Draft of work-in-progress for 4/24/18 Mind & Language session [revision of 4/20/18]. Comments invited.

When Meaning Meets Vagueness

(Accommodating Vagueness in Semantics and Metasemantics)

Stephen Schiffer
New York University

A noteworthy feature of presentations of foundational theories of meaning is that they nearly always ignore vagueness, even though virtually every utterance is vague. This needn't be a fault; perhaps ignoring vagueness can be justified as a useful idealization akin to Galileo's ignoring friction in his idealized model of bodies in motion. Perhaps not. Either way there remains the question of what shape a theory of meaning must take if it's to accommodate the vagueness of vague expressions and vague speech acts. That is the overarching question this paper addresses in three parts: I. Vague Speaker-Meaning; II. Vague Expression-Meaning: What It Isn't; and III. Vague Expression-Meaning: What Might It Be?

I. Vague Speaker-Meaning

Assertoric speaker-meaning—henceforth, for present purposes, simply speaker-meaning—is the notion of a speaker's meaning that such-and-such, as when, for example, in uttering 'He's ready' Jill meant that Jack was ready to go to dinner. It's the most general kind of assertoric illocutionary act, the genus of which all other kinds of assertoric illocutionary acts—saying that such-and-such, asserting that such-and-such, denying that such-and-such, objecting that such-and-such, telling so-and-so that such-and-such, etc.—are species. The dominant conception of speaker-meaning in both philosophy of language and linguistic semantics is that of a relation, *S meant p*, between a person *S* and a proposition *p* that she meant, where a proposition is an abstract entity that has truth conditions, has those truth conditions necessarily, and is mind- and language-independent in that it belongs to no language and wasn't created by what anyone said or thought.[1] Let's call this

[1] For present purposes it doesn't matter to which kind of proposition—Fregean, Russellian, functions from possible worlds into truth-values, whatever—the propositions we mean are taken to belong, provided (for a reason that emerges in Part II) they can be assigned in a compositional truth-theoretic semantics, but for simplicity of exposition I will sometimes write as though they are taken to be Russellian propositions, i.e. structured entities whose basic components are the objects and properties our speech acts are about.

dominant conception the *propositionalist account of speaker-meaning* and begin by testing the account's ability to say what a speaker means in an act of vague speaker-meaning.

Here is an unexceptional example of vague speaker-meaning. Tom is reading in the park when a woman appears, calling 'Billy, where are you? We have to leave now'. Intending to tell her something she might find helpful, Tom says to the woman 'A boy was here a little while ago'. We would certainly regard Tom as having told the woman something, and therefore as having meant and said something in producing his utterance. If the woman to whom he spoke didn't catch his words and asked him what he said, Tom wouldn't hesitate to say 'I said that a boy was here a little while ago', and we, knowing what we do, would accept Tom's report of what he said as true. We would unhesitatingly take Tom's utterance to be an act of speaker-meaning. Tom's utterance was also vague. His utterance was vague because even if it was definitely true or definitely false, it might have been neither; it might have been borderline true/false. It's three-ways overdetermined that Tom's utterance was vague, for its contained utterances of 'boy', 'here', and 'a little while ago' were vague, and the vagueness of any one of those utterances sufficed to make Tom's utterance of 'A boy was here a little while ago' vague. And each of those utterances was vague because, even if the application to which Tom put it was definitely correct or definitely incorrect, it might have been neither. I will use Tom's utterance as an exemplar of vague speaker-meaning.

The first question to be addressed is: On the assumption that Tom meant something in uttering 'A boy was here a little while ago' and that the propositionalist account of speaker-meaning is correct, *what did Tom mean in uttering 'A boy was here a little while ago'*?[2] I will

[2] My use of 'vague', 'definite', 'indefinite' and 'borderline' throughout this essay is pretheoretic, which isn't to say it's not discussable (for pretty much any assumption one makes about a philosophically important notion, it's a safe bet that some philosophers will find it debatable). So, for example, as far as this paper is concerned, it's not assumed that 'It's indefinite whether Harold is bald' entails 'It's neither true nor false that Harold is bald'. I will, however, assume that: (i) '*x* is borderline *F*' entails 'It's indefinite whether *x* is *F*'. (ii) 'It's indefinite whether *x* is *F*' doesn't entail '*x* is borderline *F*'. For example, if, as many suppose, it's indefinite whether the continuum hypothesis is true, that wouldn't entail that it was borderline true or even vague. In this essay, however, the only indefiniteness that will be at issue is vagueness-induced indefiniteness, so for all that presently matters 'It's indefinite whether *x* is *F*' is equivalent to '*x* is borderline *F*'. (iii) If in a person *a*'s idiolect it's indefinite whether *x* is *F*, then it's impossible for *a* to know whether *x* is *F*. If, for example, 'bald' in my idiolect is such that Harold is borderline bald, then nothing can count as my discovering that Harold is in fact bald or that he is in fact not bald. If I did come to know that Harold was bald, then I would thereby come to know that he's not borderlines bald. [Some philosophers would evidently disagree and would say that if, as 'bald' is used in my idiolect, Harold is borderline bald then, notwithstanding that, it's possible for it to be indefinite whether I know that Harold is bald (see e.g. Dorr 2003, Wright 2003, and Barnett 2011). I disagree (see e.g.

