance of how different explanations may bear on the question of whether, or to what extent, vagueness involves a departure from classical logic.) Here I want to focus on this second part of the story. However, since the notion of validity depends on the notion of truth—or so one may argue—I also want to comment briefly on the first.

1. Precisifications and Supervaluations

I take it that the basic insight of any supervaluational semantics boils down to the following two thoughts: first, a vague language is one that admits of several

precisifications; second, when a language admits of several precisifications, its semantics is fixed only insofar as—and exactly insofar as—all those precisifications agree. In particular, the semantic value of a statement is fixed only insofar as there is complete agreement on that value: the statement is true if it is super-true, that is, true on every admissible precisification, and it is false if it is super-false, that is, false on every admissible precisification; otherwise it has no semantic value. All of this, of course, presupposes that we know how to figure out the value of a statement on a precisification, but that's part of the idea: precisifications are semantically standard, hence our standard semantic algorithms apply just fine. So the idea is that when *several* precisifications are equally admissible, we apply those algorithms several times and then see what happens: If we come up with different answers—too bad; but if the answer is always the same, if our statement always gets the same value, then we can rest content, since our lack of precision turns out to be immaterial. Different admissible precisifications induce different admissible valuations, none of which can trump the others; but the logical product of such valuations—their supervaluation—is reliable enough.

Now, there are two big questions that need to be answered before we can say we have a full-fledged supervaluational semantics for a language L . First, how exactly is the notion of an admissible *precisification* to be cashed out? Second, how exactly do we cash out the notion of an *admissible* precisification? The second question is notoriously a difficult one. It is difficult in practice (Michael Dummett, 1991, p. 74, says that here comes the 'hard work' when we attend to the semantics of a specific language) as well as in principle (since it gives rise to worries concerning higher-order vagueness). But the first question is also important, since the philosophical plausibility of the basic insight depends crucially on the answer. Just to give an idea, there are at least two main options one may consider:

- (1) One option is to construe a precisification of our vague language, L , as a precise *language* in its own right. (This is how Dummett and David Lewis put it, at least in some of their works.¹) From this point of view, to say that L

¹ See for instance Lewis 1975, p. 188: 'Our convention of language is not exactly a convention of truthfulness and trust in a single language ... Rather it is a convention of truthfulness and trust in whichever we please of some cluster of similar languages: languages with more or less the same sentences ... The convention confines us to the cluster, but leaves us with indeterminacies whenever the languages of the cluster disagree.' Burns 1991 takes this as a starting point for an account of vagueness that is pragmatic, as opposed to semantic, but Lewis's later writings indicate that he was thinking along supervaluational lines.

admits of several precisifications is to say that *L* is really many languages, a cluster of several (homophonic) precise languages whose semantics are only partially in agreement: our practices have simply failed to uniquely identify the one language that we are speaking. Correspondingly, to say that a statement of *L* is super-true (for instance) is to say that it is true no matter how we suppose *L* to be identified, that is, no matter which (homophonic) variant of our statement we consider.

- (2) A different, more popular option is to construe a precisification of a vague language *L* as a precise *interpretation* of *L*. (This is how most authors see it, from Kit Fine to Marian Przełęcki to the later David Lewis to Vann McGee and Brian McLaughlin.²) Here the idea is that the grammar of our language is in principle compatible with countless interpretations, countless models each of which is logically adequate in that each assigns an extension to every predicate constant, a denotation to every individual constant, etc. Our linguistic practices and conventions are meant to select one such interpretation as the intended one, but they may fall short of doing the job properly. Correspondingly, to say that a statement of *L* is super-true (for instance) is to say that it is true no matter how we suppose the job to *be* done properly.

Both of these options (and there are others³) may in turn be further qualified in a number of ways. In particular, each of them can be qualified by further specifying the analytic *link* between the given vague language, *L*, and its precisifications. One may:

- (a) think of *L* as being literally *defined* by its precisifications (as the above formulations suggest), or
- (b) think of *L* as being analytically *prior* to its precisifications, the latter being what we get—or would get—by replacing *L*'s vague words with precise

² See for instance McGee and McLaughlin 1995, p. 228: 'The position we are developing here does not require looking at a lot of different languages, but rather looking at a lot of different models. The models we look at are all models of the vague language whose semantics we are trying to describe'. Compare Fine 1975, p. 125, Przełęcki 1976, pp. 376f, and Lewis 1993, p. 172.

³ For instance, a third option is to construe a precisification as an *assessment* of the given language *L*, that is, as a classification of every atomic *L*-statement as either true or false. (This is how Bas van Fraassen 1966 originally conceived of it, though his concern was with lack of reference rather than vagueness; see also Herzberger 1982.) The idea, in this case, is that the semantics of a language is characteristically identified by the truth-values of its atomic statements: we come to learn the meaning of a word by learning which statements containing that word are correct, i.e., true, and which are incorrect, i.e., false, according to the beliefs of our linguistic community. To the extent that these beliefs may disagree, or fail to cover every case, our language is vague.

ones (option 1) or by sharpening the actual interpretation of those vague words (option 2).⁴

Moreover, each option can be further qualified by allowing for a certain leeway in the *scope* of the relevant precisifications. One may:

- (i) speak of *total* precisifications, that is, precisifications relative to the whole language (as in the above formulations), or
- (ii) speak of *limited* precisifications, that is, precisifications relative only to that portion of the language that shows up in the particular statement or statements that we wish to evaluate.

So there obviously are several distinct ways of spelling out the basic insight on which a semantics of this sort is erected. A supervaluation registers the pattern of agreement among the valuations induced by a certain class of admissible precisifications, but exactly what these precisifications amount to is no straightforward business.

Does it really matter which option we settle on? In a way, one may think that these are distinctions without a difference. What really matters, in the end, is the supervaluation itself, which is just a partial function from statements to truth-values; and so long as we can establish a suitable correspondence among the relevant criteria of admissibility, it is perfectly conceivable that we end up with the same supervaluation in all cases. Indeed, if we confine ourselves to a standard language—that is, a language of the sort considered in classical logical theories—then it is easy to verify that all options yield supervaluations that are, if not identical, equivalent up to isomorphism, at least under certain conditions.⁵ Gener-

⁴ Thus, the passage from Lewis in note 1 is in the spirit of option (1)(a), but Dummett's formulation is in line with (1)(b): 'For every vague predicate, say "red", we may consider the relation which a given predicate, say "rouge", will have to it when "rouge" is what I shall call an *acceptable sharpening* of "red"' (1991, p. 73). Likewise, McGee and McLaughlin's account follows option (2)(a), but there are writers, such as Hans Kamp (1975), who explicitly go for (2)(b).

⁵ To illustrate, consider options (1)(a)(i) and (2)(a)(i). Given a vague language L , it is easy to establish a correspondence (up to isomorphism) between the interpretations of the precise languages that qualify as precisifications of L in the first sense and the precise interpretations of L that qualify as precisifications in the second sense. Suppose for simplicity that L admits of just two precisifications in the first sense, two languages L_1 and L_2 that are perfect duplicates of each other except that the L_1 -interpretation of a certain predicate, F_1 , is slightly different from the L_2 -interpretation of its duplicate, F_2 . Strictly speaking, L_1 and L_2 are distinct languages, but we are supposed to think of them as determining the same vague language L , so we can construe each pair of duplicate symbols as a single L -symbol. Accordingly, we can treat the interpretations of our two languages, $I(L_1)$ and $I(L_2)$, as two interpretations of the same language, $I_1(L)$ and $I_2(L)$, which is exactly what L 's precisifications would amount to in the second sense. Conversely, given

ally speaking, however, this is not enough to conclude that they all boil down to the same thing. There are at least two sorts of consideration that suggest the opposite.

On the one hand, the identification of truth with super-truth has been attacked on several grounds, and depending on how one sees the details, the response on behalf of supervaluationism may look very different. Think, for example, of David Sanford's classic objection (1976, p. 206), emphatically echoed by Jerry Fodor and Ernie LePore (1996): the very idea of explaining the semantics of a vague language L by looking at its admissible precisifications would be wrong-headed. For how could we learn something about a language that is *in fact* vague by examining the semantics of its *possible* precisifications? Surely this objection has a strong appeal if we are thinking in terms of (b)-style precisifications, that is, precisifications construed as precise languages or interpretations that go beyond what we in fact have. But the objection loses its force if we are thinking in terms of (a)-style precisifications, that is, if we are truly identifying L with a cluster of precise languages or interpretations. For in that case, examining a precisification does not amount to examining something *else* than what we in fact have. As McGee and McLaughlin (1999) have pointed out, from this perspective the objection betrays a misconstrual of the idea that admissible interpretations must respect conceptual truths: there is no *a priori* requirement that such interpretations reflect every aspect of a word's meaning, and one may insist that the semantic features of a vague language are global. (See also Morreau 1999.) Moreover, even with respect to (b)-style precisifications, the objection loses its force if we are thinking along option (2) rather than option (1). By replacing L 's vague words with precise ones we may indeed lose track of certain distinguishing features of L : for example, it is a 'conceptual truth' of English that 'small' has borderline cases, and this conceptual truth would seem to be lost in every precise variant of English. However, considering how the vague *interpretation* of those words can be made more precise need not have that effect. One can plausibly maintain that how an expression can be made precise is already part of its meaning, as Fine (1975, p. 131) put it: the meaning of an expression is a product of both its actual meaning (the mean-

two precise interpretations of a single vague language L , $I_1(L)$ and $I_2(L)$, we can obviously split each L -symbol into two duplicates and construct two different languages, L_1 and L_2 , setting $I(L_1) = I_1(L)$ and $I(L_2) = I_2(L)$. So the two options yield isomorphic supervaluations. (The 'certain conditions' mentioned in the text concern the difference between (i)-style and (ii)-style precisifications. These can be shown to be equivalent only if we assume that all words can be simultaneously precisified; this assumption is part and parcel of option (i), but may be relaxed if one follows option (ii), hence the latter option may in principle yield supervaluations that are undefined with respect to statements that the former option treats as super-true or super-false.)

ing fixed by the partial interpretation of L) and its potential meanings (the meaning fixed by the complete extensions of that interpretation). Finally, even with respect to (1)(b)-style precisifications, the force of the objection decreases if we think in terms of option (ii) rather than option (i). Here the implausibility of the supervaluational manoeuvre may seem striking insofar as it is unrealistic to presume that a vague language *as a whole* can be matched up with a precise one: a language L may be necessarily vague in that some of its expressions cannot be precisified, individually or collectively. Indeed, a precise expression E^* cannot qualify as an admissible substitute of a vague expression E of L unless it is in principle possible for any two speakers of L to shift their standards of correctness so as to accord with the rules for the proper application of E^* ; it must in principle be possible, in other words, for any two speakers to decide to speak the language L^* in which E^* replaces E , and for it to be common knowledge that this shift has taken place. As there is no guarantee that every vague expression admits of replacements that meet these conditions,⁶ there is no reason to suppose that L admits of total precisifications in the sense of option (1)(b)(i). Yet this is not to say that we cannot learn anything about the semantics of L by considering its possible (1)(b)(ii)-precisifications in those cases where the above conditions *are* met. We may not want to replace vague expressions by precise ones, but the fact that we could—and the extent to which we could—is arguably a fact about our language that may contribute to explain the truth-conditions of our statements. (To put it differently, Fodor and LePore worry about strict identity conditions for linguistic expressions, but one could argue that (1)(b)-style precisifications are rather to be thought of as Lewisian *counterparts*.⁷ And while it may be implausible to suppose that *all* vague expressions can be matched up with admissible precise counterparts, it is a fact that some can.)

On the other hand, even the formal equivalence between the various options might break down as soon as we consider languages that are richer than standard languages. Consider, for instance the result of adding an operator corresponding to the English phrase ‘It is definitely the case that’, abbreviated as ‘D’—a very natural thing to have in a vague language. If we construe precisifications along option (a), that is, if we think of L as being literally defined by its precisifications, then it is customary to treat D in analogy with the modal operator for *necessity*: assuming a relation of accessibility to be defined on the space of all given precisi-

⁶ John Collins and I (2000) have argued that certain rationality predicates, such as ‘rationally obliged to take the money on the table in a game of take-it-or-leave-it’, are a case in point.

