

Vague, So Untrue*

DAVID BRAUN AND THEODORE SIDER

Noûs 41 (2007): 133–156

According to an old and attractive view, vagueness must be eliminated before semantic notions (truth, implication, and so on) may be applied. This view — call it *semantic nihilism* — was accepted by Frege, but is rarely defended nowadays.¹ This recent neglect is unjustified: the thorny nest of problems surrounding vagueness is best untangled by accepting something like the old

*The authors formulated the main idea for this paper in 1996. Sider wrote a first draft in the Spring of 2001 and presented talks at various places, some under the title “Vagueness, Ambiguity and the Application of Logic”. Braun and Sider revised the paper significantly during 2002–2004. We would like to thank the following for helpful comments: JC Beall, Jiri Benovsky, Sylvain Bromberger, Cian Dorr, Delia Graff, Hilary Greaves, John Hawthorne, Mark Heller, Benj Hellie, Hud Hudson, Jeff King, Karson Kovakovich, Sarah-Jane Leslie, Europa Malynicz, Ned Markosian, Gail Mauner, Vann McGee, Brian McLaughlin, Daniel Nolan, Teresa Robertson, Jennifer Saul, Roy Sorensen, Jason Stanley, Gabriel Uzquiano, Ryan Wasserman, Brian Weatherson, Timothy Williamson, anonymous referees, the participants in Tim Maudlin’s seminar on truth, and audiences at Alabama, Alberta, the ANU, Boise State, Calvin, the Creighton Club, UC Davis, Kentucky, Leeds, Massachusetts, Massey, M.I.T., Miami, the New Jersey Regional Association conference, Oxford, the Princeton-Rutgers graduate student conference, Richmond, Southern Illinois at Edwardsville, Stockholm, Syracuse, Virginia, and Western Washington University.

¹Frege (1903, section 56); see also Williamson (1994, section 2.2). The recent literature does contain some related views. Mark Heller’s view is somewhat like ours (1990, chapters 3 and 4). Peter Unger (1979*a,b,c*) claimed that vague sentences like ‘I exist’ are untrue, but his view was far more radical than ours (e.g.: “our existing expressions, at least by and large, fail to make any contact with whatever is there” (1979*a*, p. 249)). See also Dummett (1975) and Wheeler (1975, 1979). An interesting case is David Lewis. Mostly he seems to endorse standard supervaluationism, but some intriguing remarks sound closer to semantic nihilism: “Super-truth, with respect to a language interpreted in an imperfectly decisive way, **replaces** truth *simpliciter* as the goal of a cooperative speaker attempting to impart information” (1993, p. 29 — our **boldface**). Sorensen (2002) suggests that supervaluationists should reject the identification of supertruth with truth, and should hold that no sentence is true. Perhaps he has something like semantic nihilism in mind. Vann McGee and Brian McLaughlin (1995) also reject the identification, but unlike us, try to use supertruth to capture part of our ordinary notion of truth. Closest of all is the view presented by Kirk Ludwig and Greg Ray (2002). Like us, Ludwig and Ray say that vague sentences are untrue, and that vagueness does not require abandoning classical logic. But there are many differences between our theories, the most crucial of which stem from their (implausible, in our view) denial of vagueness in semantic vocabulary. Consequently, they do not develop a theory of ignoring, which is crucial to our account (section 1.2), nor do they address the worries that the view is self-defeating (section 1.4) and inexpressible (section 3.2).

Fregean view.

If semantic notions such as truth apply only to completely precise sentences, they do not apply to English or any other natural language. Thus, almost no English sentences are true (or false). We defend this seemingly radical and self-refuting conclusion by developing a theory of how vagueness is typically and harmlessly *ignored*.

Section one sets out semantic nihilism. Section two argues for its superiority to the structurally similar theory of supervaluationism.² Section three concerns truth.

1. Semantic nihilism

1.1 Semantic indeterminacy and truth

An expression is vague if it can be unclear to a speaker informed of all relevant facts whether the expression correctly applies. Imagine a series of patches of color varying continuously from red to pink. A speaker who can see the patches clearly will nevertheless be unsure whether ‘red’ applies to certain intermediate patches.

Like many, we think that vagueness occurs when there exist multiple equally good candidates to be the meaning of a given linguistic expression. ‘Red’ is vague because there are many color properties that equally deserve to be expressed by ‘red’. On our usage, any sentence containing a vague word such as ‘red’ counts as vague, even if it does not concern a borderline case. We choose this usage because the multiplicity of candidate meanings for ‘red’ results in a multiplicity of propositions that equally deserve to be the meaning of such a sentence.³ We assume that the properties, relations, and propositions that are candidates for being the meanings of linguistic expressions are precise: any n -tuple of objects either definitely instantiates or definitely fails to instantiate a given n -place relation, and any proposition is either definitely true or definitely false.⁴ But the facts that determine meaning (for instance, facts about use, naturalness of properties, and causal relations between speakers and properties)

²We make no attempt at an exhaustive comparison with other rival theories. The literature contains powerful critiques of these theories (see especially Williamson (1994)), which provide our theory with ample motivation, if it can be adequately defended.

³The multiplicity of propositions follows automatically, given a “structured” conception of propositions.

⁴Our theory thus cannot capture so-called “worldly” vagueness, if it exists.

do not determine a unique property to be the meaning of ‘red’.⁵ There is no property that ‘red’ uniquely expresses, and therefore no unique proposition that a sentence containing ‘red’ expresses. Vagueness is a type of semantic indeterminacy.⁶

Many agree that vagueness results from a multiplicity of candidate meanings. Our distinctive claim concerns the impact of this multiplicity on semantics. The leading idea is that vagueness is a lot like ambiguity.⁷ To be either true or false, a sentence must have a unique meaning. Ambiguous sentences do not have unique meanings. Therefore, they are neither true nor false.⁸ Similarly, sentences containing vague expressions do not have unique meanings; therefore, they too are neither true nor false.⁹ Ambiguity is usually taken to be resolved in context by the intentions of speaker and audience. If vagueness could be eliminated by context then truth and falsity could indeed be achieved. But unlike ambiguity, vagueness is rarely (if ever) totally eliminated in context; and so on our view, utterances are rarely (if ever) true or false.

Most agree that vague sentences concerning “borderline cases”, for instance ‘Five piled stones are a heap’, are neither true nor false. Semantic nihilism goes much further: *no* sentences containing vague terms are either true or false, not even sentences that do not concern borderline cases. None of the following sentences is either true or false:

Zero piled stones are a heap.

⁵See Lewis (1983, 1984) and Devitt (1984, section 12.4) on content determination. We therefore oppose epistemicists like Sorensen (1988; 2001, chapter 11) and Williamson (1994, section 7.5), who hold that linguistic expressions generally do have unique meanings. We also oppose those like Fara (2000) and Soames (1999, chapter 7), who hold that vague expressions have meanings relative either to contexts, or sets of interests, that generate unique extensions. If the facts that determine meaning are not sufficiently “fine-grained” to secure sharp cutoffs, supplementing them with facts about the context will not suffice either.

⁶Similar points hold for mental items.

⁷Vagueness and ambiguity are merely analogous, not identical. The meanings of an ambiguous term “cluster” into a small number of disjoint groups which must be mastered individually by a competent speaker; in any context of use one has one of these clusters in mind; etc. And perhaps vagueness and ambiguity should be treated at different levels of linguistic analysis, the former being a property of linguistic expressions, the latter a property of sounds. (Thanks to Sylvain Bromberger for discussion.)

⁸This is the dominant view about ambiguous sentences and truth values. See Fine (1975, p. 284), for a contrary opinion, and Tye (1989) for a criticism of Fine.