consider four answers, but before stating them I should explain a distinction to which those answers will advert. This is a distinction between what I will call *metaphysically-precise* propositions, properties and things, on the one hand, and, on the other hand, *metaphysically-vague* propositions, properties and things. I'll say that a proposition is metaphysically precise provided that, necessarily, there is a fact of the matter as to what truth-value, if any, it has. If bivalence holds for propositions—i.e. if it's necessarily the case that every proposition is such that it's either a fact that it's true or else a fact that it's false—then every proposition is metaphysically precise, but its being metaphysically precise is compatible with it's being vague on an epistemicist reading of 'vague'. If, as Frege held, it's necessarily the case that every proposition is such that it's a fact that it's true, a fact that it's false, or else a fact that it's neither true nor false, then that, too, would make every proposition metaphysically precise. And if there are $n \geq 3$ truth-values such that it's necessarily the case that there is a fact of the matter as to which of the n truth-values a proposition has, then again every proposition is metaphysically precise. For simplicity of exposition I will for present purposes assume that a proposition is metaphysically precise just in case it's necessarily bivalent, either true or else false. The assumption of bivalence for metaphysically-precise propositions induces an assumption of bivalence for metaphysically-precise properties: I shall say that a property (or n-ary relation) is metaphysically precise provided that, necessarily, every thing (or n-tuple) either instantiates the property (or n-ary relation) or else fails to instantiate it. A thing that is neither a property nor a proposition—e.g. an apple, a dog, a geographical area, or a period of time—is metaphysically precise provided that it has precise conditions of individuation. A proposition, property, or particular is metaphysically vague just in case it's not metaphysically

---

Schiffer forthcoming) and will for the purposes of this paper assume that I'm right; but I suspect that nothing essential to the argument of this paper would be undermined if my working assumption were instead that, if, as I use 'bald', Harold is borderline bald, then it's impossible for me to definitely know that he's bald.] ((iv) One might wonder what the difference is supposed to be between *x*'s being *F* and x's being definitely *F*. What, for example, is the difference between its being *true* that it's raining and its being *definitely true* that it's raining? If they are two distinct facts, in what can the difference between them consist? Are we to imagine it's raining harder if we suppose that the proposition *that it's raining* is definitely true, as opposed to being merely true? Of course not. It's not that one might know that Sadie is lethargic and then wonder whether she's also definitely lethargic. What is really going on is that we want our pretheoretic use of 'indefinite' and 'definite' to capture facts that a theory of vagueness needs to explain but to *leave open* the issues that a theory of vagueness should be expected to resolve. The most important feature of my pretheoretic use of 'indefinite' is that if it's indefinite whether …, then it's impossible to know whether …, but we want to leave open the consequences of this for such things as bivalence and excluded middle. The primary role of 'definite' is that if it's definite that …, then it's not indefinite whether …., which is why '*x* is definitely *F*' entails '*x* is *F*'. But the reason '*x* is *F*' doesn't entail '*x* is definitely *F*' isn't that we think '*x* is *F* and it's indefinite whether *x* is *F*' is possibly true, but rather that we want to leave that question open for a theory of vagueness to resolve.

precise. Henceforth, unless unclarity threatens, 'metaphysically' will occur invisibly and inaudibly in the phrase 'metaphysically precise/vague'.

The four answers to be considered to the question before us—namely, "What did Tom mean in uttering 'A boy was here a little while ago', given that he meant something and that the propositionalist account of speaker-meaning is correct?"—are as follows:

(A) For some precise proposition *p*, Tom meant *p* in uttering 'A boy was here a little while ago', but it's indefinite whether he meant *p* in uttering that sentence.

(B) There is no precise proposition that Tom definitely meant in uttering 'A boy was here a little while ago', but there are two or more precise propositions each such that it's indefinite whether he meant it in uttering that sentence.

(C) For some vague proposition *p*, Tom meant *p* in uttering 'A boy was here a little while ago'.

(D) There is no vague proposition that Tom definitely meant in uttering 'A boy was here a little while ago', but there are two or more vague propositions each such that it's indefinite whether he meant it in uttering that sentence.

Although (A)-(D) are the only answers I will consider, logical space contains infinitely more possible answers. In fact, for every answer *X* contained in logical space, logical space also contains the answer that it's indefinite whether *X*, so that, for example, in addition to (A), logical space also contains the answers

(A_2) It's indefinite whether (A).

(A_3) It's indefinite whether (A_2).

(A_4) It's indefinite whether (A_3)

⋮

There is an important question as to how we are to understand so-called higher-order vagueness,[3] but it's not one that present purposes require me to try to answer; for once we finish our critical examination of (A)-(D) it will be clear that we have no need to consider the permutations of them induced by higher-order vagueness and that, consequently, if the propositionalist has a correct answer to the question before us, it will be one of (A)-(D). So the big question now is whether any of (A)-(D) can survive scrutiny.

[3] See Wright (2010).

Re (A) [For some precise proposition p, Tom meant p in uttering 'A boy was here a little while ago', but it's indefinite whether he meant p in uttering that sentence]. If it's indefinite whether such-and-such is the case, then it's impossible for anyone to know whether such-and-such is the case. How can there be any proposition that Tom meant in uttering 'A boy was here a little while ago' if it's impossible for him or anyone else to know that he meant it? Yet (A) would be true if the epistemic theory of vagueness were true, for that theory expounds the thesis that "the proposition a vague sentence expresses in a borderline case is true or false, and we cannot know which."[4] At the same time, we can know that (A) is false if we can know that there couldn't have been a precise proposition that Tom meant in uttering 'A boy was here a little while ago', and I believe we can know that in the following way.