⁷ In this sense, the worry parallels Kripke’s ‘Humphrey objection’ to counterpart theory (1972, p. 45, n. 13), and the (1)(b)-supervaluationist’s reply can mimic Lewis’s (1986, p. 196).

fications, a statement of the form $D\phi$ ⁸ will be evaluated as true on a precisification P if and only if ϕ is true on all those precisifications that are accessible from P . Accordingly, the logic of D will depend on the conditions imposed on the accessibility relation. Since the minimum requirement is that it be reflexive, the minimal logic for D will correspond to the modal logic for \Box known as KT (modulo certain concerns about the entailment relation to be discussed shortly).⁹ By contrast, if we construe precisifications along option (b), that is, if we think of L as being analytically prior to its precisifications, then there is more flexibility. We can still treat D in analogy to the necessity operator; but we may also treat it in analogy with the *actuality* operator, for we may want to say that the truth-value of $D\phi$ on a precisification is determined by the actual truth-conditions of ϕ , which is to say by the truth-conditions of ϕ as initially determined by our vague semantic conventions. (As far as ordinary connectives and quantifiers are concerned, such conventions may be modelled by some partial truth-value semantics, e.g. in accordance with the weak/strong truth-conditions of Kleene 1952.) The intuition would be that statements of the form $D\phi$ are not necessarily made more precise through making ϕ more precise. If ϕ suffers from first-order vagueness, then $D\phi$ is, in a way, already perfectly precise—it is false. And if ϕ suffers from $(n+1)$ -th order vagueness, then $D\phi$ will only be n -th-order vague. Thus, on this view $D\phi$ would be true on a precisification if and only if ϕ is ‘already true’ before we embark in the precisification business ($D\phi$ itself qualifying as already true if and only if so is ϕ). And the resulting logic for D would be stronger than KT : it would be at least as strong as the modal logic known as $S5$ (i.e., $KT5$).¹⁰

⁸ To simplify notation, I shall freely treat symbols as names of themselves, using concatenation to indicate the concatenation of various symbols. For example, if ϕ is any formula, I shall write ‘ $D\phi$ ’ for the result of concatenating ‘ D ’ and ϕ , setting $D\phi = \lceil D\phi \rceil$.

⁹ See Williamson 1994, Sect. 5.6. Strictly speaking, there are two different options here. One is described in the text, where the accessibility relation is somehow imposed upon a given space of precisifications. The other is to identify accessibility with admissibility, in the following sense. Every language comes with a set of precisifications, corresponding to the various ways in which first-order vagueness can be resolved. Each precisification, in turn, comes with a set of admissible alternative precisifications, all of which may also come with sets of alternative precisifications, and so on. Super-truth is truth on all initial precisifications; definite truth at a precisification P is truth at all precisifications admissible from the point of view of P . The two options yield different logics. In particular, unless admissibility is required to be transitive, on this alternative strategy the super-truth of a statement ϕ would not entail the super-truth of $D\phi$, while the entailment holds on the approach described in the text.

¹⁰ Again, strictly speaking there is room for other options here. For instance, Fine (1975, pp. 141–143) equates being ‘already true’ with being true at the ‘base specification point’, which is to say super-true. This is still in the spirit of an actuality-like construal of D , though the outcome is obviously different: we still get $S5$, but on this account the super-truth of ϕ entails that of $D\phi$.

2. Validity: Global, Local, and Collective

So much for the building blocks of supervaluationism. The basic insight is clear enough, but its implementation is no straightforward business and a lot depends on matters of detail. I now want to consider more closely what happens when we proceed to the task of explaining the supervaluational *logic* of a vague language, that is, not the notion of truth as applied to individual statements, but the notion of validity as applied to whole arguments. For the sake of generality, and not to beg any questions, it pays to work within the broadest possible setting, allowing for multiple-conclusion patterns of reasoning. Thus, by an argument I mean quite generally a set Σ of premises followed by a set Γ of conclusions, and to say that an argument $\Sigma \vdash \Gamma$ is valid is to say that the premises in Σ jointly entail at least one conclusion in Γ .¹¹ What exactly this means, and on what conditions the entailment obtains (hence, what logic we get), are the two questions I wish to address.

It is important to begin with the first question. No matter how we cash out the idea of a precisification, it is obvious that supervaluations need not be bivalent: perhaps every statement *can* be super-true (T) or super-false (F), but some statements may in fact be neither—they may be indeterminate (I). It follows that supervaluationally we cannot identify being T with not being F, or being F with not being T, hence the standard notion of argument validity does not automatically carry over to a supervaluational scenario. Standardly, one says that an argument is valid if and only if it is truth preserving: whenever all the premises are true, one of the conclusions must be true. One also says that an argument is valid if and only if it is not possible for all the conclusions to be false when the premises are all true. In the presence of bivalence, the two characterizations are equivalent.¹² Indeed, there are four equivalent ways of cashing out the same intuition:

while the entailment may fail on the approach described in the text. Moreover, on the account in the text the super-truth of $D(\phi \vee \psi)$ entails that of $D\phi \vee D\psi$, while on Fine's account it does not. (Fine says the latter entailment is unacceptable, which it really is *if* 'definitely' is to express, in the material mode, what 'super-true' expresses in the formal mode.)

¹¹ This general setting is especially important if one is interested in dualizing the analysis so as to apply it to what I have called 'subvaluationism'—the view according to which a statement is true/false if and only if it is true/false on *some* admissible precisification (Varzi 1997, 1999, 2000). An application of subvaluational semantics to vagueness is outlined and defended in Hyde 1997.

¹² This is not to say that they express the same conception of validity. For instance, often one explains the rationale behind these characterizations in terms of commitments, or warranties, and there is no obvious equivalence between being committed to *accept* (or being warranted in *asserting*) a conclusion and being committed to *reject* (or being warranted in *denying*) a premise. As my focus here is mostly on the formal semantic features of the entailment relation, I will ignore such concerns, as I will ignore any worries that might be raised on such grounds against the notion of a multiple-conclusion argument (referring to Restall 2005 for discussion).

- (A) An argument is valid iff, necessarily, if every premise is T, then some conclusion is T.
- (B) An argument is valid iff, necessarily, if every conclusion is F, then some premise is F.
- (C) An argument is valid iff, necessarily, if every premise is T, then some conclusion is not F.
- (D) An argument is valid iff, necessarily, if every conclusion is not T, then some premise is F.

In the absence of bivalence, however, there is no guarantee that the equivalence is preserved. In particular, on the most natural supervaluational construal, according to which truth/falsity is super-truth/falsity, the above conditions are all distinct:

- (A) An argument is valid iff, necessarily, if every premise is: T on all precisifications, then some conclusion is: T on all precisifications.
- (B) An argument is valid iff, necessarily, if every conclusion is: F on all precisifications, then some premise is: F on all precisifications.
- (C) An argument is valid iff, necessarily, if every premise is: T on all precisifications, then some conclusion is not: F on all precisifications.
- (D) An argument is valid iff, necessarily, if every conclusion is not: T on all precisification, then some premise is: F on all precisifications.

(From now on, to simplify terminology I shall generally speak of precisifications meaning *admissible* precisifications.)

To see that these four conditions are pairwise distinct, it is sufficient to consider the following two argument forms:

- [1] $\phi, \neg\phi \vdash \phi \wedge \neg\phi$
- [2] $\phi \vee \neg\phi \vdash \phi, \neg\phi.$

Inspection shows that [1] is valid according to conditions (A) and (C), though not according to (B) and (D) (just let the value of ϕ be indeterminate). Similarly, [2] is valid according to conditions (B) and (C) but not according to (A) and (D). If we only considered single-conclusion arguments, then it's easy to verify that A- and C-validity would coincide, as would B- and D-validity—but only if the language does not contain the D operator. Otherwise we can still test the pairwise non-equivalence of all four conditions by considering the following:

- [3] $\phi \vdash D\phi$
- [4] $\neg D\phi \vdash \neg\phi.$

Again, inspection shows that [3] is only A- and C-valid, whereas [4] is only B- and C-valid, the counterexamples arising once again when ϕ is indeterminate. (To be more precise, here and below I am assuming that D is treated in accordance with the first policy mentioned at the end of section 1, that is, as an operator analogous to the modal necessity operator.¹³ If D is interpreted according to the alternative policy, in analogy with the actuality operator, then [3] would be neither A-valid nor C-valid. For example, if x is a borderline case of F , then ' $Fx \vee \neg Fx$ ' fails to be 'already true' on the partial interpretation of the language, at least if we rely on a partial semantics *à la* Kleene. Hence ' $D(Fx \vee \neg Fx)$ ' is not super-true although ' $Fx \vee \neg Fx$ ' is. Likewise, [4] would be neither B- nor C-valid. We can already see here that the details of the basic framework can make a difference in the overall logic of the language.)

So supervaluationism allows for a multiplicity of entailment relations, that is, notions of validity. In fact, these are not the only options, either, for a supervaluational perspective allows for different ways of understanding the relationship between the premises and the conclusions of a valid argument. Conditions (A)–(D) would be the only options if we blindly imported the standard conditions, taking T to be super-truth and F to be super-falsity¹⁴; but a supervaluationist might want to exploit a different intuition. She might want to say that just as questions of truth may only be answered upon considering the precisifications of the language, so questions of validity may be answered only upon considering those precisifications. Just as a statement is rated true, supervaluationally, if and only if it

¹³ Modulo the qualification at note 9.

¹⁴ To be sure, there are additional possibilities. For one thing, in the absence of bivalence it is natural to consider double-barrelled notions of validity (Scott 1975). Combining (A) and (B), for instance, one might require both transmission of (super-)truth from the premises to at least one conclusion *and* re-transmission of (super-)falsity from all conclusions to at least one premise (see e.g. Kremer and Kremer 2003). Since the results presented below can easily be extended to such notions, I will not examine them explicitly. Secondly, one might consider variants of (A)–(D) obtained by contraposition. For instance, the contrapositive of (A) would read: (A') An argument is valid iff, necessarily, if every conclusion is not: T on all precisifications, then some premise is not: T on all precisifications. Ordinarily, contraposition is a logically invariant operation, so in a way (A') reduces to (A). However, just as there are many notions of entailment, so there are many notions of equivalence (understood as two-way entailment). In particular, we shall see below that in a vague language with a supervaluational semantics contraposition may fail to be A-valid, which is to say that a statement and its contrapositive may fail to be A-equivalent. To the extent that the notion of an 'admissible precisification' is vague, the metalanguage in which the semantics is formulated is itself vague, hence the (metalinguistic) A-equivalence between (A) and (A') cannot be proved by mere appeal to (metalinguistic) contraposition. (Thanks to Patrick Greenough for raising this point.) Nonetheless, I fail to see any counterexamples, so in the following I will ignore (A') and focus exclusively on (A). Ditto for the contrapositives of (B)–(D).

is true on all admissible precisifications, so an argument may be rated valid if and only if, necessarily, its premises and conclusions stand in the appropriate relation on all admissible precisifications. Formally, this amounts to a different way of fixing the scope of the relevant quantification over precisifications, corresponding to the following variants of (A)–(D):

- (α) An argument is valid iff, necessarily, on all precisifications: if every premise is T, then some conclusion is T.
- (β) An argument is valid iff, necessarily, on all precisifications: if every conclusion is F, then some premise is F.
- (γ) An argument is valid iff, necessarily, on all precisifications: if every premise is T, then some conclusion is not F.
- (δ) An argument is valid iff, necessarily, on all precisifications: if every conclusion is not T, then some premise is F.

Tim Williamson and others have objected that these variants of (A)–(D) would betray a disloyalty to supervaluationism, since here super-truth plays no role in the definitia.¹⁵ That strikes me as unfair. For one thing, when we are dealing with a vague language, it seems perfectly reasonable to suppose that we may want to reason from premises that lack a definite truth-value, in which case super-truth cannot be our guidance. Indeed, one might suggest that it is precisely by reasoning according to (α)–(δ) that a supervaluationist finds it natural to accept so-called principles of penumbral connection: ‘Look, I’m not sure what ‘small’ exactly means, so I am not sure whether x is truly small. But I certainly know this: Assuming x is small, since y ’s height is less than x ’s, y must be small, too.’ Moreover, the intuitive rationale for these conditions may vary significantly according to how we construe precisifications. If we construe them according to option (1), specifically in its (a)(i)-variant, the intuition behind (α)–(δ) seems straightforward precisely insofar as truth is identified with super-truth: if our language is truly a cluster of totally precise languages, then it is natural to think that we should check the status of our arguments by checking their status in each language in the cluster (and for each logically possible way of defining the cluster). To put it differently, to the extent that supervaluationism construes vagueness as ambiguity on a grand scale, as Kit Fine originally put it (1975, p. 136), type-(1) precisifications are like disambiguations, so to assess the validity of an argument amounts to checking whether the argument is valid no matter how we systematically disambiguate its premises and conclusions. By contrast, if we construe pre-

¹⁵ See Williamson 1994, p. 148. Rosanna Keefe (2000, p. 174, n. 10) agrees.

cisifications according to option (2), again on its (a)(i)-variant, the intuition is different. On this construal, the total precisifications admitted by our language are akin to the possible worlds countenanced in the semantics of modal logic: we interpret a vague language by means of a cluster of classical models just as we interpret a modal language by means of a cluster of possible worlds. So when it comes to argument validity, the analogy delivers exactly the account under examination: conditions (α) – (δ) match the four conditions that may be considered in modal logic, with ‘precisification’ in place of ‘possible world’. (This becomes particularly attractive if we think that vagueness is, in fact, a modal phenomenon, a phenomenon that induces a ‘mode of truth’ not reducible to assertoric truth, as Josh Dever *et al.* 2004, have recently argued.) Neither rationale would, I think, be equally appealing if we worked with precisifications of type (b) or (ii), so here Williamson’s misgivings may be warranted. Yet this may be debatable, too. For example, working with precisifications of type (2)(b), Fine opted for an A-style definition of argument validity (1975, p. 136), but Dummett opted for an α -style definition (1975, p. 108).