⁹Do *any* sentences lack vagueness altogether? Perhaps mathematical sentences, or sentences constructed out of purely logical vocabulary, such as “Everything is self-identical” (uttered in a context in which the quantifier is unrestricted).

One million piled stones are a heap.

If five piled stones are a heap then five piled stones are a heap.

Five piled stones are a heap and it is not the case that five piled stones are a heap.

These sentences lack truth value for the same reason ambiguous sentences do: lack of a unique meaning.

1.2 Approximate truth and ignoring vagueness

We say that an utterance of ‘A man with no hairs on his head is bald’ is vague, and therefore untrue. ‘There exist some people’, ‘Snow is white’, and ‘George W. Bush was President of the United States in 2002’ are all untrue. Despite this, we do not think of ourselves as defending a radical view. We do not recommend ceasing to speak vague languages.¹⁰ We do not view vagueness as a defect, as Frege did.¹¹ For all we say, vagueness may be essential for language to be useful.¹² We do not recommend wholesale changes in linguistic practice; our semantic nihilism is intended as a rational reconstruction of existing practice. On our view, ordinary speakers typically and harmlessly *ignore* vagueness. And when doing so, it is reasonable to speak, in a sense to be defined, the *approximate truth*.

While the facts that determine meaning do not determine a unique proposition as being meant by a vague sentence, they are not entirely impotent. There is typically a cloud of propositions in the neighborhood of a sentence uttered by a vague speaker. Vagueness prevents the speaker from singling out one of these propositions uniquely, but does not banish the cloud. Speaking vaguely (as always), there is a range of *legitimate disambiguations* for a vague expression. These are objects, properties, relations, or propositions, depending on whether that expression is a name, monadic predicate, polyadic predicate, or sentence (we ignore other grammatical categories for simplicity). When all the legitimate disambiguations of a sentence are true, call that sentence *approximately true*.¹³ An ordinary utterance of ‘A man with no hairs on his head

¹⁰See Unger (1979a).

¹¹See Dummett (1981, pp. 32–35).

¹²See Wright (1975, p. 330, p. 335).

¹³Calling a sentence “approximately true” may misleadingly suggest that it uniquely expresses a proposition that is “close to” or “similar to” a true proposition. But sentences that

is bald' is approximately true, despite failing to be true.¹⁴

Truth is usually thought to play a pervasive role in our cognitive lives. Truth (or perhaps known truth) is the goal of successful inquiry: it is what a diligent inquirer must strive to accept. Truth (or perhaps known truth) is the norm of assertion: it is what a cooperative speaker must strive to communicate.¹⁵ Given our picture of semantics, truth is an impossible standard that we never achieve. But it would be pointlessly fussy to enforce this standard to the letter, requiring the (exact) truth. It would rarely be possible to live up to this standard, nor would it be desirable to try, for the differences between the legitimate disambiguations of our sentences are rarely significant to us. As a result, it is usually harmless to *ignore* vagueness, set it aside, and act as if one's sentence is not vague, but rather expresses a unique proposition. When vagueness is being ignored, the cooperative communicator satisfies her communicative obligations well enough by uttering sentences that are approximately true; a diligent inquirer satisfies her intellectual obligations well enough by accepting sentences that are approximately true. For in such cases, the differences between the legitimate disambiguations do not matter, and each is true. A typical speaker will, for instance, ignore the vagueness in 'A man with no hairs on his head is bald' and treat that sentence as true, since it is approximately true. In so doing she would not, strictly speaking, be living up to her obligation to speak only the truth; but since the sentence is approximately true, she would closely approximate satisfaction of that obligation — as closely as anyone ever does.

Likewise, cooperative speakers are usually thought to be obliged to *deny* sentences only when they are false or suffer from some other semantic defect,

are approximately true in our technical sense need not express a single proposition uniquely. Moreover, some sentences that are approximately true in the ordinary sense are not approximately true in our technical sense, for example 'The number of pennies in the jar is 256,291', said of a jar containing 256,292 pennies.

¹⁴Note the parallel with the supervaluationists. Our legitimate disambiguations and approximate truth are structurally similar to their precisifications and supertruth. As with precisifications, the range of legitimate disambiguations of an expression may well vary from context to context. Moreover, following Fine (1975, p. 276), an assignment of legitimate disambiguations to an entire language must coordinate what it assigns to distinct terms. 'Bob Dylan is identical with Robert Zimmerman' should turn out approximately true; thus, 'Bob Dylan' and 'Robert Zimmerman' cannot be simultaneously disambiguated differently. Despite these structural similarities between our theoretical tools, we put those tools to a very different *use* than do supervaluationists. Most importantly, we do not identify supertruth with truth. See below.

¹⁵At least one of us is suspicious of talk of "the norm of assertion". We might speak instead of various goals speakers have, one of which is to communicate truths.

such as presupposition-failure or (other) lack of truth value. On our view, when vagueness is ignored, this norm is satisfied well enough if the denied sentence, or its presupposition, is approximately false — that is, false on all legitimate disambiguations — or the sentence suffers from some failure of truth value unrelated to vagueness.¹⁶

Ordinary speakers do *not* ignore the vagueness of sentences that concern borderline cases, for instance, ‘Five piled stones are a heap’. Here it is evident to any competent speaker that the sentence is untrue because of its vagueness.¹⁷ ‘Heap’ would have to be more precise than it in fact is in order for the sentence to be true. Speakers thus refrain from uttering or accepting this sentence. For the same reason they refrain from uttering its negation, ‘Five piled stones are not a heap’.¹⁸ Suppose that, despite the manifest vagueness and untruth of ‘Five piled stones are a heap’, a speaker decided to ignore its vagueness and utter it anyway. That speaker would then badly fail in her duty as a cooperative communicator, for she would not even be approximately satisfying the norm of assertion to speak only the truth: while some legitimate disambiguations of her sentence are true, others are false.

“Truth is the norm of assertion” is an oversimplification. As Grice (1989) points out, cooperative speakers sometimes do not strive to speak the truth. A speaker may ironically utter ‘Smith is a fine friend’ despite not accepting that sentence, expecting his hearer to accept the sentence ‘Smith is *not* a fine friend’. On Grice’s view, the speaker implicates a truth even though the sentence uttered is false. Consider also the phenomenon of “loose talk” explored by relevance theorists.¹⁹ A cooperative speaker may utter ‘His face was square’, believing that this sentence is false, but intending to communicate some other true sentence. Grice and the relevance theorists thus provide a refined norm of assertion, according to which speakers are obliged to *communicate* truths in a broad sense. On our view, the sentences communicated in this broad sense are not true. Nevertheless, the communicated sentences are often approximately

¹⁶When vagueness is *not* ignored, a speaker might deny a sentence she takes to be untrue because of vagueness.

¹⁷This will be justified on nearly any proto-theory of vagueness the speaker may have internalized — the speaker need not implicitly accept semantic nihilism.

¹⁸They may *deny* the sentence, and may also deny its negation, and may even express these denials with the seemingly contradictory sentence ‘Five piled stones is not a heap, and it’s not *not* a heap either’. Thanks to Brian Weatherston for this example. This may be an instance of metalinguistic negation. See Horn (1989).

¹⁹See Sperber and Wilson (1986).

true. Thus, speakers who ignore vagueness closely approximate satisfaction of the refined norm by communicating approximate truths.

We use the term ‘ignore’ in a technical way. To ignore the vagueness of an expression is to fail to take account of its vagueness and the impact of its vagueness on its truth. When ignoring the vagueness of a sentence, a speaker uses it (and accepts or rejects it) as if it both lacked vagueness and had a truth value. Ordinary speakers usually do not think about vagueness at all, and therefore usually ignore vagueness, in our technical sense. We call this *unconscious* ignoring. Even philosophers regularly fail to think about vagueness, and so in daily life regularly unconsciously ignore vagueness.