The vagueness of Tom's utterance, we have noticed, was three-ways overdetermined: by the vagueness of his utterance of 'boy', the vagueness of his utterance of 'here', and the vagueness of his utterance of 'a little while ago'. Consequently, if Tom meant a precise proposition of any kind (Russellian, Fregean, etc.) in uttering 'A boy was here a little while ago', then it must also have been the case that:

(1) for some precise property φ, Tom referred to φ with the token of 'boy' he uttered (if we pretend that 'human' and 'male' express precise properties and that there is a precise moment at which a person comes into existence, then φ might be the property of being a human male whose age in milliseconds $\leq$ 531,066,240,000);[5]

(2) for some precise area α—i.e. area that has precise boundaries, and thus comprises a precise number of square millimeters, and a precise location relative to Tom's location—Tom referred to α with 'here';

(3) for some precise period of time π, Tom referred to π with 'a little while ago', where in order for that to have been the case there must have been instants of time of 0 duration t, t', t'', and real numbers n, n', n'', such that (i) t was the instant π began, (ii) t' was the instant π ended, (iii) t'' was the instant of time a "little while" before which Tom was saying the boy was in α, the precise area to which Tom referred with 'here'; (iv) n

[4] Williamson (1997: 921).

[5] 'Expressed φ' would conform better with philosophical usage than does 'refers to φ', but I use the latter in order to unify my discussion of what (A) requires of 'boy' with my discussion of what it requires of 'here' and 'a little while ago'. Also, as an expository convenience, instead of saying e.g. *Tom referred to φ with the token of 'boy' he uttered*, I'll say *Tom referred to φ with 'boy'*, where that will be shorthand for the longer way of speaking. Likewise, for *Tom referred to area α with 'here'* and *Tom referred to period of time π with 'a little while ago'*.

> is the precise number of milliseconds between t and t', (v) n' is the precise number of milliseconds between t'', the instant of time from which all measurements of time relevant to the reference of 'a little while ago' emanate, and t, the instant of time such that if the boy's appearance in α was so much as one yoctosecond before t,[6] then it was too long before t'' to count as "a *little* while ago," and (vi) n'' is the precise number of milliseconds between t'' and t', the end of π and thus the instant of time such that if the boy's appearance in α was so much as one yoctosecond after t', then it was too soon before t'' to count as a "little while *ago*."

We can show that Tom didn't mean any precise proposition in uttering 'A boy was here a little while ago' if we can show that any one of (1)-(3) is false. I think we can show that each of (1)-(3) is false. I'll begin with (2) since it seems to be the simplest of the three requirements. There are at least the following reasons why Tom couldn't have referred to a precise area with 'here'.

(i) Acts of speaker-reference are *intentional* acts in that, if for some *S*, *o* and *x*, *S* referred to *o* with *x*, then *S* uttered *x* with those *o*-directed intentions that are constitutive, or at least partly constitutive, of her referring to *o* with *x*, and if *S* has the concept of speaker-reference, then, in the normal case, she intended to refer to *o* with *x*. Tom's utterance was a normal case, but it ought to be obvious that there was no precise area to which he intended to refer when he produced it, nor would he have thought there was any need to refer to such an area. We should expect a fuller description of the imagined scenario to include the fact that when he uttered 'A boy was here a little while ago' Tom was confident that he was saying something true because he was confident that a human male child no older than six was within four meters of him no more than five minutes before he spoke. Tom gave no thought to which of the indefinitely many precise areas containing the boy was the one he wanted to make a statement about, because there was no precise area about which he wanted to make a statement. Careful and considerate speakers try not to use a vague term unless they are confident that their use of it would be recognized as correct, but they have no reason to consider what would have to have been the case for their use of the term to have been borderline correct, or just barely correct or incorrect. Tom, if he understands it and has his wits about him, ought to be rendered speechless by the question "Exactly which area did you intend to refer when you uttered 'here'?" Nor did Tom think there was any need to refer to any particular area that contained the boy, for he would think he succeeded in giving the woman the information he intended to give her if, as a result of his utterance, she believed that a boy was at a location within the vicinity of Tom which made his utterance true, and that didn't require her to think that

[6] One yoctosecond = one trillionth of a trillionth of a second.

any particular area in which the boy was contained was the area to which Tom referred in producing his utterance.

(ii) We refer to things in order to make known to our hearers what we are talking about. In a normal case, such as Tom's, a speaker can't refer to a thing if she knows that her hearer won't be able to know to what she was referring. Tom would know that even if there were a precise location to which he wanted to refer, his hearer would have no way of knowing which of the indefinitely many eligible precise areas was the one to which he was referring. Given that, he couldn't have intended to refer to any precise area.

(iii) There is a deeper explanation of why Tom couldn't have referred to any precise area. The statement

There is a precise area α such that Tom intended to refer to α with 'here'

ascribes to Tom an intention that is *de re* with respect to an unspecified precise area α, and just as one can have a belief that is *de re* with respect to a thing under one way of thinking of it but not under another, so one can have an intention that is *de re* with respect to a thing under one way of thinking of it but not under another. Let α be any precise area in the vicinity of Tom. What way of thinking of α might Tom have under which it would be possible for him to intend to refer to α with 'here'? He will have no perceptual way of thinking of α that would do the job, and I am unable to conceive of any knowledge by description of α he might have that would enable him to intend to refer to α under it. *If* a precise area α could be secured as the referent of the token of 'here' Tom uttered without his having any intention that was *de re* with respect to α, then perhaps Tom could have intended to refer to α under the description *the area to which the token of 'here' I uttered refers*. Yes, but it's every bit as difficult to see what feature one of indefinitely many indistinguishable precise areas could make it alone the referent of the token of 'here' as it is to see what feature could make it alone the one to which Tom referred with that token of 'here'. True, there are theorists who would claim both that the speaker's referential intentions are never a factor in determining the referents of indexicals or demonstratives and that there is a precise area to which the token of 'here' Tom uttered uniquely refers, but the things these theorists say are the things that determine the reference of a pronoun or demonstrative relative to "contexts of utterance" are such things as salience, common ground, conversational relevance, and discourse-coherence relations, and it ought to be obvious that no such features are available to secure a *precise area* as the referent of 'here' relative to Tom's context of utterance. Consider salience, for example: if it's *given* that there is a precise area to which 'here' in the Tom example uniquely refers, then none of the *indefinitely many* precise areas that would be eligible to be that referent would be any more salient than any of the others, and it's equally obvious that the same goes for common ground, conversational relevance, and whatever other contextual features might be thought to be reference-