Be that as it may, there is no question that (α) – (δ) suggest themselves as obvious alternatives to (A)–(D). In fact, we may just focus on (α) , since (β) , (γ) , and (δ) are trivially equivalent. This follows from the fact that all precisifications are bivalent, which means that on all precisifications being T coincides with not being F and being F with not being T.¹⁶ Nonetheless, inspection shows that this new sense of argument validity is indeed logically distinct from the four senses defined in (A)–(D). As it turns out, if we confine ourselves to D-free, single-conclusion patterns, an argument is bound to be α -valid if and only if it is also valid in the sense of conditions (A) and (C), but it may fail according to (B) and (D) (consider [1]).¹⁷ If we allow for multiple-conclusion patterns, some α -valid arguments may also fail according to condition (A) (consider [2]). And in the presence of the D-operator, there are arguments that are not α -valid in spite of being C-valid (consider [3] and [4]). So α -validity is generally different from validity in any of the other four senses. Adapting Williamson’s terminology, we may say that conditions (A)–(D) afford *global* notions of validity, whereas (α) affords a *local* notion. In the same spirit, Stewart Shapiro (2006, Ch. 4) speaks of *external* and *internal* validity, respectively: the former, but not the latter, requires that we take into account the external factors that influence our way of determining the actual truth-conditions of our statements.

¹⁶ Double-barrelled variants of (α) – (δ) (see note 14) will similarly collapse to (α) .

¹⁷ The proof of the equivalence between α -validity and A-validity, in D-free contexts, can be gathered from Shapiro 2006, Ch. 4, Theorems 13 and 15.

We may, in addition, consider the following variants, which reflect a third, different way of collecting the quantification over precisifications:

- (X) An argument is valid iff, necessarily, if on all precisifications every premise is T, then on all precisifications some conclusion is T.
- (Y) An argument is valid iff, necessarily, if on all precisifications every conclusion is F, then on all precisifications some premise is F.
- (Z) An argument is valid iff, necessarily, if on all precisifications every premise is T, then on all precisifications some conclusion is not F.
- (W) An argument is valid iff, necessarily, if on all precisifications every conclusion is not T, then on all precisifications some premise is F.

Here we may quickly note that (Z) is equivalent to (X), since on all precisifications being T coincides with not being F, and (W) is equivalent to (Y), since on all precisifications being F coincides with not being T. However, conditions (X) and (Y) are distinct, since only (X) validates [3] and only (Y) validates [4], and both conditions are distinct from any of the other conditions considered so far: both (X) and (Y) validate [1] (thus differing from (B) and (D)) and [2] (thus differing from (A)), but neither validates both [3] and [4] (thus differing from (C)) and both validate either [3] or [4] (thus differing from (α)).

The rationale for these two additional notions of validity might appear artificial, but it isn't. In both cases, it reflects the intuition that a valid argument is one in which the *conjunction* of the premises is related in the appropriate way to the *disjunction* of the conclusions. In classical logic, this intuition is perfectly captured by the standard definitions considered at the beginning, since that logic is truth-functional. But supervaluationism is not truth-functional; in particular, there is a difference between super-falsifying a conjunction and super-falsifying at least one conjunct, just as there is a difference between super-verifying a disjunction and super-verifying at least one disjunct. That is precisely why [1] may fail to be B- or D-valid, while [2] may fail to be A- or D-valid, respectively.¹⁸ This feature of supervaluationism may be controversial, and to some critics that is already enough to look elsewhere for a good semantics of vagueness. (That's the famous objection from upper-case letters, as Jamie Tappenden calls it: 'You say that "either ϕ or ψ " is true, so EITHER ϕ OR ψ [stamp the foot, bang the table] must be true', 1993, p. 564.) But never mind that; every supervaluationist must come to

¹⁸ Likewise, super-falsifying a universal generalization differs from super-falsifying one of its instances, and super-verifying an existential generalization differs from super-verifying one of its instances. This is why, when it comes to (A)–(D), the logical status of [1] and [2] is inherited by arguments involving quantifiers—whence the supervaluational way out of the sorites paradox.

terms with this feature of their semantics—take it or leave it. What is relevant, from the present perspective, is that precisely because of this feature there are two ways of understanding the intuition behind the standard definitions of validity, depending on whether we understand the relevant quantifications over premises and conclusions collectively (‘in the same breath’) or distributively. Global and local notions of validity reflect a distributive reading, for they all require that each premise and conclusion be evaluated in its own terms. Conditions (X) and (Y), by contrast, reflect a collective reading: to consider whether all precisifications verify every premise, or falsify some premise, is to consider whether they verify or falsify the relevant (possibly infinitary) conjunction, that is, whether such a conjunction is super-true or super-false, respectively; and to consider whether all precisifications verify some conclusion, or falsify every conclusion, is to consider whether they verify or falsify the relevant (possibly infinitary) disjunction. I can see why such conditions may not be *prima facie* appealing in the absence of truth-functionality. But they are legitimate conditions to consider, and it is a fact that some theories that broadly qualify as supervaluational (e.g. Rescher and Brandom 1980, Sect. 5) are built around such a *collective* notion of argument-validity.

To recapitulate, then, a supervaluational semantics makes room for at least seven distinct notions of argument validity: four global, one local, and two collective notions. Writing ‘ $\Sigma \models_i \Gamma$ ’ to indicate that the argument $\Sigma \vdash \Gamma$ is *i*-valid, that is, valid according to condition (*i*), we can summarize the picture as follows:

- (A) $\Sigma \models_A \Gamma =_{df}$ Necessarily, if every $\phi \in \Sigma$ is: T on all precisifications, then some $\psi \in \Gamma$ is: T on all precisifications.
- (B) $\Sigma \models_B \Gamma =_{df}$ Necessarily, if every $\psi \in \Gamma$ is: F on all precisifications, then some $\phi \in \Sigma$ is: F on all precisifications.
- (C) $\Sigma \models_C \Gamma =_{df}$ Necessarily, if every $\phi \in \Sigma$ is: T on all precisifications, then some $\psi \in \Gamma$ is not: F on all precisifications.
- (D) $\Sigma \models_D \Gamma =_{df}$ Necessarily, if every $\psi \in \Gamma$ is not: T on all precisifications, then some $\phi \in \Sigma$ is: F on all precisifications.
- (α) $\Sigma \models_\alpha \Gamma =_{df}$ Necessarily, on all precisifications: if every $\phi \in \Sigma$ is T, then some $\psi \in \Gamma$ is T.
- (X) $\Sigma \models_X \Gamma =_{df}$ Necessarily, if on all precisifications every $\phi \in \Sigma$ is T, then on all precisifications some $\psi \in \Gamma$ is T.
- (Y) $\Sigma \models_Y \Gamma =_{df}$ Necessarily, if on all precisifications every $\psi \in \Gamma$ is F, then on all precisifications some $\phi \in \Sigma$ is F.

It would of course be nice to complete the picture with some account of the relative strengths of these entailment relations, but the account is rather intricate as things change significantly depending on whether Σ and Γ contain several, one, or

zero elements, and on whether and how the D operator is admitted into the language. We have already seen an example of this intricacy in discussing the relationships between the local and global senses of validity. (Besides, a systematic comparison would call for a full formal treatment, so as to attach a precise meaning to the locution ‘necessarily’ that appears in the definiens: intuitively, the locution means ‘in every logically possible situation’, but of course this may signify different things depending on the details of the overall semantic machinery.) The only general relationships that can be asserted with no qualification are that A-validity implies C-validity and D-validity implies B-validity, whereas α -validity implies both X-validity and Y-validity (since the universal quantifier ‘on all precisifications’ distributes over the ‘if... then’ conditional). Moreover, all seven entailment relations coincide in the two limit cases: when Σ is empty and Γ is a singleton, and when Σ is a singleton and Γ is empty. For in those cases all conditions amount to the same thing: the argument $\Sigma \vdash \Gamma$ is valid if and only if the unique element of Γ is necessarily true on all precisifications, or if and only if the unique element of Σ is necessarily false on all precisifications, respectively. Thus, logical truth and logical falsity do not depend on the particular notion of validity that one considers. All other cases, however, require careful examination.

3. Comparisons

What is the best notion of validity from a supervenience perspective? Or: is there a best notion? To address questions such as these, I want to take a look at how the options behave *vis-à-vis* a number of worries that have been voiced against the sort of logic that emerges from supervenience.

One immediate consequence of the last remark of the previous section is that all seven notions of validity coincide with the classical notion when it comes to identifying logical truths and logical falsities, at least if we confine ourselves to supervenience semantics based on type-(i) (i.e., total) precisifications. For, on the one hand, if a statement ψ is necessarily true on all such precisifications, then ψ is true on all precise models of the language, hence logically true in the sense of classical logic. On the other hand, if ψ is not necessarily true on all precisifications, then there must be a model such that ψ is false on some relevant precisifications, which implies that ψ must be false on some precise models and cannot, therefore, qualify as a classical logical truth. Similarly for logical falsity. This is a well-known result, and in one form or other it has fuelled the best-selling claim of supervenience: you can stick to classical logic even in the presence of vagueness. (Why can this claim not be extended to semantics based on type-(ii), partial precisifications? Because one motivation for such semantics is to allow for the

possibility that some expressions be unsharpenable, and a statement involving unsharpenable expressions will be indeterminate even if it is an instance of a classical logical truth/falsity. In fact, in the absence of formally *ad hoc* constraints, it may well turn out that such semantics deliver a notion of logical truth/falsity that is not even recursively axiomatizable.¹⁹ From now on, however, I shall for simplicity ignore such semantics.)

One thing is logical truth, though, and quite another is logical validity broadly understood. And it is precisely here that one begins to worry. Just how classical is the logic delivered by supervaluational semantics? And where it goes non-classical, just how adequate is it to dealing with the phenomenon of vagueness? Let me go through this sort of worry by briefly considering three objections that have attracted a great deal of attention in the recent literature. I shall phrase the objections in general terms, as if there were just one notion of argument validity available to supervaluationism, and then I shall try to disentangle the picture by examining how the objections persist or dissolve depending on which specific notion one considers. For the sake of precision, I shall also assume that the D operator is always handled in accordance with the first policy considered earlier, namely, as an operator akin to the necessity operator axiomatized by a modal logic at least as strong as KT. This is fair enough, since this policy is compatible with all supervaluational accounts that we have been considering and is, in fact, a favorite option in the literature. Later we shall see whether treating D as an actuality operator can make a difference.

Objection 1. Supervaluationism may well deliver a classical notion of logical truth, or even a classical notion of entailment relative to single-conclusion arguments. But as soon as we look at the large picture, we find multi-conclusion argument forms that are classically valid and yet may fail in a vague supervaluational language. Argument [2] above is a case in point. Even disregarding the objection from upper-case letters, there are many other instances of the same phenomenon—for example:

- [5] $\phi \vdash \psi, \neg(\phi \rightarrow \psi)$
 [6] $\forall xFx \vdash \forall xGx, \neg \forall x(Fx \rightarrow Gx).$

[Proof.²⁰ For [5], let ϕ be T and ψ be I. For [6], suppose everything in the domain is F , whereas some things are G and the rest is borderline G .]

¹⁹ A negative result of this sort is known to hold for supervaluational treatments of non-denoting singular terms; see Bencivenga 1978.