Not all ignoring is unconscious. A theorist can consciously realize that ‘chair’ is vague, and say ‘There is a chair in the next room’, though he knows full well that its vagueness deprives it of truth. He might utter it because he is speaking to hearers who are not aware of the impact of vagueness on truth. He might utter it consciously trying to approximate the norm of uttering truths. His uttering the sentence may lead his hearers to accept it, which is a result he may well desire, since he knows that it is impractical to get them to accept (exactly) true sentences. His uttering the sentence may also lead his hearers to perform desirable actions, such as fetching a chair. Altogether, uttering the sentence is likely to be preferable to lapsing into complete silence. Such a theorist is engaging in what we call *conscious* ignoring.²⁰ Of course there are intermediate cases — for instance, people who know that a term is vague but do not share our view about the nature of vagueness. We still count them as ignoring vagueness, for they use, accept, and deny sentences as if their vagueness did not undermine their truth.

We close this section with some final remarks about the relationship between our theory and pragmatics. Our theory is not intended to be a “pragmatic theory of vagueness”. Ignoring, in our sense, should not be subsumed under any of the standard mechanisms posited by semanticists or pragmaticists to account for communication in context. Ignoring vagueness is merely using sentences as if they were not vague (and so as if they were not truth-valueless). This is usually done unconsciously, by speakers who are oblivious to vagueness. Nor is there any need for theorists like us to exploit any of the standardly discussed pragmatic mechanisms, when we shift from ignoring vagueness to attending to

²⁰Like the teacher who ignores the raised hand of a persistent student, one can decide to ignore vagueness, in our sense, even when vagueness has been made salient, contrary to Lewis’s (1996) Rule of Attention.

vagueness, and then back again. Semanticists and pragmaticists can theorize in peace, without taking account of our theory (and our theory need not bend to fit theirs). These theorists usually ignore vagueness, and this is harmless. For as we have explained in this section, when vagueness is being ignored, the assertion and communication, by semantic or pragmatic mechanisms, of approximate truths closely approximates the assertion and communication of truths that semanticists and pragmaticists describe.

1.3 Higher order vagueness

An expression is vague if it can be unclear to a speaker informed of all the relevant facts whether it applies. An expression exhibits *higher-order vagueness* if it can be unclear to such a speaker whether it clearly applies. It is evident that higher-order vagueness exists. Given a series of patches continuously varying in color from clearly red to clearly non-red, it would be no easier to pick out the last clearly red patch than to pick out the last red patch.

One might suspect that higher-order vagueness refutes semantic nihilism. We say that there is a range of legitimate disambiguations for 'red'. So there must exist a unique *set* of legitimate disambiguations of 'red'. So, the objection continues, there must be an abrupt transition between the patches that fall under each legitimate disambiguation of 'red' and those that fall under only some legitimate disambiguations of 'red'. But then there would be an abrupt transition between the clearly red patches and those patches that are neither clearly red nor clearly non-red. For, surely, the clearly red patches are those in the extension of every legitimate disambiguation of 'red', and the patches that are neither clearly red nor clearly non-red are those patches that are in the extension of some, but not all, of the legitimate disambiguations of 'red'.²¹ And even if 'clearly' and 'legitimate disambiguation' are not thus connected, it would be intolerable to admit an abrupt transition between patches that fall under each legitimate disambiguation of 'red' and those that do not, for 'legitimate disambiguation' is surely vague. The facts that determine meaning no more determine a unique meaning for 'legitimate disambiguation' than for 'red'.²² Either way, semantic nihilism must be incorrect.

²¹Some assumption of this sort connecting 'clear' and 'legitimate disambiguations' seems plausible; however, we take no stand on the semantics of 'clear' and 'clearly'. We mean to use these expressions in their ordinary senses, in order to give intuitive characterizations of vagueness and higher-order vagueness.

²²Ludwig and Ray (2002) do not agree that semantic vocabulary is vague; they claim that

The objection overlooks the ways in which vagueness can be ignored. When presenting our theory, we did indeed write sentences similar to (L).

(L) There is a set of legitimate disambiguations of ‘red’.

But we were then ignoring the vagueness in ‘legitimate disambiguation’. Therefore, (L) is vague, and so untrue, by our lights.

So, we object to the above criticism when it makes use of (L). Granted, (L) is approximately true; that is why we asserted it while ignoring the vagueness of ‘legitimate disambiguation’. But that expression is vague, and this vagueness may no longer be ignored when (L) is put forward as being inconsistent with higher-order vagueness. Its vagueness must then be attended to, and (L)’s lack of truth acknowledged.

Our theory applies to higher-order vagueness in the same way that it applies to first-order vagueness. Higher-order vagueness arises from vagueness in expressions used to talk about vagueness. These expressions include the theoretical expression ‘legitimate disambiguation’ as well as (certain uses of) ‘clearly’, ‘definitely’, and other ordinary English expressions. On our view, these expressions are indeed vague: they have many legitimate disambiguations. Sentences containing those expressions are therefore untrue. Nevertheless, in many cases either they or their negations can be asserted, namely, when their vagueness is ignored and their legitimate disambiguations are uniformly true or false. But in other cases their legitimate disambiguations are a mixture of true and false. These include cases that can be thought of as “borderline cases of borderline cases”.

1.4 Self-defeating?²³

At this point we may appear to be in a bind. We say that vague sentences are untrue. But the key sentences in this very paper are vague, and so are untrue by our lights.

‘applies’ (as in “The term ‘bald’ applies to Bob”) is precise. (This claim is crucial to their attempted proof of semantic nihilism (2002, appendix).) In our view, their claim is incredible. In a sorites series of gradually varying color patches, how could the facts that determine meaning secure a first patch of which one can say that ‘red’ applies to it? And if, contrary to appearances, meaning-determining facts are really that powerful, why not go all the way to epistemicism?

²³See Williamson (1994, section 6.2).

We admit, for instance, that ‘legitimate disambiguation’ is vague. So we must admit that ‘approximately true’ is vague as well, since it was defined in terms of ‘legitimate disambiguation’. We must further admit that the sentential truth predicate is vague, for a sentence is true (in English) iff the proposition it expresses (in English) is (propositionally) true, and ‘expresses (in English)’ is clearly vague. Thus, we must admit the untruth of each of the following sentences:

- (o) ‘A man with no hairs on his head’ is true on all legitimate disambiguations
- ‘A man with no hairs on his head is bald’ is approximately true.
- ‘A man with no hairs on his head is bald’ is not true

But we wrote sentences like each of these in section 1.2. Indeed, we must acknowledge that all the sentences in this paper are untrue, since all of them are vague.²⁴ There is no getting around this: the theory applies to its own statement just as surely as to sentences about heaps of stones and bald men.

We make theoretical assertions and then take them back when attending to their vagueness — what of it? We do the same for nontheoretical assertions. In ordinary life we are happy to assert ‘A man with no hairs on his head is bald’. Later, we attend to the vagueness in ‘bald’, acknowledge that the ordinary sentence is untrue, but defend the propriety of having asserted it in an ordinary context by pointing out that it is approximately true. We say the same about our section 1.2 assertions of (o). The sentences in (o) are untrue, but it was proper for us to assert them since we were ignoring their vagueness (more on this below). The sentences in (o) are approximately true, after all. (Of course, that very claim is untrue, only approximately true! Etcetera!)