determining. It seems impossible to conjure up any way of thinking of a precise area that would yield a way of thinking of it under which Tom might have *any* intention (or belief) that is *de re* with respect to it. It seems, in short, that there couldn't have been anything about any particular precise area that would explain what made it, rather than any of the indefinitely many other precise areas that are indistinguishable from it, the area to which Tom referred with 'here'.

I conclude that there is no precise area to which Tom referred with 'here', and therefore no precise proposition that he meant in uttering 'A boy was here a little while ago'. So (A) is false. The same sort of considerations used to show that Tom couldn't have referred to a precise area with 'here' can also be used to show that he couldn't have referred to a precise period of time with 'a little while ago'. In fact, given the complexity of what would have to be the case in order for Tom to have referred to a precise period of time, it should be more intuitively obvious that he couldn't have referred to a precise period of time than it is that he couldn't have referred to a precise area. So that's another way to show that there couldn't have been a precise proposition that Tom meant in uttering 'A boy was here a little while ago', and therefore another way to show that (A) is false.

Can the same sort of considerations used to show that Tom couldn't have referred to a precise area with 'here', or to a precise period of time with 'a little while ago', also be used to show that Tom couldn't have referred to a precise property with 'boy'? I believe so, but it first has to be ruled out (i) that there is such a thing as *the* property of being a boy and that it's the meaning of 'boy', (ii) that Tom knows that 'boy' means the property of being a boy, and may therefore be said to refer to that property with the token of 'boy' he uttered, (iii) that the property of being a boy = (say) the property of being a human male whose age in milliseconds seconds ≤ 531,066,240,000.07229; and (iv) that the vagueness of 'boy' is explained in part by its being impossible for anyone to know that the property of being a boy = the property of being a human male whose age in milliseconds seconds ≤ 531,066,240,000.07229, and therefore impossible to know that 'boy' means the property of being a human male whose age in milliseconds seconds ≤ 531,066,240,000.07229, in which case 'Tom intended to refer to the property of being a boy' would be true while 'Tom referred to the property of being a human male whose age in milliseconds seconds ≤ 531,066,240,000.07229' was false. This line, however, can't work: (i)-(iv) are ruled out by a well-known feature of every vague expression that I call *penumbral shift*.[7]

Penumbral shift is a feature of every vague expression, but I will explain it only with respect to vague predicates like 'boy', 'violinist', 'mathematician', 'communist', and 'triangular' which, unlike such vague predicates as 'local', 'ready', 'tall' or 'intelligent', are apt to be regarded as having constant characters if vagueness is being ignored. I'll refer to such vague predicates as

[7] See e.g. Schiffer (2010) and (2016).

predicates*. Now, every token of a predicate* has a *penumbral profile*, and two tokens of a predicate* have the same penumbral profile just in case if either token is true/false of a thing, then likewise for the other; if either token is such that it's indefinite whether it's true/false of a thing, then likewise for the other; if either token is such that it's indefinite whether it's indefinite whether it's true/false of a thing, then likewise for the other; and so on. Then we may say that int*ra*personal penumbral shift (when restricted to predicates*) is the fact that the penumbral profiles of tokens of a predicate* in a person's idiolect may shift somewhat from one token of the predicate* to another; that is to say, two tokens of a predicate* in the person's idiolect may have somewhat different penumbral profiles. The "somewhat" qualification is important. For example, if Clyde is a man whose scalp is as hairless as a billiard ball and on whose scalp no hair can grow, then every token of 'bald man' must be true of Clyde, and if Clyde is blessed with a head of hair like the one Tom Cruise appears to have, then every token of 'bald man' must be false of him. At the same time, penumbral shift makes it possible for a person to produce three utterances in three different contexts in which the same predicate* is applied to the same thing such that one utterance is true, a second false, and a third neither definitely true nor definitely false, even though the thing underwent no relevant change between the time of the first utterance and the time of third. Here are three examples:

- At a party George is asked whether Henrietta came to the party with anyone, and he replies, gesturing at a man with no apparent hair on his scalp, 'She came with that bald guy'. That utterance would most likely be accepted in the context in which it occurred as true, even if it transpired that the man in question shaves his scalp but wouldn't be said to be bald if he let his hair grow out. In another conversation, however, in which the discussion is about hereditary baldness, George might correct a remark about the same man by saying, 'No; he's not bald; he just shaves his scalp', and that utterance, in that context, would very likely count as true. In still another context the question is raised whether a man who shaved his scalp would be bald if no one would take him to be bald if he stopped shaving his scalp and let his hair grow out, and in that context it might be true for George to say, 'That's undetermined by the use of "bald" in everyday speech; such a man would be neither definitely bald nor definitely not bald'.
- In a community in which people typically marry before the age of twenty, an utterance of 'He's a bachelor' may count as true when said of an unmarried eighteen-year-old male, whereas in a conversation about New Yorkers, where for both men and women the average age for

a first marriage is between thirty and thirty-five, an utterance of 'He's a bachelor' would most likely not count as definitely true when said of an unmarried eighteen-year-old male, and may even count as false.