²⁰ I write ‘proof’ meaning ‘purported proof’. As will be obvious shortly, the proof only goes through on some understandings of ‘valid argument’.

Objection 2. Even with respect to single-conclusion arguments, the claim that supervaluationism preserves classical logic is only true on a narrow conception of ‘logic’. For example, as Williamson (1994, pp. 151f) pointed out, the following *rules* of inference are classically valid, yet they may fail in a vague language with a supervaluational semantics:²¹

- | | | |
|------|--|--------------------------|
| [7] | From $\Sigma, \phi \models \psi$ infer $\Sigma \models \phi \rightarrow \psi$ | <i>Conditional proof</i> |
| [8] | From $\Sigma, \phi \models \psi$ infer $\Sigma, \neg\psi \models \neg\phi$ | <i>Contraposition</i> |
| [9] | From $\Sigma, \phi \models \psi \wedge \neg\psi$ infer $\Sigma \models \neg\phi$ | <i>Indirect proof</i> |
| [10] | From $\Sigma, \phi \models \sigma$ and $\Sigma, \psi \models \sigma$ infer $\Sigma, \phi \vee \psi \models \sigma$ | <i>Proof by cases</i> |

[Proof: For [7] and [8], let ‘ ψ ’ be ‘ $D\phi$ ’; for [9], let ‘ ϕ ’ be ‘ $\sigma \wedge \neg D\sigma$ ’ and ‘ ψ ’ be ‘ $D\sigma$ ’; for [10], let ‘ ψ ’ be ‘ $\neg\phi$ ’, and ‘ σ ’ be ‘ $D\phi \vee D\neg\phi$ ’.]

Objection 3. The supervaluational account of the D operator is inconsistent with unrestricted higher-order vagueness. For, on the one hand, if unrestricted higher-order vagueness is admitted, then the relation of accessibility among precisifications must be such as to verify the following entailment whenever x_j and x_{j+1} are adjacent elements of a sorites series:

$$[11] \quad DD^n Fx_j \models \neg D\neg D^n Fx_{j+1} \quad \text{D-gap}$$

(‘ D^n ’ stands for n repetitions of ‘D’, $n \geq 0$). On the other hand, super-truth entails definite truth:

$$[12] \quad \phi \models D\phi. \quad \text{D-introduction}$$

Yet [11] and [12] are logically inconsistent. As Crispin Wright (1987, p. 233) and Delia Graff (2003, p. 201) have shown,²² given a sorites series of n objects $x_1 \dots x_n$ such that ‘ Fx_1 ’ is super-true and ‘ Fx_n ’ is super-false, those principles jointly imply the contradiction:

$$[13] \quad D^n Fx_1 \wedge \neg D^n Fx_1.$$

[Proof: The first conjunct follows from Fx_1 by n applications of [12]; the second follows from $\neg Fx_n$ by repeated applications of [12] and [11].]

Now, there are several things that supervaluationists have said (or could say) in response to these objections, but that is not my main concern here. For the re-

²¹ The failure of [8] and [9] is already noted in Fine 1975 and Machina 1976, respectively.

²² The proof mentioned here is Graff’s, and differs from Wright’s in a significant way to which I shall briefly return below. Strictly speaking, Graff does not rely on [12] but on the rule: From $\Sigma \models \phi$ infer $\Sigma \models D\phi$, from which [12] follows (since $\phi \models \phi$).

cord, I don't think faithfulness to classical logic is such a big deal. It's not that supervaluationism has been put forward as a semantics for vagueness that retains classical logic *holus bolus*. It has been put forward as a semantics for vagueness in its own right, one that reflects a certain understanding of what vagueness is and of how the truth conditions of statements involving vague words can be specified without abandoning the *terra firma* of our standard semantic algorithms. As Fine put it, there is but one rule linking super-truth to classical truth, so the truth conditions are, if not classical, 'classical at a remove' (1975, p. 132). If it turns out that supervaluational semantics yields classical logic—good; for classical logic is a nice thing in spite of the fact that it has been developed on the Fregean assumption that precision is a *sine qua non* condition (see Frege 1903, §56). If it turns out that the logic is not fully classical—so be it; after all, classical logic *has* been developed under the Fregean assumption. In short, I agree with Stewart Shapiro (2006, Sect. 4.5): We should first determine how vague expressions function, and figure out the logic from there.

Anyway, this is not my main concern here. My main concern is whether and to what extent the classicality issue depends on the notion of validity one considers. For although there is but one rule linking super-truth to classical truth, at least relative to any particular way of spelling out the details of the machinery, there are several notions of validity that suggest themselves, all of which have equal claim to being a natural extension of our classical understanding of this notion. Perhaps here is where supervaluationism makes room for battles of intuitions. Or perhaps this is just a sign of the fact that once bivalence is abandoned, validity ceases to be an all-or-nothing affair. In any event, it is obvious that the charge of non-classicality warrants further investigation in the light of the multiplicity of meanings that we can attach to the notion of a valid argument. Moreover, some of the above-mentioned objections do not concern the classicality of supervaluational logic but rather its independent adequacy *vis-à-vis* the phenomenon of vagueness. Supervaluationists have dealt extensively with some basic misgivings in this regard, such as the objection from upper-case letters and its relevance to the sorites paradox.²³ They have also explained how their semantics can make room for higher-order vagueness: just as our beliefs and linguistic practices do not succeed in fixing a unique language, or a unique interpretation of the language, they may not succeed in fixing a unique cluster of languages, or a unique cluster of interpretations, which means that the notion of an admissible precisification may itself be vague. This is enough to say that the framework allows for the pos-

²³ See, for instance, McGee and McLaughlin 1995, pp. 207ff, and Keefe 2000a, Sect. 7.5. My own views on this objection may be found in Varzi 2003a, Sect. 2, and 2004, Sect. 4.

sibility that a statement be indeterminately indeterminate, or indeterminately indeterminately indeterminate, and so on.²⁴ However, objection 3 says that higher-order vagueness leads to serious problems independently of this general sort of consideration, at least insofar as it can be represented in the object language by means of the D operator, so the worry cannot be dismissed so easily. In fact, it is clear that the issue of higher-order vagueness ties in directly with the question of what is the logic delivered by a supervaluational semantics. For otherwise we could just rely on the vagueness of our metalanguage (which affects our pre-analytic notion of truth just as it affects the notion of an admissible precisification) and let vagueness be taken care of by the Tarski biconditional:

[14] ‘ ϕ ’ is true if and only if ϕ .

The vagueness of the object language would be reflected in the vagueness of the truth-predicate, and that would be it. This is what Rosanna Keefe calls the ‘Simple Theory’ of vagueness (2000a, pp. 205f). And the reason why this theory won’t do is precisely that we are left with the task of explaining the logic of the ‘if and only if’ connective in the metalanguage. If its logic were fully classical, then we could run Williamson’s argument (1994, pp. 187–189) and conclude that the excluded middle entails bivalence:

[15] ‘ $\phi \vee \neg\phi$ ’ is true if and only if ‘ ϕ ’ is true or ‘ $\neg\phi$ ’ is true.

Thus, to the extent that vagueness involves semantic indeterminacy, either the object language or the metalanguage *must* fail to obey to classical logic. Let me therefore go through each objection in turn.

Consider *Objection 1*, to the effect that multi-conclusion arguments may not retain their logical status in supervaluational logic. We have already seen that although [1] is not valid according to conditions (B) and (D), it is valid according to (A) and (C). Likewise, we have seen that [2] is invalid according to conditions (A) and (D) but valid according to (B) and (C)—and inspection shows that the same holds of [5] and [6]. Moreover, inspection shows that all such arguments are valid according to condition (α), hence according to (X) and (Y). Thus, there are four notions of super-validity according to which these arguments *are* valid—global C-validity, local α -validity, and collective X- and Y-validity—and it is easy to verify that this applies to every argument that is rated valid in classical logic. So obviously the charge depends crucially on what ‘valid’ means in a supervaluational context.

²⁴ This popular line of reaction is fully spelled out in Keefe 2000a, Sect. 8.1.

With *Objection 2* the picture is more intricate. On the one hand, one can again check that all rules of inference retain their classical status if ‘ \models ’ is given a local reading,²⁵ while they may fail on the global reading corresponding to condition (A): the rules do not preserve super-truth. On the other hand, this time the global readings corresponding to conditions (B) and (D) are unaffected by the objection, since both verify all the rules, while the global reading corresponding to condition (C) *is* affected by the objection, since it verifies [7]–[9] but not [10]. Furthermore, although the rules are verified by one of the collective readings, corresponding to condition (Y), each of [7]–[10] may fail on the other reading, corresponding to condition (X). In short, the picture is this: conditional proof, contraposition, and indirect proof fail to be *i*-valid if and only if $i = A, X$; proof by cases fails if and only if $i = A, C, X$. This means that there are again four notions of supervaluational validity that behave classically—global B- and D-validity, local α -validity, collective Y-validity—though these notions do not quite coincide with the ones that resist *Objection 1*. Putting the two objections together, one could therefore conclude that there are still *two* senses of validity on which supervaluationists can claim full faithfulness to classical logic: local α -validity and collective Y-validity. Yet this may begin to sound as though the number of options is getting slim and unattractive: no notion of global validity can do the job.

So let us finally consider *Objection 3*, to the effect that supervaluational accounts of the D operator do not allow for unrestricted higher-order vagueness due to the inconsistency of [11] (D-gap) and [12] (D-introduction). The proof I have offered of this claim is due to Delia Graff, and is significantly stronger than Crispin Wright’s original proof (which makes use of D-introduction within subproofs²⁶). Nonetheless, it should now be obvious that the proof goes through only on some construals. For one thing, although D-gap may be accepted across the board,²⁷ we have already seen that D-introduction is not *i*-valid for $i = B, D, \alpha, Y$, so on those readings of ‘ \models ’ the objection does not get off the ground. Second, the proof depends on repeated applications of both D-gap and D-introduction, so it relies on the following classical rule of inference:

[16] From $\Sigma, \phi \models \psi$ and $\Sigma, \psi \models \sigma$ infer $\Sigma, \phi \models \sigma$. *Transitivity*

As it turns out, this rule is generally valid, but it fails on the C-reading of ‘ \models ’ (let ‘ ϕ ’ be ‘ $\neg D\gamma$ ’, ‘ ψ ’ be ‘ $\neg\gamma$ ’, and ‘ σ ’ be ‘ $D\neg\gamma$ ’, and suppose γ is indeterminate).

²⁵ This is emphasized in McGee and McLaughlin 1998, though in a slightly different jargon.

²⁶ Edgington (1993) and Heck (1993) have argued that this is illegitimate.

²⁷ This is not to say that D-gap is *i*-valid for all *i*; only that it can be *i*-valid as a result of imposing suitable conditions on the accessibility relation.

Thus, the proof does not go through on this reading, either. We are therefore left with only two readings to which the objection applies, namely when ‘ \models ’ stands for A-validity or for X-validity. In other words, all the proof establishes is the following claim, for $i = A, X$:²⁸

$$[17] \quad Fx_1, \neg Fx_n \models_i D^n Fx_1 \wedge \neg D^n Fx_1.$$

Of course, since neither value of i verifies the rule of indirect proof [9], this result is not enough to reach either of the following conclusions, each of which would amount to a sorites paradox:

$$[18] \quad Fx_1 \models_i Fx_n$$

$$[19] \quad \neg Fx_n \models_i \neg Fx_1,$$

However, this is of little consolation since [17] does imply that ‘ Fx_1 ’ and ‘ $\neg Fx_n$ ’ cannot be both super-true, contrary to the facts, and that is enough to conclude that A-validity and X-validity are inconsistent with higher-order vagueness.

In conclusion, the moral we can draw from the three objections is that A-validity fares badly on all scores whereas α -validity and Y-validity fair perfectly well. The other notions are somewhere in between:

	\models_A	\models_B	\models_C	\models_D	\models_α	\models_X	\models_Y
<i>Objection 1</i>	x	x	✓	x	✓	✓	✓
<i>Objection 2</i>	x	✓	x	✓	✓	x	✓
<i>Objection 3</i>	x	✓	✓	✓	✓	x	✓

4. A-Validity to the Rescue?

The bad thing about this moral is that A-validity is, after all, a most natural way of extending the standard notion of argument validity to a supervaluational framework. Never mind the fact that the framework admits of other ways. There surely is an important sense in which we want to say that a good argument is one that never lets you go from true premises to false conclusions, and if truth is super-truth this means that the global notion of validity captured by condition (A) is an important one, if not the only one. Indeed, supervaluationists might easily concede Objection 1, on the grounds of the fact that multiple-conclusion arguments

²⁸ Indeed, the proof only establishes this fact on the assumption that D works as a necessity operator. We have already pointed out that when D is treated in analogy to the actuality operator, D-introduction may fail also on these readings of ‘ \models ’. It may also fail if D is treated along the lines mentioned in note 9. I’ll briefly come back to this point below.

are not paradigmatic, or because such arguments justify a collective understanding of what validity amounts to. But succumbing to the other objections is certainly something to worry about, since it would mean that A-validity yields a logic that is both far from classical and far from acceptable in the presence of higher-order vagueness.