This all has an air of paradox about it. Some will see in this an objection to semantic nihilism: since it entails its own untruth, it is untrue, and therefore unacceptable. Notice that the objector cannot stop with the claim that the theory is untrue, for we freely grant *that*: on our view, nearly all the sentences uttered and accepted by successful communicators and inquirers are untrue. Thus, the final part of the objection, that our theory’s unacceptability follows

²⁴We ignore the vagueness of ‘Everything we say is vague’ and ‘Everything we say is untrue’ when we utter them. So we are, in effect, restricting the domain of quantification to everything *else* that we say. Of course, we can later reflect on those sentences and say ‘They too are vague and untrue’. (Thanks to Timothy Williamson for discussion.)

from this, is crucial. But as we said above, responsible communicators and inquirers generally ignore vagueness, uttering and accepting sentences that are only approximately true. That is what we are doing in writing this paper.

Of course, we are not ignoring *all* vagueness in writing this paper. The paper is about vagueness, after all. When we write the sentence:

(oa) ‘A man with no hairs on his head is bald’ is approximately true.

for instance, we are attending to the vagueness in ‘bald’, ‘man’, ‘hair’, ‘head’ and ‘on’, while ignoring the vagueness in ‘approximately true’ and the quotation name of the sentence ‘A man with no hairs on his head is bald’. Now, a challenger may point out that the latter terms are vague, and so (oa) is untrue, on our view. In response we could be uncooperative and continue to ignore this vagueness. Alternatively, we could play along, cease ignoring that vagueness, and so cease uttering (oa). In defense of our earlier use of (oa), we would point out that even though (oa) is not true, we were ignoring (parts of) its vagueness, which was harmless since (oa) is approximately true. In thus uttering the sentence:

Sentence (oa) is approximately true

we would of course be ignoring the vagueness in ‘(oa)’ and (the final occurrence of) ‘approximately true’. In response to a further challenge, we might cease to ignore this, and so cease to utter ‘(oa) is approximately true’. In defense of having uttered it we would note that it is approximately true. There might then be a further challenge. And so on.

Those who continue to fear paradox may find a theological illustration helpful.²⁵ We are blocked from stating our theory truly because we cannot eliminate the vagueness in semantic words like ‘refer’, ‘express’, and ‘legitimate disambiguation’. But God could arbitrarily choose legitimate disambiguations of these words and state a version of our thesis accordingly. God would say: “You humans_i never uniquely refer_j to anything with your names_k and never uniquely express_l any properties with your predicates_m; therefore you never uniquely express_l any propositions”. This utterance of God’s would express a true proposition, if we are on the right track. The appeal to God is vivid but inessential: the point is that there exist true propositions at which this paper gestures, even though we have no way to express one uniquely. The differences among them are not especially important, no more important

²⁵Thanks to Hud Hudson here.

than the differences among the legitimate disambiguations of a typical vague sentence such as ‘There exist some chairs’. It would be picayune to insist on speaking the truth with such a sentence. Likewise for the insistence that we speak the truth in this paper. (More on this in section 3.)

1.5 Validity

The now-standard conception of validity is model-theoretic. A sentence is valid iff it is true in all models. An argument is valid iff its conclusion is true in any model in which its premises are all true. But consider an *ambiguous* argument (and ignore vagueness for the moment):

- (A) All bats have wings.
Barry Bonds swung a bat.
Therefore, Barry Bonds swung something with wings.

Blind application of the model-theoretic notion of validity to (A) yields the conclusion that it is valid. But in some *tokens* of (A) the premises are true and conclusion is false — these are cases in which the speaker shifts mid-argument between different senses of the ambiguous word ‘bat’. This result clashes with the ordinary notion of validity. Similarly, blind application of the model-theoretic notion of validity to (B) yields the conclusion that it is valid, even though some of its tokens are false (those with different disambiguations of the tokens of ‘bat’):

- (B) If all bats have wings, then all bats have wings.

We should rethink the application of model-theoretic validity to ambiguous arguments.

As students, we are told to *disambiguate* an argument (or sentence) before evaluating for validity. Replacing each occurrence of ‘bat’ in (A) with either ‘chiropteran’ or ‘baseball bat’ yields four disambiguated arguments. The evaluation of each is straightforward: some are invalid, some have false premises; none is sound. If asked “Is the *original* argument (A) valid or invalid?”, we could reasonably regard the question as confused. Just as ambiguous sentences are not true or false *simpliciter*, only true or false under a disambiguation, ambiguous arguments are not valid or invalid *simpliciter*, only valid or invalid with respect to disambiguations. Similar points hold for sentences such as (B).

Return now to vagueness. Given semantic nihilism, similar difficulties arise in applying the standard model-theoretic notion of validity to arguments and sentences, for instance:

- (A_v) All birds have wings.
Tweety is a bird.
Therefore, Tweety has wings.
- (B_v) If all birds have wings, then all birds have wings.

(B_v) is model-theoretically valid, even though it contains vague expressions, and so fails to express a unique proposition, and so on our view fails to be either true or false. (A_v) is also valid on the standard model-theoretic conception, even though its premises and conclusion fail to express unique propositions, and so fail to be true or false. This should make us hesitate to apply the standard, unrelativized, model-theoretic notion of validity to vague arguments and sentences. We should, instead, treat vague arguments and sentences like ambiguous arguments and sentences. Though vague sentences are neither true nor false, they are true or false with respect to legitimate disambiguations. Therefore, we should assess the validity of vague arguments with respect to legitimate disambiguations.²⁶

We can, however, introduce an *unrelativized* notion of validity for vague arguments. (We could do the same for ambiguous arguments as well.) Say that an argument is “valid *simpliciter*” iff it is valid with respect to all *uniform* legitimate disambiguations — disambiguations in which all occurrences of a given word receive the same legitimate disambiguation throughout the argument. (A_v) is valid *simpliciter* in this sense. We can thus achieve a non-relativized notion of validity for a vague language, but only by detouring through the notion of validity with respect to a legitimate disambiguation. In fact, all arguments that are classically valid will be valid *simpliciter* in this new sense.²⁷

²⁶More carefully: for any assignment f of legitimate disambiguations to the occurrences of vague words in an argument, consider the “sentences” that result from the sentences in the argument by replacing the occurrences with ordered pairs of those expressions and their legitimate disambiguations under f . Assessing the validity of the argument relative to f means assessing the validity, model-theoretically, of this new argument, treating the ordered pairs as the words (non-logical constants) of its “sentences”. *Mutatis mutandis* for sentence-validity.

²⁷The sentence ‘All birds are birds’ is valid *simpliciter*, but not true, because it is vague. Thus our notion of validity *simpliciter* differs from the supervaluationist’s notion of validity, which entails that all valid sentences are true (that is, supertrue). Ludwig and Ray (2002, pp. 444–446) also invoke the concept of validity under all precisifications.

Validity *simpliciter* preserves approximate truth: if the premises of a valid *simpliciter* argument are approximately true then its conclusion will be approximately true as well.²⁸ Typical applications of logic concern arguments that are valid *simpliciter* and contain approximately true premises whose vagueness is being ignored; in such cases, the conclusions are approximately true, which is all we can reasonably demand.

A core claim of this paper has been that vagueness is usually ignored. In section 1.2 we saw that untrue *sentences* whose vagueness is ignored may be reasonably asserted; they need only be approximately true. Now we have seen that *arguments* may be put forward while ignoring their vagueness. For their vagueness is harmless: if they are valid *simpliciter* then they preserve approximate truth. Since classically valid arguments are valid *simpliciter*, classical logic suffices for ordinary reasoning in which vagueness is ignored.

1.6 The sorites

Classical logic apparently leads to trouble in sorites reasoning:

1. A “pile” of a single stone is not a heap
2. If a “pile” of a single stone is not a heap, then a pile of 2 stones is not a heap
3. If a pile of 2 stones is not a heap, then a pile of 3 stones is not a heap
-
-
100. If a pile of 99 stones is not a heap, then a pile of 100 stones is not a heap

C. Therefore, a pile of 100 stones is not a heap

Intuitively true premises lead validly to an intuitively false conclusion. What went wrong?