- Al is boiling water to make pasta, which he will serve with store-bought marinara sauce. Asked what he's doing, Al replies 'I'm cooking dinner'. On another occasion he is doing the same thing, but when asked whether he's cooking dinner, he replies 'No, I'm too tired to cook; I'm just putting together some pasta that I'll serve with some store-bought marinara sauce'. Both utterances are accepted as true.

The foregoing has been about intrapersonal penumbral shift. There is also int*er*personal penumbral shift, which makes it possible for there to be three simultaneous utterances by three different people in which the same predicate* is applied to the same thing in all three utterances, but one utterance is true, another is false, and still another is neither definitely true nor definitely false. Interpersonal penumbral shift is often due to a vague expression's having somewhat different uses in somewhat different idiolects, but it may also occur when the same language is the idiolect of different people.

We can explain in the following way how penumbral shift precludes any of (i)-(iv) from being true. Penumbral shift makes it possible for there to a true utterance of, say, 'Sasha is a boy' in one context and a false utterance of that sentence by the same person in a different context, even though there was no relevant change in Sasha's age between the two utterances. That would be impossible if some property φ was the meaning of 'boy', for if that were so, φ would be the content of every token of 'boy', in which case the two imagined utterances of 'Sasha is a boy' would have the same truth-value.[8]

I conclude that we know that there was no precise proposition that Tom meant in uttering 'A boy was here a little while ago', and therefore know that (A) is false.

Re (B) [There is no precise proposition that Tom definitely meant in uttering 'A boy was here a little while ago', but there are two or more precise propositions each such that it's indefinite whether he meant it in uttering that sentence]. We know from the discussion of (A) that there was no precise proposition that Tom meant in uttering 'A boy was here a little while ago', and if we know that, then it's definitely the case. Therefore, it's not the case that there is any proposition such that it's indefinite whether Tom meant it: the considerations adduced to show that (A) is false also show that every precise proposition is such that Tom definitely did not mean it. Therefore, (B), as well as (A), is false.

[8] See Part III for an elaboration of this point.

Re (C) [For some vague—i.e. metaphysically-vague—proposition p, Tom meant p in uttering 'A boy was here a little while ago']. This answer presupposes that things other than words and concepts are vague—more specifically, that there are propositions for which there is no fact of the matter as to what truth-value, if any, they have; properties for which there is no fact of the matter as to which things, if any, they are true of; and things such that there is no fact of the matter as to what their conditions of individuation are. Some philosophers think it's incoherent to suppose that there might be such things. If I could assume that there were no vague propositions, properties or things, it would follow from my responses to (A) and (B) that in uttering 'A boy was here a little while ago' Tom definitely didn't mean any proposition, and that therefore the propositional account of speaker-meaning is false. But in order to give that account its best run for the money, I will for present purposes assume that there are vague propositions, properties and things. What we will see is that, even given that assumption, (C) fails for the same sort of reason that (A) fails.

The hypothesis that Tom meant a *precise* proposition in uttering 'A boy was here a little while ago' entails that he referred to a precise property with 'boy', referred to a precise area with 'here', *and* referred to a precise period of time with 'a little while ago'. The hypothesis that Tom meant a *vague* proposition in producing his utterance entails only that he referred to a vague property with 'boy', referred to a vague area with 'here', *or* referred to a vague period of time with 'a little while ago'. For example, it's compatible with Tom's having meant a vague proposition that he referred to a precise property with 'boy' and referred to a precise period of time with 'a little while ago', but referred to a vague area with 'here'. But since, as we've seen, Tom can't have referred to a precise property with 'boy', or to a precise area with 'here', or to a precise period of time with 'a little while ago', he would have meant a vague proposition only if he referred to a vague property with 'boy', referred to a vague area with 'here', *and* referred to a vague period of time with 'a little while ago'. Consequently, we can see that (C) is false by seeing that any one of those three things wasn't possible. In fact, we can see that none of them was possible.

Here is why no vague area could have been the area to which Tom referred with 'here'. The fundamental reason Tom couldn't have referred to any *precise* area with 'here' was that he had no way of thinking about any precise area under which he might have intended to refer to it, and this because each such area was for him indistinguishable from the indefinitely many precise areas that differed from it only in some imperceptible way. In other words (more or less), there is nothing to explain how just one of indefinitely many precise areas could have had a feature not possessed by any of the other areas that would have made it alone the area to which Tom referred. The same is true of the indefinitely many vague areas (assuming there are such things) each of which was as likely as any of the others to have been the vague area to which Tom uniquely referred with 'here'. For example, there would have been among the eligible vague areas two areas, α_1 and α_2, such that if any location is definitely in α_1, then it's also definitely in α_2, and vice

versa, the only difference between the two areas being that there are locations such that, while it's indefinite whether they are in α_1, it's merely indefinite whether it's indefinite whether they are in α_2, and that's not a difference that would enable Tom to intend to refer to either area. And, of course, what goes for vague areas also goes, *mutatis mutandis*, for vague periods of time and vague properties.

Re (D) [There is no vague proposition that Tom definitely meant in uttering 'A boy was here a little while ago', but there are two or more vague propositions each such that it's indefinite whether he meant it in uttering that sentence]. (D) stands to (C) as (B) stands to (A). We know from the discussion of (C) that there was no vague proposition that Tom meant in uttering 'A boy was here a little while ago', and if we know that, then it's definitely the case. Therefore, it's not the case that there are two or more vague propositions each such that it's indefinite whether Tom meant it in uttering 'A boy was here a little while ago', for the considerations adduced to show that (C) is false also show that, if there are vague propositions, then every vague proposition is such that Tom definitely did not mean it. Therefore, (D), as well as (C), is false.