Let us focus again on Objection 2 and the worry about classicality. There are two additional remarks that are worth making. First of all, one could observe that the presence of the D operator in the relevant counterexamples is crucial. In fact, all possible counterexamples to the inference rules [7]–[10] appear to be restricted to arguments involving the D operator, or similar devices (such as the truth predicate or vagueifiers such as ‘roughly’ or ‘-ish’) whose semantics requires truth on a precisification to call upon truth-values across a whole family of precisifications.²⁹ This is important, at least relative to D. For it means that whether or not [7]–[10] turn out to be valid depends crucially on the interpretation of D. Indeed, we have been assuming that D be interpreted as a necessity-like operator. But if we switch to the alternative interpretation mentioned earlier, that is, if we interpret D in analogy with the actuality operator, then [7]–[10] do preserve their classical status. For example, we have seen that on that interpretation D-introduction [12] fails, so [7] and [8] would both be safe. The counterexample to [9] would be blocked for similar reasons. Finally, [10] would be safe because the argument $\phi \vdash D\phi \vee D\neg\phi$ turns out not to be A-valid, since ϕ can be super-true even if it fails to be ‘already true’ on the partial interpretation of the language (e.g. when ϕ is ‘ $Fx \vee \neg Fx$ ’ and x is a borderline case of F , again assuming a partial semantics *à la* Kleene). So isn’t this enough to save A-validity from Objection 2? Isn’t it enough to switch from the necessity-like reading of D to its actuality-like reading?³⁰

²⁹ This point has been advertised in Fine 1975, p. 290, and Keefe 2000a, p. 178, though the only proof I am aware of (with regard to [7]) is in Graff 2003, p. 211. Conversely, one could contend that [7]–[10] may fail in classical logic, too, provided ‘ ϕ ’, ‘ ψ ’, etc. are allowed to range over open formulas. For instance, Williams 2005 points out that classically we have $Fx \models \forall xFx$ but not $\neg\forall xFx \models \neg Fx$. This may be right, but not enough to counter the objection. For it suffices to say that *when restricted to statements*, that is, closed formulas, [7]–[10] are classically but not super-valuationally valid. In any event, even the classical validity of $Fx \models \forall xFx$ is controversial. If ‘ $\phi \models \psi$ ’ is read as ‘if all assignments of values to variables satisfy ϕ , then all assignments of values to variables satisfy ψ ’, then the entailment holds. But not all textbooks follow this reading (which calls for a restriction of the Deduction Theorem). On the contrary, many classics read ‘ $\phi \models \psi$ ’ as ‘any assignment of values to variables that satisfies ϕ also satisfies ψ ’ (see e.g. Mendelson 1987, p. 52; Enderton 2001, p. 88). On that reading, the entailment $Fx \models \forall xFx$ fails.

³⁰ Equivalently, isn’t it enough to switch to the alternative semantics mentioned in note 9 (which blocks D-introduction)? See Williams 2005 for a proposal in this sense.

It isn't. Never mind the question of whether the actuality-like reading is a better candidate for modeling the English phrase 'It is definitely the case that'. Even if we switched to that reading, nothing would prevent us from introducing a different operator, D_{\square} , whose semantics is exactly the one we have so far been assuming for D , that is, the semantics akin to the necessity-like reading.³¹ Evidently, such an operator would satisfy the analogue of D -introduction and counterexamples to [7]–[10] could therefore be constructed with ' D_{\square} ' in place of ' D ', so Objection 2 would strike back. There just is no way to evade the objection simply by wangling with the semantics of this or that operator, for there is no way of banning all trouble-makers once and for all.

The second remark is this. Since D and similar devices play a crucial role in the counterexamples to [7]–[10], one might conclude that the lesson to be learned from Objection 2 is simply that A -validity requires extra caution in the presence of such special operators—hardly a surprising lesson given their characteristically metalinguistic flavor. Unfortunately, however, we cannot leave it at that. Surely, if extra caution is needed, we should better come up with suitably revised rules of inference that do work for us in the problematic cases as well. Can that be done?

Rosanna Keefe (2000a, pp. 179f) has offered the following positive answer: if the language contains the D operator, for instance (and let us stick to its necessity-like interpretation for simplicity), all we have to do is replace [7]–[10] with the following:

- [7*] From $\Sigma, \phi \models \psi$ infer $\Sigma \models D\phi \rightarrow \psi$
- [8*] From $\Sigma, \phi \models \psi$ infer $\Sigma, \neg\psi \models \neg D\phi$
- [9*] From $\Sigma, \phi \models \psi \wedge \neg\psi$ infer $\Sigma \models \neg D\phi$
- [10*] From $\Sigma, \phi \models \sigma$ and $\Sigma, \psi \models \sigma$ infer $\Sigma, D\phi \vee D\psi \models \sigma$.

These rules are indeed valid across the board, that is, on all seven readings of ' \models ', while reducing to [7]–[10] in the presence of bivalence, so in a way we have what we want. On the other hand, at this point we may still ask (and I think we ought to ask): What is the status of such 'rules'?

³¹ Well, perhaps this is a bit hasty. The referee pointed out that the combination of ' D_{\square} ' and ' D ' may raise odd cases. For example, one could create analogues of what Zalta (1988) called 'logical truths that are not necessary'. With the alethic actuality operator ' A ', when ϕ is only contingently true one gets that $\phi \leftrightarrow A\phi$ is logically true (for it is true no matter how the actual world is) but not necessarily true (for it is false at any possible world at which ϕ is false). Dual to this are logical falsehoods that are not necessarily false, like $\neg\phi \leftrightarrow A\phi$ (see Sobel 2004, p. 556, n. 13). It is not clear to me, however, whether such oddities, and especially their analogues for the 'definitely' operators, are a reason to *prevent* the combination.

If we think of them literally as *rules of inference*, then we are not done at all; we still have to show that the revised rules are enough to preserve semantic completeness, and as Brian Weatherson (2002, p. 34) has pointed out, this is far from obvious. For example, the following argument forms are supervaluationally valid in every sense of the term, yet they cannot be proved if [7]–[10] are replaced by [7*]–[10*]:

- [20] $\vdash p \rightarrow p$
- [21] $p \vdash \neg\neg p$
- [22] $p \vee q \vdash q \vee p$
- [23] $p \rightarrow q \vdash (p \wedge r) \rightarrow q.$

Perhaps this problem could be taken care of by relying on *both* sets of rules under suitable restrictions: use [7]–[10] for D-free arguments, otherwise use [7*]–[10*]. I am not sure this would work, but even if it did, things would again begin to look ugly and one might as well think that the right thing to do is to bite the bullet and give up A-validity altogether.

On the other hand, perhaps we should not think of the rules as rules of inference in the strict sense. It is customary to phrase them using the locution ‘From ... infer’, but really the locution ‘If ... then’ would be just as fine: rules are just metalinguistic conditionals whose application involves a metalinguistic *modus ponens*. (This is why logicians worry about what the Tortoise said to Achilles.³²) If so, then the problem appears to dissolve. Or rather, the initial objection turns into something else than what it was meant to be and Keefe’s cure seems perfectly all right. For it’s not that supervaluationism preserves classical logic only on a narrow conception of ‘logic’, where inferences count and rules of inference don’t. More simply, supervaluationism preserves classical logic (with respect to single-conclusion arguments) only insofar as we consider the object language. That some statements in the metalanguage fail to be preserved is no news: we already know that the bivalence principle, for instance, breaks down.

Unfortunately, this line of reasoning does not suffice to justify the cure. For whether or not we think of [7*]–[10*] as rules of inference, if such statements are true, then their truth will contribute to determining the logic of the object language. And if there is a sense in which this determination is too weak, as mentioned above, there is also a sense in which it is too strong. For example, inspection shows that the analogue of the S4 modal principle for the D operator turns out to be A-valid:

³² I am thinking of the ‘justification of deduction’ problem raised in Carroll 1895 and brought to current attention in Dummett 1973 and Haack 1976.

$$[24] \quad D\phi \rightarrow DD\phi.$$

(This follows from [7*] by taking ‘ ψ ’ to be ‘ $DD\phi$ ’, given the A-validity of D-introduction and transitivity.) And in the presence of second-order vagueness, the validity of [24] is just as problematic as the first-order validity of

$$[25] \quad \phi \rightarrow D\phi,$$

which is exactly what [7*] was designed to avoid. In other words, accepting Keefe’s cure is tantamount to ruling out higher-order vagueness, at least insofar as this can be expressed using the D operator on its necessity-like reading.

This brings us to Objection 3, which effectively shows that A-validity is incompatible with higher-order vagueness *regardless* of this maneuver to save it from Objection 2. For Objection 3 does not depend on any of the critical principles in [7]–[10]. It only depends on D-introduction, D-gap, and transitivity, all of which are intrinsically plausible.³³ Is there any way of deflating this result without giving up A-validity altogether? It seems to me that there are only two options: denying that the D operator can fully represent the concept of determinacy in the object language, or making sense of the idea that all vagueness is indeed first-order. (A third option would be to reject D-introduction by switching to the actuality reading of D. However, such a reading yields an S5 logic, so [24] would still hold.³⁴)

The first option can be spelled out as follows. Williamson (1994, p. 160) pointed out that there is a problem with the D operator in that hidden sharp boundaries are bound to pop up even if we allowed for non-trivial infinite iterations of D. For let us suppose we introduce a new operator to express determinacy at every level of the hierarchy:

$$[26] \quad D^*\phi =_{df} D\phi \wedge DD\phi \wedge DDD\phi \wedge \dots$$

This new operator is automatically transitive, hence it satisfies the analogue of the S4 axiom [24]:

$$[27] \quad D^*\phi \rightarrow D^*D^*\phi.$$

³³ Actually, in the presence of D-introduction and transitivity, D-gap amounts to the principle $Fx_j \models \neg D\neg D^n Fx_{j+1}$, and it is unclear whether this principle really expresses a non-negotiable intuition about higher-order vagueness. (Thanks to Richard Heck Jr for pointing this out.) Nonetheless, to reject D-gap on such grounds strikes me as exceedingly dismissive of the objection.

³⁴ Again, strictly speaking there is still a further option, corresponding to the alternative semantics for D outlined in note 9. On that semantics D-introduction fails, too, and so does [24]. However, this option eventually leads to the same issue discussed in the text with reference to the D^* operator, so I shall not consider it.