In the previous section we upheld the use of valid *simpliciter* arguments when vagueness is ignored. But the sorites argument is precisely a case in which vagueness cannot be ignored, since many of its premises manifestly

²⁸Assuming there exists at least one uniform legitimate disambiguation of premises and conclusion.

concern borderline cases. It is, after all, the vagueness in ‘heap’ that leads to the trouble! When we attend to vagueness, we must treat a vague argument like an ambiguous one. An ambiguous argument is neither sound nor unsound *simpliciter*, for its premises are neither true nor false *simpliciter*, only true or false relative to disambiguations. To evaluate for soundness we must first disambiguate; *before* disambiguation there is no question of soundness. When we attend to its vagueness we will say the same about the sorites argument. Before disambiguation there is no question of soundness, for each premise before disambiguation lacks a unique meaning, so lacks a truth value, and so, *a fortiori*, is untrue. The argument is valid with respect to some legitimate disambiguations, but not all. It is valid with respect to all *uniform* legitimate disambiguations, but all uniform legitimate disambiguations of the argument contain a false premise.

Press a non-philosopher on exactly when a non-heap turns into a heap, and she will invariably reply, “It depends on what you mean by ‘heap’”. Confront her with the sorites, and she will say: “Define your terms. Tell me what ‘heap’ means, and I’ll tell you which premise of the argument is false.” Semantic nihilism vindicates these reactions. In a disambiguated argument in which ‘heap’ uniformly means *pile of at least 40 stones*, the premise ‘If a pile of 39 stones is not a heap then a pile of 40 stones is not a heap’ is false; the rest are true. A disambiguated argument in which ‘heap’ uniformly means *pile of at least 41 stones* has a different false premise. A disambiguated argument in which ‘heap’ has different meanings in different places may have all true premises, but it is invalid.²⁹

This natural resolution of the sorites paradox is, we feel, one of our theory’s most attractive features. We are taught not to resolve the paradox in this way — not to treat vague arguments like ambiguous ones — because unlike ambiguous sentences, vague sentences are thought capable of truth-evaluation *before disambiguation*. Whereas most of us are happy to say that the sentences in argument (A) lack truth value before disambiguation, few will say this about vague sentences since nearly all sentences we typically utter are vague. If vague sentences cannot be true then almost nothing we ever say is true. But as we have explained, this consequence is acceptable.

²⁹This is not to say that speakers are capable of expressing (completely) disambiguated arguments. Unlike Fara (2000) and Soames (1999, chapter 7), we think that, even taking the context of utterance and the interests of speakers into account, speakers’ utterances have many legitimate disambiguations.

The sorites argument we have considered contains a series of conditionals, but one can formulate sorites arguments in other ways, for example with a quantified conditional, a series of disjunctions, or a series of negated conjunctions. Concerning *soundness*, we say the same about each of these as about the series-of-conditionals sorites: before disambiguation none of the premises is true; disambiguation results in either invalidity or a false premise.

Accounting for speakers' *intuitive judgments* about the premises of sorites arguments is a separate and difficult task. Interestingly, speakers react differently to different versions of the argument. Speakers do *not* judge the premises of the series-of-disjunctions sorites to be true. In the premise "Either a pile of five stones is a heap or a pile of six stones is not a heap", neither disjunct seems true (as a consequence of vagueness), and a disjunction of untruths seems untrue — contrary to what some supervaluationists say. Contrast the series of negated conjunctions, which has perhaps the most intuitively appealing premises. It is easy to see why speakers think that the first negated conjunction in the argument is true: i) it is natural to ignore the vagueness of 'A pile of one stone is not a heap' and 'A pile of two stones is a heap'; ii) the second is approximately false; so, the first premise appears to be the negation of a false conjunction. But why do speakers tend to think that premises concerning the border of heapdom, such as the following, are true?

($\sim\&$) Not: a pile of 5 stones is not a heap and a pile of six stones is a heap.

We suspect the following: speakers attend to the vagueness of the constituent sentences, and so judge that one or both are not true, and therefore judge that the conjunction is not true. They then express this by uttering ($\sim\&$).³⁰ We suspect that similarly mistaken (though understandable) reasoning lies behind the intuitive appeal of the borderline premises in the series-of-conditionals:

(\rightarrow) If a pile of five stones is not a heap, then a pile of six stones is not a heap.

Speakers think that in order for (\rightarrow) to be false, the embedded conjunction of ($\sim\&$) must be true. They see that this conjunction is not true, and conclude that (\rightarrow) is true.

³⁰Perhaps the negation is metalinguistic (Horn, 1989), or perhaps the speakers are just making a mistake.

2. Supervaluationism

We have given an account of vagueness motivated by the idea that vagueness is semantic indeterminacy. But another popular view, supervaluationism, purports to do the same, without implying the conclusion that no vague sentence has a truth value.

Like our theory, supervaluationism makes use of legitimate disambiguations of vague sentences; the supervaluationist calls these ‘precisifications’. We call truth-on-all-legitimate-disambiguations ‘approximate truth’, and make use of approximate truth in explaining language use. The supervaluationist calls truth-on-all-precisifications ‘supertruth’, and assigns it a somewhat similar explanatory role. But the supervaluationist goes further, and identifies supertruth with truth. A vague sentence, for the supervaluationist, is true iff true relative to each of its precisifications. Likewise, a sentence is false iff *superfalse* — false on all precisifications.

The additional claim that supertruth is truth is more trouble than it is worth. The observation that ignoring semantic indeterminacy is harmless (because typical sentences we treat as true are true on all legitimate disambiguations) provides an adequate explanation of our use of sentences; there is no need to add that truth on all legitimate disambiguations is *truth*. Indeed, that addition leads to trouble. The trouble is well-known; what is news is that the trouble does not come from the use of legitimate disambiguations in theorizing about vagueness, only from the identification of truth-on-all-legitimate-disambiguations with truth. Finally, some question the relevance of truth under precisifications, which are admittedly not actual semantic values of vague sentences, to a sentence’s actual truth.³¹ We preserve what is right about supervaluationism while discarding the problematic baggage.

Supervaluationism’s apparent advantage over semantic nihilism is that it allegedly assigns truth values that match the truth values ordinary speakers intuit (factual error and pragmatic noise aside). ‘A man with no hair on his head is bald’ is supertrue, so true. But standard objections to supervaluationism show that supertruth does not match intuited truth in all cases. Famously, supervaluationism implies that the sentence:

- (\exists) There is some number, n , such that n piled stones is not a heap whereas $n + 1$ piled stones is a heap

³¹Cf. Sanford (1976, section 6); Fodor and Lepore (1996).

is true, since it is supertrue — each precisification of ‘heap’ sets a precise cutoff somewhere. But (\exists) does *not* seem true.³² Likewise, said of a borderline pink/red patch of color, the following sentence turns out true:

(\vee) The patch is pink or the patch is red

Any precisifications of ‘pink’ and ‘red’ place the patch in the extension of one or the other. The sentence is therefore supertrue, and hence true simpliciter, given supervaluationism. But intuitively, (\vee) does not seem true. One is inclined to ask which disjunct is true, and the question cannot be answered, nor is there any acceptable explanation of why we would be ignorant of the answer. Some supervaluationists disagree, claiming that the intuitive truth of sentences like (\vee) is a datum that supports supervaluationism. We reject the alleged datum: (\vee) does *not* seem true, and does not seem typically to be assertable. The reader should consult his or her semantic intuitions here. We suspect most will agree with us.