Some Upshots

So, I submit, none of (A)-(D) survives scrutiny. None survives scrutiny because in each case we see that, for every proposition *p*, Tom definitely didn't mean *p*. That result also shows that any permutation of any one of (A)-(D) induced by higher-order vagueness must also be false. Tom's utterance is an arbitrary example of vague speaker-meaning, and since virtually every utterance is vague, virtually every utterance is an instance of vague speaker-meaning. If there was no proposition that Tom meant in producing his utterance, then speakers virtually never mean propositions when they speak. The propositionalist conception of speaker-meaning collapses in the face of vague speaker-meaning, if what I've argued in the preceding section is correct.

Might some non-propositionalist conception of speaker-meaning escape the objections leveled against the propositionalist conception? Such an account might take speaker-meaning to be a relation to contentful things other than propositions—say, to formulae in a language of thought or, as on Davidson's "paratactic" theory of saying-that, to utterances of public-language sentences.[9] In fact, it's not easy to see how *any* account of speaker-meaning can be correct. For no matter what account a theorist accepts, won't she want to say, for example, that what Tom meant in uttering 'A boy was here a little while ago' was determined in part by what area he was referring to with 'here', so that if the theorist accepts that Tom definitely wasn't referring to any area, then she will have no way to assign even indefinite truth conditions to what Tom meant? At

[9] Davidson (1968).

the same time, we surely don't want to say that speakers never mean anything when they speak, and we surely do want to say that an utterance of a sentence such as

> In uttering 'Mon professeur de philosophie est léthargique' Pierre said, and thus meant, that his professor of philosophy was lethargic

might well be true. What we have at this point is a *puzzle*—in fact, a couple of related puzzles: one about what speaker-meaning is and another about the semantic properties of speech-act reports. I'll touch on this again in Part III.

The propositionalist conception of speaker-meaning is the dominant conception of it largely, if not entirely, because the dominant conception of propositional attitudes is that they are … well, *propositional* attitudes. The dominant view of believing, for example, is that it's a relation between a believer and a proposition she believes. But if the considerations adduced to show that there was no proposition that Tom meant in producing his utterance really do show that, then they also show that vague propositional attitudes aren't relations to propositions. This is an important point. Other philosophers have made their own trouble for the view that communication involves a speaker's uttering words that encode the proposition she wants to communicate, and that her attempted communication is successful just in case her hearer successfully decodes the encoded proposition. But the philosophers who have argued against this view of communication have held that, while successful communication doesn't consist in a hearer's entertaining or believing the very same proposition that is the content of the belief the speaker expressed in producing her utterance, it does consist in a certain similarity relation's obtaining between the proposition the speaker believed and the one the hearer entertained or believed as a result of the speaker's utterance. That was the view to which Frege was giving voice when he wrote:

> In the case of an actual proper name such as 'Aristotle' opinions as to the sense may differ. It might, for instance, be taken to be the following: the pupil of Plato and teacher of Alexander the Great. Anybody who does this will attach another sense to the sentence 'Aristotle was born in Stagira' than will a man who takes as the sense of the name: the teacher of Alexander the Great who was born in Stagira. So long as the reference remains the same, such variations of sense may be tolerated, although they are to be avoided in the theoretical structure of a demonstrative science and ought not to occur in a perfect language.[10]

[10] Frege (1892). For contemporary expressions of the view see e.g. McDowell (1984a), Heck (2002), Buchanan and Ostertag (2005), and Buchanan (2010). NYU Ph.D. student Martin Abreu defends a novel version of this line in his nearly completed doctoral dissertation.

But if what I have been arguing is on the right track, not only was there no proposition that Tom meant in uttering 'A boy was here a little while ago', there were also no propositions that were the contents of the beliefs and intentions that lead him to utter that sentence. This should be clear, for if what I said about Tom's utterance of e.g. 'here' was correct, then the reason he couldn't refer either to a precise or to a vague area with 'here' is that there was nothing about any area of either kind that could explain how Tom could have an intention or belief that was *de re* with respect it. At the same time, to redirect to propositional-attitude reports a point already made about speech-act reports, it's true that what led to his utterance was his knowing that a boy had been in his vicinity a little while before he spoke, and that he said what he did to the woman to whom he spoke because he wanted to impart that knowledge to her. It's just that, although the propositional-attitude report I just made in writing the preceding sentence is true, neither its 'that'-clause nor the occurrence of 'that knowledge' in it refers to a proposition.

Anyway, these are issues to which we'll presently return. Let's turn now to how vagueness affects expression-meaning.

II. Vague Expression-Meaning: What It Isn't

Virtually every sentence we encounter—e.g. this one—is vague. At the same time, familiar foundational models of expression-meaning are developed and presented as idealizations that prescind from the vagueness of vague expressions, theorizing about the semantics of those expressions as though they weren't vague. Later in this part (i.e. Part II) I will claim that there is an important commonality among these models: they all require our language to have what I will call a *compositional truth-theoretic semantics*, and at that point the crucial question will be whether vague expressions can be accommodated within such a semantics in a way that doesn't ignore their vagueness.