Does this mean that $D^*\phi$ cannot be indeterminate? Well, not necessarily: one could say that the vagueness of D^* is just not something that can be expressed using that very operator itself. As Williamson puts it, one could say that the D^* operator cannot be used to measure its own vagueness: that would be ‘like a cloud said to have an exact length because it is exactly as long as itself’ (1994, p. 160). To do the job properly, we would need to resort to a new operator, D^{**} , defined in a similar fashion, and the vagueness of D^* would manifest itself not in the failure of [27] but rather in the failure of

$$[28] \quad D^*\phi \rightarrow D^{**}D^*\phi.$$

Now—one could argue—if this line of reasoning is accepted, then why not apply it to the D operator itself? This operator satisfies the S4 axiom [24], but that need not amount to saying that D cannot be vague; it’s just that the vagueness of D cannot be expressed using that very operator itself. To do the job properly we would need to resort to a new operator D' , an operator capturing the first-order vagueness of the metalanguage, and point out that the truth of [24] does not imply the truth of

$$[29] \quad D\phi \rightarrow D'D\phi.$$

And so on and so forth. Informally, this is just a different way of saying that we need to ascend to the metalanguage to capture the vagueness of our talk of ‘admissible precisifications’, hence of super-truth. Supervaluationists who sympathize with the latter idea have no difficulty in accepting this line of response, thus saving A-validity from the charge of inadequacy *vis-à-vis* the phenomenon of higher-order vagueness (see Keefe 2000a, p. 210). In particular, the D-gap principle [11] on which Objection 3 relies would be rejected for similar reasons. If x_j and x_{j+1} are two adjacent members of a sorites series for a predicate F , then there are two possibilities: either F does not suffer from higher-order vagueness, or it does. If it doesn’t, then the objection misfires. If it does, then D-gap is not unrestrictedly A-valid. It is only first-order A-valid, so to say—that is, valid for $n = 0$. In other words, we can accept [30] but not [31], thus blocking the objection at the second step of the proof:

$$[30] \quad DFx_j \models_A \neg D\neg Fx_{j+1}$$

$$[31] \quad DDFx_j \models_A \neg D\neg DFx_{j+1}.$$

I think this way of handling the problem is technically correct and perfectly defensible on formal grounds. Indeed, we have already said that supervaluationism must come to terms with the fact that it may be vague which precisifications should count as admissible, and this is just the same idea spelled out in re-

gard to the logic of D.³⁵ I myself have sympathized with this view for a long time. However, I am no longer convinced it is fully acceptable. One problem is simply that a hierarchy of definiteness operators faces the same sort of objections that have been raised against Tarskian theories of truth. Keefe says that the multiplicity of operators is still compatible with a univocal account of the meaning of ‘definitely’ insofar as ‘its formalization within a sentence can depend *systematically* on its position in the embedding’ (2000a, p. 210). That may be right, and it might even be right with regard to the truth predicate.³⁶ On the face of it, however, it appears that we can use ‘true’ and ‘definitely’ to pick out concepts outside the hierarchy, as when we say that everything John believes is true, or when we object to ontological vagueness by claiming that everything in the world is determinately one way or the other. In this sense, the analogy between D and D* does not seem to be fair.

Another problem is that there is something fishy in a hierarchy that works that way. Supervaluationism accounts for first-order vagueness by representing the object language *via* its precisifications. But if higher-order vagueness can only be accounted for by representing our metalanguage(s) in a similar fashion, then it means that the higher-order vagueness of an object-language predicate such as ‘small’ is really the first-order vagueness of ‘definitely’ (or ‘true’). Now, there certainly is a correlation between the former and the latter. But as already Fine (1975, p. 148) pointed out, it would seem that the latter arises from the former, not vice versa. ‘Definitely’ and ‘true’ supervene upon the object language: there can be no independent grounds for their having borderline cases. Another way of phrasing this worry is that on this approach first-order vagueness is unwarrantably *sui generis*. As Brian Weatherson (2002, p. 47) pointed out, the approach accounts for first-order vagueness in terms of multiplicity of precisifications, and second- and higher-order vagueness in terms of indeterminacy of the multiplicity, indeterminacy of this indeterminacy, and so on. Why the bifurcation?—Weatherson asks. If at every higher level vagueness amounts to there being indeterminacy as to what the relevant precisifications are, why not represent the object language itself through *one* precisification and regard all vagueness as coming about from there being indeterminacy as to what *that* precisification looks like?

³⁵ Patrick Greenough (2005, pp. 185f) raises the following worry. For each $n \geq 1$, let D_n be the operator capturing the n -th order vagueness of the language (so, in particular, $D_1 = D$ and $D_2 = D'$). Then, in a finite sorites for F , there is bound to be some n such that, for every x , either $D_{n+1}D_n \dots D_1 Fx$ or $D_{n+1} \neg D_n \dots D_1 Fx$, which is to say that the series is bound to involve a sharp cut-off at the $n+1$ -th level. However, this strikes me as all right, unless higher-order vagueness is identified with what Sainsbury (1990) calls ‘boundarylessness’ (as urged e.g. in Horgan 1998).

³⁶ But see McGee 1991.

The answer, of course, is that going this way would take us back to the Simple Theory of vagueness reflected in the acceptance of the Tarski biconditional [14]: in Fine's words, the vagueness of truth would wax and wane with the vagueness of the statement to be evaluated (1975, p. 149). We saw that the reason why this theory won't suffice is that it leaves us with the task of explaining the logic of the metalanguage—the 'if and only if' in Tarski's biconditional. But that is precisely what this approach says we should do; we should attend to the vagueness of the metalanguage. So either we don't do that or, if we do it, then we should stop worrying about A-validity in the first place. Global validity has no room in the Simple Theory of vagueness, for super-truth vanishes: on the Simple Theory, the multiplicity of admissible precisifications is only relevant to the semantics of 'definitely' and the only supervenient notion of validity is local (our α -validity).

All of this speaks against the first option mentioned above—denying that D can fully represent definiteness in the object language. What about the second option? Can one save A-validity by making sense of the idea that all vagueness is in fact first-order? To most people, this would just be biting the bullet. Of course there is higher-order vagueness, for just as our linguistic practices and beliefs have failed to draw a precise boundary between the small and the non-small, they have failed to draw a precise boundary between the small and smallish, or between the smallish and the smallish-ish. The trouble with vagueness is that vague predicates—and vague expressions at large—draw no boundaries at all, not that they merely fail to draw the boundaries presupposed by the Fregean ideal of precision. I agree. Nonetheless it is legitimate to ask: Why so? What is the argument to the effect that accepting vagueness *entails* accepting higher-order vagueness?

One argument to this effect—in fact, the only argument I am aware of—comes from Dominic Hyde (1994). According to Hyde, the existence of higher orders of vagueness follows from the twofold consideration that (i) there are vague predicates and (ii) 'vague' is one of them. Now, there is no question about (i), so the argument depends entirely on (ii). And what reasons do we have to accept (ii)? Hyde appeals to a proof by Roy Sorensen (1985), to the effect that 'vague' is prone to the sorites paradox. However, as I pointed out elsewhere (Varzi 2003b), Sorensen's proof depends on the assumption that some predicates (specifically the predicate 'small') are higher-order vague, so one cannot rely on *that* proof in order to establish the existence of such predicates. For the record, elsewhere (Varzi 2005) I have also defended Sorensen's proof against a number of formal misgivings that have been raised even recently, so I do accept (ii), hence I do find Hyde's argument sound. Nonetheless the argument is circular and cannot be used to establish the point under examination, namely the existence of higher-order vagueness. So unless one can do better, the option of denying

higher-order vagueness is in principle available, and the friend of supervenient A-validity might just resort to it.

As I said, to most of us this sounds like biting the bullet. But let me spend a word on behalf of this option, since at this point so much depends on it. Let us focus on a vague predicate F . The intuition is that it is hard to accept the existence of a sharp boundary between the borderline cases of F and its clear positive instances (or its clear negative instances). Why so? The informal answer is that in a sorites series for F , the last positive instance and the first borderline case are just as indiscernible (in the relevant respects) as any other adjacent cases. However, this is not to say that there is no such boundary; it's just that it is impossible to draw it. More precisely, one could argue that it is impossible to draw a sharp line demarcating those items that are relevantly indiscernible from the end members of the series from those items that are not so indiscernible, and that the source of this impossibility is epistemic. Such reasons, as Robert Koons (1994) pointed out (and Williamson, 1990, Sect. 6.3, before him), would lie in the fact that the property of being indiscernible with respect to every property that is relevant for the application of the predicate F is not *decidable*. More precisely, is only semi-decidable. If a and b are discernible, then we can eventually find out—either by direct comparison, or by finding a third element c so that a is discernible from c while b isn't. But we may not be able to tell if two given objects a and b are indiscernible, for the task of going through all possible relevant comparisons with other objects is a never-ending task. To illustrate, suppose a_j is any member of the series that is discernible—with respect to F —from the initial member, a_0 . We can explain this fact by directly comparing a_j to a_0 or, if they look relevantly similar, by noting that a_j , but not a_0 , is indiscernible from some later item a_{j+k} . However, we may not be able to tell whether a_j is the *first* such object, for we may not be able to verify whether its predecessor is indiscernible from a_0 . In short, the idea is that a semantic account of the vagueness of F can be combined with an epistemicist account for the vagueness of indiscernibility (i.e. the predicate 'indiscernible' or the underlying concept). And this would suffice to justify the claim that all vagueness is indeed first-order.

Of course, one may respond that if we are willing to appeal to epistemicism when it comes to indiscernibility, we might as well buy into wholesale epistemicism with regard to vagueness: what's the point of splitting the account? But there is an answer: 'indiscernible' is a relational predicate, and there are good reasons to think that in many cases relational predicates are not semantically vague even if they yield *prima facie* semantic indeterminacy. The identity predicate is arguably a case in point. Surely 'identical' is not semantically vague, yet one could argue that it is epistemically vague: if x 's sharing the same properties as y is a necessary

condition to determine whether $x = y$, the truth of this identity statement may be epistemically indeterminate owing to the impossibility of surveying all the properties. Such is the drawback of Leibniz's law—the *indiscernibility* of identicals—when it comes to its application to concrete identity issues. No serious supervaluationist would claim that the *prima facie* indeterminacy of the mind-body problem is a sign of semantic vagueness, at least not without buying into a whole metaphysical package that calls for independent justification.

Be that as it may, let me emphasize that I am not defending this account. I'm just outlining it on behalf of the view that all vagueness is first-order, which is what the friend of A-validity is forced to hold. That it is a good account is a different story. In fact, my overall assessment is that going this way simply shifts the burden of proof: the argument in favour of higher-order vagueness calls for non-question-begging evidence for the claim that 'vague' is semantically vague, like 'small'; the argument against higher-order vagueness calls for non-question-begging evidence for the claim that 'indiscernible' is epistemically vague, like 'identical'. Strictly speaking we are at a deadlock. But in all honesty, it is awkward that in order to rescue the notion of A-validity and the primacy of supertruth, a supervaluationist should find herself in the business of arguing for an epistemic way out.

5. α -Validity for Everybody

With all this, I think supervaluationists should better take seriously the idea that when it comes to the logic of a vague language, global validity—and A-validity in particular—is not as good a notion as one might initially think. And since we have seen that the only two notions of validity that do not suffer from the objections are the local notion of α -validity and the collective notion of Y-validity, supervaluationists should take seriously the idea that these notions are not as bad as one might have thought. In fact, since Y-validity does not correspond to the intuition that validity is preservation of truth (but only the weaker intuition that validity is preservation of non-falsity), supervaluationists should take seriously the idea of 'going local'.³⁷ As I have mentioned, this idea is particularly attractive if we back

³⁷ I have focused on the notion of validity, but of course going local will have an impact on other notions, too. For instance, standardly an argument is *sound* if and only if it is valid and its premises are true. We may stick to this definition on the understanding that truth is super-truth. However, if validity is construed locally, a sound argument need not have a true conclusion, otherwise the sorites paradox would strike back. (This can be verified by considering the paradox in the form: $Fx_1 \vdash \neg(Fx_1 \rightarrow Fx_2), \dots, \neg(Fx_{n-1} \rightarrow Fx_n), Fx_n$, where $x_1 \dots x_n$ form a sorites series with respect to F ; this argument is sound, yet none of the conclusions is super-true.) Locally, from the

it up with the thought that vagueness is ultimately a modal phenomenon, as Dever *et al.* 2004 have argued. But I have also pointed out that the notion of α -validity is reasonable in its own right, both with respect to type-(1) precisifications and with respect to type-(2) precisifications (at least insofar as these are construed in conformity to options (a) and (i), i.e., as forming a cluster of completely precise languages or interpretations that truly and fully define our vague language). So let me conclude with a general consideration to the effect that there is in fact a good reason to focus primarily on this notion of argument validity from a general super-evaluative perspective.

In a nutshell, the reason is simply that α -validity, which is fully classical and yet sensitive to vagueness of any order, allows us to recast global validity as a defined notion. For let us introduce two operators T and F to express super-truth and super-falsity, respectively, with the obvious semantics: $T\phi$ is true on a precisification if and only if ϕ is true on every precisification, and $F\phi$ is true on a precisification if and only if ϕ is false on every precisification. Evidently both operators satisfy all the axioms of the modal logic S5, so there is no room for non-trivial iterations. But that's fine: there may be indeterminacy as to what qualifies as the correct precisification space, yet relative to any precisification space super-truth and super-falsity are perfectly precise, though not exhaustive. Using these operators, we can then recast each notion of global validity via the following equivalences:³⁸

- [32] $\Sigma \models_A \Gamma$ if and only if $\{T\phi: \phi \in \Sigma\} \models_\alpha \{T\psi: \psi \in \Gamma\}$
- [33] $\Sigma \models_B \Gamma$ if and only if $\{F\psi: \psi \in \Gamma\} \models_\alpha \{F\phi: \phi \in \Sigma\}$
- [34] $\Sigma \models_C \Gamma$ if and only if $\{T\phi: \phi \in \Sigma\} \models_\alpha \{\neg F\psi: \psi \in \Gamma\}$
- [35] $\Sigma \models_D \Gamma$ if and only if $\{\neg T\psi: \psi \in \Gamma\} \models_\alpha \{F\phi: \phi \in \Sigma\}$.