Sentences (\exists) and (\vee) bring up a second objectionable feature of supervaluationism. Neither disjunct of (\vee) is supertrue. ‘The patch is pink’ is not supertrue, since there are precisifications on which it is false. Likewise, ‘The patch is red’ is not supertrue. But as noted, (\vee) itself *is* supertrue. Similarly, although the existential sentence (\exists) is supertrue, it has no supertrue instance; for no numeral α is the sentence

α piled stones is not a heap whereas $\alpha+1$ piled stones is a heap

supertrue. These consequences show that the formal properties of supertruth clash with our ordinary conception of truth. That conception requires that a true existential have a true instance (when there is a term for each member of the domain), and that a true disjunction have a true disjunct.

Our first objection concerning (\exists) and (\vee) must be reconciled with our use of the concept of approximate truth, which after all is essentially the concept of supertruth. We must admit that (\exists) and (\vee) are approximately true. Moreover, we use the concept of approximate truth to characterize satisfaction of the norm of assertion, attainment of the goal of inquiry, and the like. Our explanation of the propriety of uttering ‘A man with no hair on his head is bald’ appealed to the approximate truth of this sentence. This explanation threatens to also

³²Cf. Sanford (1976, section 6); Williamson (1994, section 5.4).

imply the propriety of uttering (\exists) and (\forall) . Yet it is *not* proper to utter (\exists) or (\forall) .

But our explanation of the propriety of uttering ‘A man with no hair on his head is bald’ was not *merely* that this sentence is approximately true. The explanation was that vagueness is typically *ignored* when uttering this sentence; and when the vagueness of a sentence is ignored, and the sentence is approximately true, one’s utterance of it acceptably approximates satisfaction of the norm of assertion (since the differences among its disambiguations are unimportant.) This explanation is blocked in the case of (\exists) and (\forall) , for they *draw attention to their vague nature*.

An utterance of (\forall) , concerning a clearly visible borderline pink/red patch, normally prompts the question “Which is it, pink or red?” The speaker cannot answer this question, and sees that her failure is not due to ignorance. This makes the vagueness of ‘pink’ and ‘red’ obvious. She will then go through something like the following reasoning: “Clearly, neither disjunct of (\forall) is true. A disjunction is true only if it has a true disjunct. So, (\forall) is not true. Therefore, I will not assert it.” Likewise, a normal utterance of (\exists) naturally prompts the question “Which number marks the boundary between heaps and non-heaps?”. The number cannot be located, and it is perfectly obvious that this is due to vagueness in ‘heap’ and not ignorance. This makes the vagueness of ‘heap’ manifest. The speaker then engages in the following pattern of reasoning. Given the manifest vagueness of the component open sentences of (\exists) :

n piled stones is not a heap

$n+1$ piled stones is a heap

there is no number to assign to ‘ n ’, relative to which each is true. (\exists) is true only if there is such a number. So, (\exists) is not true. Consequently, the speaker does not assert (\exists) . Therefore, unlike the vagueness in ‘A man with no hair on his head is bald’, the vagueness of (\exists) and (\forall) is *not* typically ignored, and hence their approximate truth does not imply their assertability.

We do *not* endorse the rule that all approximately true sentences are assertable. The rule is simpler: Assert the truth! One may acceptably fall slightly short of strict adherence to this rule by ignoring vagueness and asserting approximate truths, but the vagueness of sentences like (\exists) and (\forall) is hard to ignore.

Hard, but perhaps not impossible. One *can* cook up cases in which the vagueness of (\exists) and (\forall) is natural to ignore, and in these cases our resistance

to the sentences is decreased. Suppose James Bond has been captured by Goldfinger, who chooses a typically creative and inefficient method for killing Bond. Bond must draw a card, and his method of death depends on the color of the card, Goldfinger says. If it is blue then Bond will be drowned in the ocean; if it is brown then Bond will be buried alive. Red means he will be gored to death by a bull; pink means he will be forced to drink poison. But Bond has a chance: he has an antidote for the poison and he is actually a skilled bullfighter. He awaits the card, and look: it comes up borderline red/pink! Bond is relieved: *the card is either pink or red*, and either way he will be saved. Here the sentence (V) does not sound so bad if the focus is just on the fact that in either case Bond will survive. But the acceptability of the sentence can easily be dispelled, for example by asking how Goldfinger will decide whether the card is red or pink.³³

As mentioned, supervaluationists purport to vindicate ordinary speakers' intuitions about the truth values of sentences. One major advantage of their view, they think, is that sentences that express definitional or analytic connections (*penumbral* connections, in Fine's terminology) turn out true.³⁴ Consider, for instance, (Bald).

(Bald) If x is bald, and y has fewer hairs than x , then y is bald.

Since this sentence is partially definitive of 'bald', it constrains the simultaneous precisification of 'has fewer hairs than' and 'bald': no precisification can include u in the extension of 'bald' and include $\langle v, u \rangle$ in the extension of 'has fewer hairs than', but fail to include v in the extension of 'bald'. (Bald) turns out true under all precisifications, so supertrue, so true.

³³Ryan Wasserman pointed out that Bond may be willing to say this because he is assuming that Goldfinger has some particular classification of the card in mind; he may be taking 'pink' to mean 'pink as classified by Goldfinger'. But Bond might still utter the disjunction even if he doesn't take Goldfinger to have decided on some cutoff.

Example for (∃). An immortal god has committed two crimes, and is given a double sentence: he must first take some boulders from one valley and make a giant heap of them in another valley miles away; then he must count every grain of sand on a certain beach. At first he is daunted by the enormity of the tasks, but his friends console him. "You are immortal", they remind him, "and what is a long time compared with forever? The first task is finite. Just keep at it, piling up boulders one by one. Though it is hard to imagine this now, one day you will put a stone on the pile, and though it was not a heap before that stone it will be afterwards. And then all that will remain will be the finite task of counting the grains of sand on the beach."

³⁴Fine (1975, sections 1 and 2). We discuss such sentences in the next section.

The T-schema has a similarly strong claim to being analytic and constitutive of truth:

T-schema: ‘ ϕ ’ is true iff ϕ

One might therefore expect that all of its instances would be true given supervaluationism. But this is not the case. Return to the borderline red/pink patch, and consider instance (T) of the T-schema.

(T) ‘The patch is red’ is true iff the patch is red

The right-hand side of (T), namely ‘The patch is red’, is true on some precisifications and false on others. Therefore it is neither supertrue nor superfalse, and hence neither true nor false. The left-hand side of (T), however, is just plain false, since ‘The patch is red’ is not true, as just noted. Therefore, (T) is not true, given supervaluationism.³⁵

Thus the supervaluationist’s claim that his theory vindicates our intuitions about all (seemingly) analytic sentences is mistaken. Moreover, the failure to validate the T-schema, which is apparently definitive of truth, calls into question the supervaluationist’s identification of truth with supertruth.³⁶

3. More on truth

3.1 Truth, ‘true’ and the T-schema

What is the status of truth, ‘true’, and the T-schema on our theory? The short answer: sentential truth is unique expression of a true proposition. ‘True’ expresses sentential truth itself, not approximate truth. As a result, any instance of the T-schema that expresses a unique proposition is true. Any instance that fails to express a unique proposition fails to be true (or false), but is approximately true and may be asserted when vagueness is ignored. We consider these points in order.

As we saw in the previous section, we differ from the supervaluationists by refusing to identify truth with approximate truth (supertruth). Rather, a sentence is true iff it expresses a unique proposition, which is true. Naturally,

³⁵(T) is a material biconditional. Supervaluationism implies that such sentences are not truth functional, but nevertheless are untrue when one constituent is untrue and the other is false.

³⁶Williamson (1994, sections 5.7 and 7.2) among others has pressed this objection against supervaluationism and other theories that deny bivalence.

then, we hold that the sentential truth predicate does not express approximate truth; instead, it expresses the property of expressing a unique proposition, which is true.