I believe we can get some leverage on answering this question just by focusing on a certain toy model, which I'll call *Model T* (*T*, for short). *T* is an idealized model that, like other idealized semantic models, prescinds from the vagueness of vague expressions. To help keep inessential complexities in check I will suppose that *T* takes every sentence to be truth evaluable, the kind of sentence that can be used to perform an assertoric act of speaker-meaning. The brochure for *T* reads as follows.

Our language has infinitely many sentences, each with its own meaning and truth conditions, and we are able to understand indefinitely many of these sentences upon encountering them for the first time. *T* supposes that to explain these facts our language must have a compositional *meaning* theory and a compositional *truth* theory. A compositional semantics for a language *L* is any finite theory of *L* that generates for each complex expression *e* of *L* a theorem that assigns a semantic value to *e* as a syntax-determined function of the semantic values the theory

assigns to *e*'s constituent morphemes (every morpheme is a "constituent" of itself). A compositional meaning theory for *L* assigns whatever it takes "meanings" to be, according to the role it requires that technical notion to play, and a compositional truth theory for *L* specifies for each sentence of *L* what must be the case in order for an utterance of the sentence to be true. *T* recognizes two kinds of meanings, *characters* and *contents*. Characters are meanings of expression-*types*, contents of expression-*tokens* (or utterances of expressions, or expressions-relative-to-"contexts-of-utterance"/occasions-of-use—these distinctions won't presently matter).

T takes contents to be Russellian propositions or their constituents. The content of a sentence-token is the proposition a speaker would say, and thus mean, in a literal and unembedded utterance of the token. The content of every expression-token is its contribution to the determination of the content of the unembedded sentence-token in which it occurs, and the content of every complex expression-token is a syntax-determined function of the contents of the token's constituent morphemes. An expression is either context-sensitive or context-insensitive. It's context-sensitive just in case its tokens may have different contents, context-insensitive just in case its tokens must have the same content. The character of an expression-type *e* tells us what must be the case in order for something to be the content of a token of *e*. It's common to represent characters as functions that take tokens of an expression to their contents. Thus, the character of a context-*in*sensitive expression is a constant (typically partial) function that takes every token of the expression (for which the function is defined) to the same content.[11] We'll say that context-insensitive expressions have constant characters. The character of a context-sensitive expression is a non-constant function and therefore doesn't map all tokens of the expression onto the same content. We'll say that context-sensitive expressions have non-constant characters. *T* holds a distinctive view about how the contents of context-sensitive expression-tokens are determined. It holds that, but for a few exceptions, the content of a context-sensitive expression-token is at least partly determined by the speaker's intentions. *T* allows that the content of a token of 'I' is the speaker who uttered it and that the content of a token of 'today' is the day on which it was uttered, but after one or two more examples like that (e.g. 'yesterday'), the character of every other context-sensitive expression *e* will entail that in order for something *c* to be the content of a token of *e* the speaker who produced the token must have done so with certain communicative intentions about *c*. For example, the character of 'she' in its use as a demonstrative might be taken to be, roughly speaking, that function *f* such that, for every token τ of 'she' and every x, $f(\tau) = x$ iff x = the female to whom the speaker referred with τ.

[11] To enhance readability I will take the qualification about partial functions to be implicit and simply speak of the character of an expression as mapping the expression's tokens onto their contents.

Although for various purposes it's convenient to represent an expression's character as a function that maps tokens of the expressions onto their contents, it's not necessary. The role of a character is to say what must be the case for something to be the content of a token of an expression, and it can do this simply by providing instances of the schema

x is the content of token τ of expression *e* iff ….

Such instances may be expressed as instances of the more revealing schema

x is the content of a token τ of expression *e* iff

(1) [eligibility condition] …;

(2) [selection condition] ….

The eligibility condition for an expression specifies the condition that something must satisfy in order for it to be eligible to be the content of a token of the expression, while the selection condition selects which of the things eligible to be the content is the content. For example, we might specify the character of the demonstrative 'she' thus:

x is the content of a token τ of expression *e* iff

(1) [eligibility condition] *x* is female;

(2) [selection condition] *x* is uniquely such that the speaker who produced τ referred to *x* with τ.

The schema admits the possibility of two limiting cases. At one limit, the eligibility condition is such that only one thing can satisfy it, and in this case the eligibility condition also does the work of the selection condition. The character of 'I' is like this if the content of a token of 'I' is the speaker who produced it. At the other limit, the eligibility condition imposes no restriction on the kind of thing that can be a content, and in this case the selection condition also does the work of the eligibility condition. I suspect no context-sensitive expression is at this limit, although 'it', when used as a demonstrative pronoun, would be if all it takes for a thing to be the content of a token of 'it' is that the speaker referred to it with the token.[12]

A compositional meaning theory for *L* of the kind *T* proposes becomes a compositional truth theory for *L* when the following two axioms are added to it:

◊ A sentence token τ is true/false iff, for some proposition *p*, *p* is the content of τ and *p* is true/false.

◊ A proposition *p* is true/false iff … [where '…' is replaced by finitely many finitely stated conditions].

[12] I realize that my gloss of the labels *eligibility condition* and *selection condition* isn't exactly precise, but I don't think it needs to be more precise to be useful in the informal intuitive way I intend the labels to be useful.

The discussion in Part I of the act of vague speaker-meaning Tom performed in uttering 'A boy was here a little while ago' turned on his utterances of 'boy', 'here', and 'a little while ago'. Since *T* is an idealized model that treats vague expressions as if they weren't vague, we should expect it to assign characters to those three expressions such that:

(a) The content of every token of 'boy' is the property of being a boy.[13]

(b) The content of a token of 'here' is an area α such that the speaker who uttered the token is within α and referred to α with that token.