In fact, we can recast the collective notions of validity, too. That would be straightforward if the language contained infinitary conjunctions and disjunctions, but even with the resources of a standard grammar we can easily do the job in view of the following equivalences:

- [36] $\Sigma \models_X \Gamma$ if and only if $\{T(\phi_1 \wedge \dots \wedge \phi_n): n > 0 \text{ and } \phi_1, \dots, \phi_n \in \Sigma\} \models_\alpha \{T(\psi_1 \vee \dots \vee \psi_n): n > 0 \text{ and } \psi_1, \dots, \psi_n \in \Gamma\}$.
- [37] $\Sigma \models_Y \Gamma$ if and only if $\{F(\psi_1 \vee \dots \vee \psi_n): n > 0 \text{ and } \psi_1, \dots, \psi_n \in \Gamma\} \models_\alpha \{F(\phi_1 \wedge \dots \wedge \phi_n): n > 0 \text{ and } \phi_1, \dots, \phi_n \in \Sigma\}$.

truth of the premises we can infer that all precisifications verify some conclusion, not that some conclusion is verified by all precisifications. (Thanks to Robert Williams for raising this point.)

³⁸ The equivalence in [32], hence the possibility of defining global, truth-preserving validity in terms of local validity, is also pointed out in Dever *et al.* 2004, Sect. 3.2, and in Shapiro 2006, Ch. 4, Theorem 16.

So there is a clear sense in which settling on α -validity does not amount to neglecting the other notions. On the contrary, each ' \models_i ' may be fully expressed in term of ' \models_α '. If desired, at this point we could even go further. We could make room for many other notions along the same lines, obtained from [32]–[37] by replacing 'T' with 'D', or with 'D"', or with 'D*', and 'F' with 'D¬', 'D"¬', or 'D*¬', respectively. And we could make room for a variety of corresponding notions even if we worked with (b)(ii)-style precisifications, focusing on a D operator characterized semantically in conformity with the second option mentioned at the end of section 1, that is, in analogy to the actuality operator of modal logic.

Of course, conceptual reduction is no big deal. After all, all notions of validity are already defined in terms of one same basic concept—precisification. It is nonetheless significant to be able to keep track of all these notions in the object language, liar paradox permitting. And it is important to be able to do so in terms of a single reading of ' \models ' that (i) captures both the truth-preserving and the falsity-avoiding features of the classical notion of validity, (ii) is not open to the objection from higher-order vagueness, and (iii) is fully classical. In discussing the availability of distinct notions of validity within the same supervaluationary framework (specifically the notions corresponding to A-, C-, and α -validity), Keefe has recommended a pluralist attitude: even granting the normative aspect of our intuitive conception of validity, the question of whether we should endorse an argument 'may itself be ambiguous or somewhat indeterminate—what we should endorse can depend on our purposes' (2000b, p. 103). With reference to the alternative between A- and C-validity, for example, Keefe observes that although we are typically concerned to infer truths from truths (so that A-validity would be appropriate), there may be contexts, such as legal contexts, in which our main concern is rather to avoid inferring falsehoods (so that C-validity would be better suited). Indeed, the failure of transitivity may even be a welcome feature when it comes to the latter sort of context: the sorites paradox may itself be construed as a context in which we want to avoid inferring a falsehood, and its paradoxical nature stems precisely from the fact that we are inclined to endorse each individual step without endorsing the big jump from the initial premise to the final conclusion. Likewise, concerning the choice between global validity (of type A or C) and local validity (of type α), Keefe observes that precisely because the conflict arises only in the presence of the D operator and similar devices, at least relative to single-conclusion arguments, the choice among the options is not up for grabs. An inference in the form of D-introduction, for instance, is globally but not locally valid, but that's because it sets a context in which the local reading is not plausible: it would be a mistake to portray such an inference as one that may fail

on certain ways of disambiguating it, that is, ways of resolving semantic indecision; for if it were decided that ϕ is true, $D\phi$ could not remain false. This does not help us choose between global and local validity, but it suggests that D-enriched arguments are not effectively interpreted in terms of ambiguity or semantic indecision, so the global reading appears to be better suited. On the other hand, I have mentioned earlier that in some cases the local reading seems better, for instance when we wish to reason under hypotheses that are not super-true: we may want to argue that the assumption that x is small, together with the fact that y 's height is less than x 's, entails that y must be small, too. Local reasonings of this sort reflect the penumbral principles that constrain the notion of an admissible precisification, so in a way they do not *establish* those principles on pain of circularity. But that is not to deny that we may find ourselves reasoning along such lines.

I don't have anything against this sort of pluralism. In principle, however, I am more sympathetic to a pluralism that stems from an ambiguity in our intuitive notion of a possible context (as JC Beall and Greg Restall 2000 have emphasized³⁹) rather than from an ambiguity in the normative aspect of our intuitive notion of validity. If there are ambiguities in the latter sense, we should welcome the opportunity to resolve them explicitly in the object language, by suitably representing the logical forms of our arguments. And that is precisely the value of the reduction afforded by [32]–[37]. Let us fix on one notion of validity, α -validity. If we want our argument to be read globally in the A-sense, let's *phrase* it accordingly, with the help of the T operator. If we want it to be read globally in the C-sense, let us phrase it accordingly. And so on. This strikes me as a better way of handling any potential pragmatic tension that may arise from the multiplicity of purposes with which we can put forward an argument. Keefe is right in saying that insofar as we can define a notion of validity that is preservation of supervaluational truth, namely global A-validity, it may look unwarranted to identify supervaluational validity with something else, namely local α -validity. Yet what looks unwarranted may not be so. Insofar as we can express A-validity in terms of α -validity, and insofar as the latter is perfectly classical, there are indeed good reasons to settle on that notion and avoid the fogs of pluralism.

Let me conclude, then, with a general remark concerning the T operator. The reduction afforded by [32]–[37] depends on the availability of this operator in the object language, which is how super-truth enters the picture, so something must be said about it. I have said that its logic is S5. But there is more. Supervaluationally, truth is not disquotational. As Williamson (1994, p. 162) famously argued, if

³⁹ In fact, my views on the matter are more radical than Beall and Restall's, but this is not the place to elaborate. See Varzi 2002.

it were, then the supervaluationist would be forced to admit bivalence, since the following argument form is supervaluationally valid:

$$[38] \quad \psi \leftrightarrow T\psi, \neg\psi \leftrightarrow T\neg\psi \models T\psi \vee T\neg\psi.$$

In fact, inspection shows that [38] is valid on every reading of ‘ \models ’, though it may be difficult to *prove* that it is using classical rules of inference (for instance, standardly one appeals to proof by cases and transitivity, and these rules are invalid on the A- and X- readings of ‘ \models ’ and on the C- reading, respectively). However, in the present context the relevant sense of validity is clearly the truth-preserving sense: to say that [38] would force a supervaluationist to ‘admit bivalence’ is to say that accepting the premises as true would commit her to accepting the conclusion as true. So really it is only the A-, α -, and X- readings of ‘ \models ’ that are relevant here. And on these readings, Williamson’s diagnosis is correct: supervaluationally we are inclined to reject the conclusion (and claim failure of bivalence) when ψ is indeterminate, and in those cases the two premises corresponding to Tarski’s schema are not super-true but indeterminate. Williamson’s takes this to be a bad thing, since he regards the disquotational property to be central to any respectable notion of truth, and the supervaluationist owes a response.

Now, with regard to the A-reading, the popular response is well known: the loss of

$$[39] \quad T\phi \leftrightarrow \phi$$

is compensated by the upholding of the equivalence

$$[40] \quad T\phi \dashv\vdash \phi,$$

and for a supervaluationist it is this equivalence that—*pace* Williamson—captures the essence of the Tarskian intuition.⁴⁰ In other words, T is not disquotational, but it still correlates in the appropriate way with the statements it applies to. I accept this line of response. Indeed, for what it’s worth, it is obvious that the response applies also on the collective, X-based conception of argument validity, which in this case coincides with the A-based conception. It does *not*, however, apply on the local conception, since on that reading the right-to-left direction of [40] may fail when ϕ is indeterminate. On the local reading both [39] *and* [40] may fail, and that is something one may worry about. Does it mean that going local forces us to give up the Tarskian intuition altogether?

⁴⁰ Compare Keefe 2000a, Sect. 8.3. Independently of supervaluational insights, McGee (1989) argues that it is indeed entailment relations that do the work to which our intuitive notion of truth is typically put.

Not quite. The local reading—we have said—corresponds to a twofold idea, depending on whether we think of precisifications according to option (1), as forming the cluster of languages that we normally construe as a single vague language, or according to option (2), as forming the cluster of interpretations that are compatible with the explicit semantic decisions characterizing our vague language. Consider first option (1). In that case, the rationale behind α -validity rests on the idea of systematic disambiguation: just as assessing the truth of a statement amounts to checking whether it is true no matter how we systematically disambiguate it, that is, no matter how we consider one of the many languages that we may be speaking, so assessing the validity of an argument amounts to checking whether it is valid no matter how we systematically disambiguate its premises and conclusions. Now, we have just seen that when the D operator is involved, it may be misleading to consider disambiguations in which the meaning of D is kept fixed: if we were speaking a language in which ϕ is true, we might want to say that in that language $D\phi$ is also true, regardless of ϕ 's truth value in other languages of the cluster. This suggests that there are *two* ways in which we can locally assess our statement, or argument: one on which 'D' has the meaning it has in our vague language, and another in which the meaning of 'D' is adapted to the context of our disambiguating scenario. On the former understanding, D-introduction fails. On the latter, D-introduction is irrelevant because we are actually treating D itself as a vague operator—we are thinking of D as an operator whose behaviour changes depending on the language we speak, and whose semantics is fixed, not by the rule in [41], but by the rule in [42]:

[41] $D\phi$ is true on a precisification P if and only if ϕ is true on all precisifications accessible from P .

[42] $D\phi$ is true on a precisification P if and only if ϕ is true on P .

Well, then: the same sort of consideration can be applied to statements or arguments involving the T operator. Generally speaking, 'T' stands for 'it is super-true that', and we can have good use for such an operator. Among other things, we need it to implement the reductive definitions in [32]–[37]. Nonetheless, when it comes to evaluating a statement or an argument in which some claims are said to be true, we should pay attention to the sense in which 'true' is being used. One sense is captured by the supervaluational semantics for 'T', as per the condition in [43] that we have been assuming so far. Another sense is the local one, captured by the condition in [44]:

[43] $T\phi$ is true on a precisification P if and only if ϕ is true on all precisifications.

[44] $T\phi$ is true on a precisification P if and only if ϕ is true on P .

On their local reading, T and D would reduce to the same operator, and plausibly so, since bivalence holds on every precisification. Accordingly, let us write ‘ Δ ’ to indicate such an operator, as fixed by [42] or [44]. (Effectively, this is the operator corresponding to Fine’s ‘ true_T ’, 1975, pp. 148f.) Then the idea is that on the second understanding of how we can go about assessing a statement ϕ or a whole argument $\Sigma \vdash \Gamma$, we do so upon reading it as involving Δ rather than T. Depending on which language we are speaking, certain statements would be true and others false *in that language*, so if ϕ or any element of $\Sigma \cup \Gamma$ contains the T operator, we should disambiguate it by replacing ‘T’ with ‘ Δ ’: disambiguation may require some adjustments in logical form. So, now, when it comes to statements of the form [39] or arguments of the form [40]—the statements and arguments that putatively reflect the Tarskian intuition—there are two ways of assessing them. We can assess them as being about truth *qua* super-truth, or we can understand them as being about local or regimented truth. In the former case, both fail to hold when ϕ is indeterminate, for super-truth dissolves in the context of a single disambiguated language; in the latter case, however, both hold unrestrictedly. In other words, the following are both α -valid:

$$[45] \quad \Delta\phi \leftrightarrow \phi$$

$$[46] \quad \Delta\phi \dashv\vdash \phi$$

If we construe precisifications according to option (2), the story is slightly different but the moral is the same. On that construal, we are supposed to interpret our vague language by means of a cluster of classical models, just as we interpret a modal language by means of a cluster of possible worlds, and the T operator registers what is true in every model of the cluster. Again, it is obvious that [39] and [40] may fail on this picture. However, we can certainly enrich our language with the Δ operator, to keep track of what is true and what is false in the individual models. In particular, we may want to keep track of the fact that every admissible model is perfectly standard, hence bivalent—a fact that we certainly cannot express by uttering the law of the excluded middle, or by uttering the principle of bivalence using D or T. So, again, when it comes to assessing a statement or an argument, we must distinguish between the global—as we may now call it—or the local sense in which we use such phrases as ‘is it definitely the case that’ or ‘it is true that’. If we go for the global sense, then these phrases correspond to D and T, since we are interested in super-truth and global validity. But if we go for the local sense, then we should go for Δ : surely a standard model is one relative to which such phrases cancel out.