Thus, instances of the T-schema that express a unique proposition are true, for example, an instance in a completely precise ideal language (spelled in capitals):

‘The PATCH is RED’ is TRUE iff the PATCH is RED.

But virtually no instances of the T-schema in ordinary English are true. A vague instance like (T) fails to express a unique proposition, and so fails to be true. (T) is not unique in this respect, on our view, for nearly all sentences fail to express unique propositions and so fail to be true.

Nevertheless, (T) is true with respect to each of its legitimate disambiguations.³⁷ Likewise for any other instance of the T-schema. Consequently, each such instance is approximately true, and therefore reasonable to assert when its vagueness is ignored. (Given the similarity between approximate truth and supertruth, how can (T) be approximately true when it is not supertrue? Answer: according to supervaluationists, the truth predicate expresses supertruth, whereas according to us it does not express approximate truth.)

For us, instances of the T-schema have a status similar to other sentences expressing definitional connections, such as (Bald). The definitional status of (Bald) constrains the simultaneous assignment of legitimate disambiguations to the terms ‘bald’ and ‘has fewer hairs than’; any legitimate disambiguation of the pair of terms must render (Bald) true.³⁸ Likewise, the definitional status of the T-schema constrains the simultaneous assignment of legitimate disambiguations to sentences, the quotation names of those sentences, and the truth predicate. (More on this in section 3.2.) Our theory accords the T-schema its rightful status, as definitional of the sentential truth predicate.

Whether a complex sentence like (T) is approximately true is not a function of the approximate-truth values of its parts. This is the correlate for our theory of the fact that, for supervaluationists, the material biconditional (and other such connectives) are not truth-functional; it results from the fact that approximate-truth value is defined directly for entire sentences, and that there are constraints on simultaneous legitimate disambiguations of definitionally related terms. Nevertheless, for us, the material biconditional (and other such connectives) remain *truth*-functional.

³⁷Although see section 3.2.

³⁸See note 14.

More should be said about the T-schema. We deny that utterances of vague sentences are true, but in ordinary life happily utter such sentences (since, we say, harmless semantic indeterminacy may typically be ignored). This leads us to produce utterances that are licensed by our theory, but which seem incompatible. For instance, in this paper we say:

(1) ‘Snow is white’ is not true

But watch us for a few days or months and you will observe us uttering (2).

(2) Snow is white

We utter (2) because (we claim) it is (typically) acceptable to ignore vagueness, and assert approximately true sentences when we do. Next, you might spot us asserting the following instance of the T-schema.

(3) ‘Snow is white’ is true iff snow is white

We would do so because (3) is approximately true, and it is usually harmless to ignore its vagueness. So our theory licenses our uttering each of (1)-(3). But (1)-(3) are syntactically contradictory. No theory should license uttering all of them, one might say. So, one might conclude, something is wrong with our theory. Since (2) and (3) are clearly acceptable (even true!), (1) must be wrong. But to give up (1) is to give up semantic nihilism.³⁹

The appearance of a problem arises because we ignore different sorts of vagueness at different times and in different contexts. In daily life we, like anyone else, ignore nearly all vagueness. It is in those contexts that we are happy to ignore the vagueness in ‘snow’ and ‘white’, and hence utter (2). Indeed, in those contexts we would also utter (3) and “‘Snow is white’ is true”. Therefore, in those contexts we would *not* utter (1). It is only in a very non-standard context — the context of writing this paper — in which we utter (1). The topic of this paper is vagueness, and we use ‘Snow is white’ as an example of a vague sentence. In this context we are *not* ignoring the vagueness in ‘snow’ and ‘white’. Our intention is to draw attention to this vagueness, to the corresponding failure of ‘Snow is white’ to express a unique proposition, and to its resulting untruth. That is why we utter (1) in the context of writing this paper. But in this context,

³⁹We thank John Hawthorne for pressing this point.

we are no longer willing to utter (2) or (3). So in no one context do we utter each of (1)-(3), nor does our theory say that doing so is acceptable.⁴⁰

3.2 Truth, expression, and legitimate disambiguations

But our utterance of (1) in the context of writing this paper does raise a more serious apparent problem. We say that it is acceptable to utter an approximately true sentence in a context in which its vagueness is ignored. We also accept that when the vagueness of a sentence is being ignored, if it is *not* approximately true then one ought *not* to utter it. But (1), which we sometimes utter, seems approximately *false*. The sentence “‘Snow is white’ is true” involves a sentential-truth predicate, and so is by our lights equivalent to (%):

(%) ‘Snow is white’ expresses a unique proposition, which is true.

The binary predicate ‘expresses’ is vague. Assume that for each of its legitimate disambiguations, *R*, ‘snow is white’ bears *R* to exactly one proposition, which is true. (%) is therefore approximately true, and (1) is therefore approximately false. How then can our utterances of (1) be licensed by our theory?

Relatedly, we *deny* ‘Snow is white’, and “‘Snow is white’ is true’ when in theoretical contexts, despite the fact that each is approximately true. This problem may be solved by pointing out that, as theorists attending to vagueness, we recognize that these sentences are not true (even though they are approximately true). Denial is an appropriate reaction to a sentence regarded as untrue. (See note 16.) But the problem of our *assertion* of (1) remains.

The reason we gave earlier in the paper for uttering (1) was that the facts that determine meaning do not secure a unique meaning for ‘Snow is white’ — ‘Snow is white’ does not have a unique legitimate disambiguation. It is natural to express this recognition by saying:

(#) ‘Snow is white’ does not express a unique proposition.

Since sentential truth requires a unique proposition expressed, (1) then follows. But the problem with this reasoning is that (#) itself is approximately

⁴⁰There are interesting further questions, which we will not take up here, about the connection between our dispositions to utter sentences in different contexts and our *beliefs*. Do our beliefs change as the context changes, or are they somehow differently manifested in different contexts?

false, assuming (as we did in the previous paragraph) that on each legitimate disambiguation of ‘express’, ‘Snow is white’ expresses a unique proposition.

The problem here is clearly generated by the assumption that each legitimate disambiguation of ‘express’ relates ‘Snow is white’ to exactly one proposition — that is, that the legitimate disambiguations of ‘express’ are *one-one* relations. If they were instead *one-many*, relating ‘Snow is white’ to each of its legitimate disambiguations, (#) and (1) would then be approximately true.

But there is pressure towards one-one legitimate disambiguations from platitudes concerning truth and communication. Consider:

(@) If a person sincerely utters a sentence, then that person believes the proposition expressed by that sentence.

(@) is just one instance of our tendency to use definite descriptions of the form “*the* proposition expressed by ‘ ϕ ’”. Assuming that definite descriptions imply uniqueness, (@) is approximately true only if “the proposition expressed by ‘Snow is white’” refers uniquely under all legitimate disambiguations; (@) is therefore approximately true only if the legitimate disambiguations of ‘express’ are one-one. Consider also any instance of the T-schema, for instance (3), which is equivalent on our view to (3’).

(3’) (‘Snow is white’ expresses a unique proposition, which is true) iff snow is white.

If there are legitimate one-many disambiguations of ‘expresses’, any choice of such a disambiguation, combined with any legitimate disambiguation of ‘Snow is white’, will render the left side of (3’) false and the right side true. (3’) would therefore not be approximately true. Apparently, the assumption that typical sentences express unique propositions is embedded deep within ordinary thought about meaning and truth.

In short, we want to utter both (#) and (1), which require the legitimate disambiguations of ‘express’ to be one-many, and also (@) and (3), which require them to be one-one. How can we have it both ways?

When we utter (@) and (3) we are ignoring all vagueness, whereas when we utter (#) and (1) we are attending to the vagueness of ‘Snow is white’. So we can have it both ways if attending to vagueness can affect the nature of the legitimate disambiguations of ‘express’. ‘Express’ has different types of legitimate disambiguations in different contexts: one-one in some, one-many in

others. Its legitimate disambiguations in a context are one-one iff the vagueness of the sentence to which it is applied is ignored in that context. Each way of taking legitimate disambiguations idealizes away the vagueness of ‘express’ along a different dimension. There is no reason to choose one, once and for all.

In ordinary contexts the penumbral connections between ‘express’ and coordinate expressions (such as ‘true’, ‘believes’, ‘says’, etc.) should be preserved. And these connections are based on ignoring vagueness; they are based on an idealized conception on which sentences express unique propositions. It is in those contexts that we utter (2) and (3); the legitimate disambiguations of ‘express’ are then one-one relations. After reflection on vagueness and the metaphysics of meaning, we no longer idealize in this way. Suppose we then attend to the vagueness of ‘Snow is white’. Uttering (4) then seems natural, and does not seem to involve an abuse of the word ‘express’. The legitimate disambiguations have now become one-many, and (4) and (5) are approximately true.

The shiftiness of the legitimate disambiguations of ‘express’ is a genuinely new element of our theory. Our earlier remarks suggested that attending to the vagueness of an expression *E* merely causes one to *refrain* from uttering sentences in which *E* is *used*. For instance, when ignoring the vagueness of ‘Snow is white’ we will readily assert that sentence. After attending to its vagueness, we will no longer assert it; nor will we assert its negation ‘Snow is not white’. But now we see that attending to the vagueness of *E* has an additional effect: it causes us to *assert* the *negations* of certain sentences we formerly asserted in which *E* is *mentioned*. When ignoring the vagueness of ‘Snow is white’ we assert “‘Snow is white’ is true”; after attending to that vagueness we assert “‘Snow is white’ is not true”.⁴¹

Though new, this element flows from our core picture of ignoring and attending to vagueness. Vagueness is usually ignored. Ignoring vagueness

⁴¹The example involving ‘approximately true’ in section 1.4 made clear that we must reserve the right to ignore the vagueness of one occurrence of a word type while attending to the vagueness of another. That is not to say that just any pattern of ignoring and attending to the vagueness of occurrences is natural. It would be odd to attend to the vagueness of the left-hand side of ‘Snow is white iff snow is white’ while ignoring the vagueness of the very same sentence on the right-hand side. It would likewise be odd to utter (3) or (3’) while attending to the vagueness of the quoted occurrences of ‘Snow is white’ on their left-hand sides (thus inducing one-many legitimate disambiguations of ‘expresses’) but ignoring the vagueness of the used occurrences on their right-hand sides. Given this pattern of ignoring, the sentences would be approximately false! (Thanks to Timothy Williamson for discussion.)

is speaking as if vagueness does not exist. Without vagueness, expression would be one-one, and utterances of sentences like (Ⓜ) and (3) would be appropriate. It is only in unusual circumstances, when we are challenged by borderline cases or when we theorize about vagueness, that we attend to vagueness. But when we attend to vagueness, we recognize the multiplicity of legitimate disambiguations, and better acknowledge the true nature of the metaphysics of meaning. Once we do so, it is natural to think of a sentence as expressing each of its legitimate disambiguations. ‘Expresses’, after all, is the word for the relation between sentences and their meanings. So when writing this paper, and thus attending to the vagueness of ‘Snow is white’, ‘expresses’ has one-many legitimate disambiguations, and we can say things like (Ⓜ) and (1).

And even when writing this paper we must ignore most vagueness — otherwise the paper would have been far shorter! Without ignoring *some* vagueness one cannot say much about anything, not even about vagueness.

University of Rochester
Rutgers University

References

- Devitt, Michael (1984). *Realism and Truth*. Oxford: Basil Blackwell.
- Dummett, Michael (1975). “Wang’s Paradox.” *Synthese* 30: 301–324.
- (1981). *The Interpretation of Frege’s Philosophy*. Cambridge, Mass.: Harvard University Press.
- Fara, Delia Graff (2000). “Shifting Sands: An Interest-Relative Theory of Vagueness.” *Philosophical Topics* 28: 45–81.
- Fine, Kit (1975). “Vagueness, Truth and Logic.” *Synthese* 30: 265–300. Reprinted in Keefe and Smith 1996: 119–150.
- Fodor, Jerry A. and Ernest Lepore (1996). “What Cannot Be Evaluated Cannot Be Evaluated, and It Cannot Be Supervalued Either.” *Journal of Philosophy* 93: 516–535.

- Frege, Gottlob (1903). *Grundgesetze der Arithmetik*, vol. 2. Jena: H. Pohle. Selections appear in Michael Beany, ed., *The Frege Reader* (Oxford: Blackwell, 1997).
- Grice, H.P. (1989). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Heller, Mark (1990). *The Ontology of Physical Objects: Four-dimensional Hunks of Matter*. Cambridge: Cambridge University Press.
- Horn, Laurence R. (1989). *A Natural History of Negation*. Chicago: University of Chicago Press.
- Keefe, Rosanna and Peter Smith (eds.) (1996). *Vagueness: A Reader*. Cambridge, Mass.: MIT Press.
- Lewis, David (1983). "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343–77. Reprinted in Lewis 1999: 8–55.
- (1984). "Putnam's Paradox." *Australasian Journal of Philosophy* 62: 221–36. Reprinted in Lewis 1999: 56–77.
- (1993). "Many, But Almost One." In John Bacon, Keith Campbell and Lloyd Reinhardt (eds.), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong*, 23–42. Cambridge: Cambridge University Press. Reprinted in Lewis 1999: 164–182.
- (1996). "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549–567. Reprinted in Lewis 1999: 418–445.
- (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Ludwig, Kirk and Greg Ray (2002). "Vagueness and the Sorites Paradox." *Philosophical Perspectives* 16: 419–461.
- McGee, Vann and Brian McLaughlin (1995). "Distinctions Without A Difference." *Southern Journal of Philosophy* 33 (Supp.): 203–251.
- Rea, Michael (ed.) (1997). *Material Constitution*. Lanham, Maryland: Rowman & Littlefield.

- Sanford, David H. (1976). "Competing Semantics of Vagueness: Many Values Versus Super-truth." *Synthese* 33: 195–210.
- Soames, Scott (1999). *Understanding Truth*. Oxford: Oxford University Press.
- Sorensen, Roy (1988). *Blindspots*. Oxford: Oxford University Press.
- (2001). *Vagueness and Contradiction*. Oxford: Clarendon.
- (2002). "Vagueness." *Stanford Encyclopedia of Philosophy*. Available online at <http://plato.stanford.edu/entries/vagueness/>.
- Sperber, Dan and Deirdre Wilson (1986). "Loose Talk." *Proceedings of the Aristotelian Society* 86: 153–171.
- Tye, Michael (1989). "Supervaluationism and the Law of the Excluded Middle." *Analysis* 49: 141–143.
- Unger, Peter (1979a). "I Do Not Exist." In G. F. Macdonald (ed.), *Perception and Identity: Essays Presented to A. J. Ayer with His Replies to Them*, 235–251. New York: Macmillan. Reprinted in Rea 1997: 175–90.
- (1979b). "There Are No Ordinary Things." *Synthese* 41: 117–154.
- (1979c). "Why There Are No People." In Peter French, Theodore E. Uehling, Jr. and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy IV*. Minneapolis: University of Minnesota Press.
- Wheeler, Samuel C. III (1975). "Reference and Vagueness." *Synthese* 30: 367–379.
- (1979). "On That Which Is Not." *Synthese* 41: 155–173.
- Williamson, Timothy (1994). *Vagueness*. London: Routledge.
- Wright, Crispin (1975). "On the Coherence of Vague Predicates." *Synthese* 30: 325–365.