So, yes, super-truth fails to be disquotational on a supervaluational logic in which validity is α -validity, both in the stronger sense expressed by Tarski’s bi-

conditional [39] and in the weaker sense expressed by the equivalence in [40]. But this is not to say that on such a logic we are giving up the disquotational intuition altogether. In fact, I have presented the picture as one involving an ambiguity between two notions of truth, only one of which is disquotational. That may sound cheap; it may sound like a concession to the sort of pluralism that I have discouraged with respect to the notion of validity. But that was just to fix the intuition. It should be clear that on this picture we can ultimately *express* the ambiguity directly in the object language, just as we can express any potential ambiguity of the notion of validity. So let me summarize the picture by rephrasing it as follows. Forget about \top . What is crucial, if we want to *argue* about truth, is that our language contains the Δ operator—and this is perfectly disquotational. The \top operator is only needed to endow such arguments with global or collective readings. And surely that is only possible if \top does not cancel out. That is why [40] cannot be locally valid. What *is* valid, as we know, is the global reading of [40], which is to say its A-reading. And this remains true. For once recast via [32], the A-reading of [40] amounts to the following claim:

$$[47] \quad \top\top\phi \dashv\vdash \top\phi.$$

And this claim is perfectly (and trivially) α -valid.⁴¹

Department of Philosophy
Columbia University
New York NY 10027
USA

ACHILLE C. VARZI

References

- Beall, J. C., and Restall, G. 2000: ‘Logical Pluralism’. *Australasian Journal of Philosophy*, 78, pp. 475–493.
- Bencivenga, E. 1978: ‘Free Semantics for Indefinite Descriptions’. *Journal of Philosophical Logic*, 7, pp. 389–405.

⁴¹ Early drafts of this paper have been presented at the *Second Workshop on Vagueness* held at the University of Navarra in Pamplona, November 14, 2005 (with commentary by Paloma Pérez-Ilzarbe), at the *Fifth Arché Vagueness Workshop* held at the University of St. Andrews, November 18–19, 2005, and at the *Workshop on Vagueness* held at the University of Buenos Aires, November 8–9, 2006 (with commentaries by Eleonora Cresto and Javier Castro Albano). I am thankful to the organizers, the commentators, and the participants in each event for their helpful suggestions and criticisms. The final version has benefited especially from comments by Richard Heck Jr, Robert Williams, and an anonymous referee of *Mind*.

- Burns, L. 1991: *Vagueness: An Investigation into Natural Languages and the Sorites Paradox*. Dordrecht: Kluwer.
- Carroll, L. 1985: 'What the Tortoise Said to Achilles'. *Mind*, 4, pp. 278–280.
- Collins, J., and Varzi A. C. 2000: 'Unsharpenable Vagueness'. *Philosophical Topics*, 28, pp. 1–10.
- Dever, J., Asher, N., and Pappas, C. 2004: 'Supervaluationism Debugged'. Unpublished.
- Dummett, M. 1973: 'The Justification of Deduction'. *Proceedings of the British Academy*, 59, pp. 201–232.
- 1975: 'Wang's Paradox', in Keefe and Smith 1996, pp. 99–118. Originally published in *Synthese*, 30.
- 1991: *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard University Press.
- Edgington, D. 1993: 'Wright and Sainsbury on Higher-Order Vagueness'. *Analysis*, 53, pp. 193–200.
- Enderton, H. B. 2001: *A Mathematical Introduction to Logic, Second Edition*. San Diego, CA: Academic Press.
- Fine, K. 1975: 'Vagueness, Truth and Logic', in Keefe and Smith 1996, pp. 119–150. Originally published in *Synthese*, 30.
- Fodor, J. A., and LePore, E. 1996: 'What Cannot Be Evaluated Cannot Be Evaluated, and It Cannot Be Supervalued Either'. *Journal of Philosophy*, 93, pp. 516–535.
- Frege, G. 1903: *Grundgesetze der Arithmetik, begriffsschriftlich abgeleitet, Band II*. Jena: Pohle. Partial Eng. trans. in *Translations from the Philosophical Writings of Gottlob Frege*, ed. by P. T. Geach and M. Black, Oxford: Blackwell, 1952, pp. 21–41.
- Graff, D. 2003: 'Gap Principles, Penumbral Consequence, and Infinitely Higher-Order Vagueness', in J. C. Beall (ed.), *Liars and Heaps: New Essays on the Semantics of Paradox*. Oxford: Clarendon Press, pp. 195–221.
- Greenough, P. 2005: 'Contextualism about Vagueness and Higher-Order Vagueness'. *Proceedings of the Aristotelian Society*, Suppl. Vol. 79, pp. 167–190.
- Haack, S. 1976: 'The Justification of Deduction'. *Mind*, 85, pp. 112–119.
- Heck, R. G., Jr. 1993: 'A Note on the Logic of (Higher-Order) Vagueness'. *Analysis*, 53, pp. 201–208.
- Herzberger, H. G. 1982: 'The Algebra of Supervaluations', *Topoi*, 1, pp. 74–81.
- Horgan, T. 1998: 'The Transvaluationist Conception of Vagueness'. *The Monist*, 81, pp. 313–330.
- Hyde, D. 1994: 'Why Higher-Order Vagueness Is a Pseudo-Problem'. *Mind*, 103, pp. 35–41.
- 1997: 'From Heaps and Gaps to Heaps and Gluts'. *Mind*, 106, pp. 641–660.
- Kamp, H. 1975: 'Two Theories about Adjectives', in E. L. Keenan (ed.), *Formal Semantics and Natural Language*. Cambridge: Cambridge University Press, pp. 45–89.
- Keefe, R. 2000a: *Theories of Vagueness*. Cambridge: Cambridge University Press.
- 2000b: 'Supervaluationism and Validity'. *Philosophical Topics*, 28, pp. 93–105.
- Keefe, R., and Smith, P. (eds) 1996: *Vagueness: A Reader*. Cambridge, MA: MIT Press.
- Kleene, S. K. 1952: *Introduction to Metamathematics*. Amsterdam: North-Holland.

- Koons, R. C. 1994: 'A New Solution to The Sorites Problem'. *Mind*, 103, pp. 439–450.
- Kremer, P., and Kremer, M. 2003: 'Some Supervaluation-Based Consequence Relations'. *Journal of Philosophical Logic*, 32, pp. 225–244.
- Kripke, S. 1972: *Naming and Necessity*, Cambridge, MA: Harvard University Press, 1980. Originally published in D. Davidson and G. Harman (eds), *Semantics of Natural Language*. Dordrecht: Reidel.
- Lewis, D. K. 1975: 'Language and Languages', in his *Philosophical Papers: Volume I*. New York: Oxford University Press, 1983, pp. 163–188. Originally published in *Minnesota Studies in the Philosophy of Science*, 7.
- 1986: *On the Plurality of Worlds*. Oxford: Blackwell.
- 1993: 'Many, but Almost One', in his *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press, 1999, pp. 164–182. Originally published in K. Campbell, J. Bacon, and L. Reinhardt (eds), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong*. Cambridge: Cambridge University Press.
- Machina, K. F. 1976: 'Truth, Belief, and Vagueness', in Keefe and Smith 1996, pp. 174–203. Originally published in *Journal of Philosophical Logic*, 5.
- McGee, V. 1989: 'Applying Kripke's Theory of Truth'. *Journal of Philosophy*, 86, pp. 530–539.
- 1991: *Truth, Vagueness, and Paradox: An Essay on the Logic of Truth*. Indianapolis, IN: Hackett.
- McGee, V., and McLaughlin, B. 1995: 'Distinctions without a Difference'. *Southern Journal of Philosophy*, 33 (Supplement), pp. 203–251.
- 1998: 'Logical Commitment and Semantic Indeterminacy: A Reply to Williamson'. *Linguistics and Philosophy*, 21, pp. 221–235.
- 1999: 'The Value of Supervaluations'. Unpublished.
- Mendelson, E. 1987: *Introduction to Mathematical Logic, Third Edition*. Monterey, CA: Wadsworth & Brooks/Cole.
- Morreau, M. 1999: 'Supervaluations Can Leave Truth-Value Gaps after All'. *Journal of Philosophy*, 96, pp. 148–156.
- Przełęcki, M. 1976: 'Fuzziness as Multiplicity'. *Erkenntnis*, 10, pp. 371–380.
- Rescher, N., and Brandom, R. 1980: *The Logic of Inconsistency: A Study in Non-Standard Possible-World Semantics and Ontology*. Oxford: Basil Blackwell.
- Restall, G., 2005: 'Multiple Conclusions', in P. Hajek, L. Valdes-Villanueva, and D. Westerstaahl (eds), *Logic, Methodology and Philosophy of Science: Proceedings of the 12th International Congress*. London: Kings' College Publications, pp. 189–205
- 1990: 'Concepts Without Boundaries', in Keefe and Smith 1996, pp. 251–264. Originally delivered as Inaugural Lecture, Department of Philosophy, King's College, London.
- Sanford, D. 1976: 'Competing Semantics of Vagueness'. *Synthese*, 33, pp. 195–210.
- Shapiro, S. 2006: *Vagueness in Context*. Oxford: Oxford University Press.
- Scott, D. 1975: 'Combinators and Classes', in C. Böhm (ed.), *λ -Calculus and Computer Science Theory*. Berlin: Springer Verlag, pp. 1–26.

- Sobel, J. W. 2004: *Logic and Theism: Arguments for and against Beliefs in God*. Cambridge: Cambridge University Press.
- Sorensen, R. A. 1985: 'An Argument for the Vagueness of "Vague"'. *Analysis*, 45, pp. 134–137.
- Tappenden, J. 1993: 'The Liar and Sorites Paradoxes: Toward a Unified Treatment'. *Journal of Philosophy*, 90, pp. 551–577.
- Van Fraassen, B. C. 1966: 'Singular Terms, Truth-Value Gaps, and Free Logic'. *Journal of Philosophy*, 63, pp. 481–495.
- Varzi, A. C. 1997: 'Inconsistency Without Contradiction'. *Notre Dame Journal of Formal Logic*, 38, pp. 621–638.
- 1999: *An Essay in Universal Semantics*. Dordrecht: Kluwer.
- 2000: 'Supervaluationism and Paraconsistency', in D. Batens, C. Mortensen, G. Priest, and J. P. Van Bendegem (eds), *Frontiers in Paraconsistent Logic*. Baldock: Research Studies Press, pp. 279–297.
- 2002: 'On Logical Relativity'. *Philosophical Issues*, 12, pp. 197–219.
- 2003a: 'Cut-Offs and their Neighbors', in J. C. Beall (ed.), *Liars and Heaps: New Essays on the Semantics of Paradox*. Oxford: Clarendon Press, pp. 24–38.
- 2003b: 'Higher-Order Vagueness and the Vagueness of "Vague"'. *Mind*, 112, pp. 295–299.
- 2004: 'Conjunction and Contradiction', in G. Priest, J. C. Beall, and B. Armour-Garb (eds), *The Law of Non-Contradiction: New Philosophical Essays*. Oxford: Oxford University Press, pp. 93–110.
- 2005: 'The Vagueness of "Vague": Rejoinder to Hull'. *Mind*, 114, pp. 695–702.
- Weatherston, B. 2002: 'Supervaluationism'. Unpublished.
- Williams, J. R. G. 2005: 'Is Global Consequence Logically Revisionary?'. Unpublished.
- Williamson, T. 1990: *Identity and Discrimination*. Oxford: Basil Blackwell.
- 1994: *Vagueness*. London: Routledge.
- Wright, C. 1987: 'Further Reflections on the Sorites Paradox', in Keefe and Smith 1996, pp. 204–250. Originally published in *Philosophical Topics*, 15.
- Zalta, E. 1988: 'Logical and Analytic Truths that Are Not Necessary'. *Journal of Philosophy*, 85, pp. 57–74.