(c) The content of a token of 'a little while ago' is a period of time π such that π occurred shortly before the token was uttered and the speaker who uttered the token referred to π with it.

(d) The content of a token of 'A boy was here a little while' is the proposition that something that had φ was in α at a time within π, where φ is the content of the uttered token of 'boy', α the content of the uttered token of 'here', and π the content of the uttered token of 'a little while ago'.

Now let's look at *T* in light of the conclusions reached in Part I about vague speaker-meaning. We immediately notice two important ways in which *T* is inconsistent with those conclusions. The first is that *T* holds both that the content of a sentence token is a proposition and that that proposition is what the speaker says, and thus means, in producing the token on its own, unembedded in another sentence token; but that can be true only if the propositional account of speaker-meaning is correct, and Part I concluded that it wasn't correct. Consequently, if a revision of *T* wants to keep propositions as contents of sentence tokens, then it won't be able to say that those contents are what speakers say. But the notion of a content is a functional notion, defined by the role a semantic theory requires contents to play. That role typically requires them to be the objects of propositional speech acts and propositional attitudes. How should we understand the notion of content when its divorced from those roles? The second inconsistency is that *T* holds that the contents of context-sensitive expressions are but for a few exceptions determined in large part by the speaker's communicative intentions. It was because of this tenet that *T*, being an idealized model that ignores vagueness, would propose characters for the expressions Tom uttered that result in (a)-(d) just above. But we can see from the following ways (a)-(d) are false that speakers' intentions aren't always available to do the work *T* requires them to do:

---

[13] A semantic theory of the kind favored in formal semantics would tell us that [[boy]] = λx . boy(x), i.e. that the meaning of 'boy' is that function of type $\langle e, t \rangle$ that maps a person to 1 (truth) if he's a boy, onto 0 (falsity) otherwise.

- (a) is shown to be false by penumbral shift. For penumbral shift makes it possible for there to be three utterances of e.g. 'Sasha is a boy', by the same speaker in three different contexts, one of which is true, another of which is false, and still another of which is borderline true/false, even though there was no relevant change in Sasha between the time of the first utterance and the time of the third. That, of course, would not be possible if the three uttered tokens of 'boy' had the same content.[14]
- (b) is shown to be false by the fact that Tom definitely didn't refer to any precise or vague area with 'here'.
- (c) is shown to be false by the fact that Tom definitely didn't refer to any precise or vague period of time with 'a little while ago'.
- And (d) is shown to be false three times over: by the falsity of (a), of (b), and of (c).

Now suppose the *T* theorist sets out to accommodate vagueness in a way that leaves intact as much of her idealized model as possible. What might she do? Let's start with the demonstrative adverb 'here'. The discussion in Part I concluded that Tom definitely didn't refer to any area, precise or vague, with the token of 'here' he uttered. If the *T* theorist wants there to be an area that is the token's referent (= its content), or even an area such that it's indefinite whether it's the referent, she will need to supply a character for 'here' that shows how referents can be assigned, if only indefinitely, to tokens of the term without any help from speakers' referential intentions. Can this be done? I think it's doubtful. Recall the hypothesis that some precise area was the area to which Tom referred with 'here'. Suppose there was such an area; call it α^*. Then, we noticed, α^* would have belonged to an extremely dense cloud of indefinitely many nearly identical precise areas that were indistinguishable to Tom, each of which was as eligible as α^* to be the area to which he referred with 'here', yet there would have been a property φ^* such that α^* was the only one of the indefinitely many equally eligible precise areas to have φ^*, and it was by virtue of having φ^* that α^* was the area to which Tom referred with 'here'. I can't conceive of what

[14] Note that it's not an option to accommodate the effect of penumbral shift on 'boy' by saying that it was indefinite which property was *the* property of being a boy, and therefore which constant function was the character of 'boy'. For suppose properties $\varphi_1, \ldots, \varphi_n$ were each such that it was indefinite whether it was the property of being a boy. Then there would be constant functions $f_1, \ldots, f_n$ such that for each $i \geq 1 \leq n$, f_i mapped every token of 'boy' onto φ_i, and each f_i would be such that it was indefinite whether it was the character of 'boy', which would entail that it was indefinite what the token's content was. But that would also be incompatible with penumbral shift, which shows that some tokens of 'boy' definitely don't have the same content.

property φ^* might be and I find it impossible to believe that there is such a magical property, so, until I'm convinced otherwise, I conclude that there definitely was no such property, and, therefore, that Tom definitely didn't refer to any precise area with 'here'. If that is right, then it's also right that no precise area was such that it's indefinite whether Tom referred to it with 'here'; he definitely didn't refer to any precise area with 'here'. Now, might it have been the case that, although not one of the indefinitely many precise areas in question had a property that made it alone out of the indefinitely many other precise areas the one to which Tom referred with the token of 'here' he uttered, nevertheless one of those precise areas had a property that made it alone the referent of the token of 'here' Tom uttered, notwithstanding that Tom definitely didn't refer to it with the token? To my mind that is no less implausible than the suggestion that some unique precise area had a property that made *it* the area to which Tom referred: the same sort of considerations that show that Tom definitely didn't refer to any precise area with the token of 'here' he uttered also show that that token definitely didn't refer to any precise area. And since, as we have already noticed, metaphysically-vague areas, if there are such things, are every bit as finely individuated as metaphysically-precise areas, the line of reasoning just concluded also applies to them: the token of 'here' Tom uttered definitely didn't refer to any area.

Some theorists, however, are evidently committed to rejecting the argument just concluded. The dominant model in formal semantics takes a compositional semantics for a language L to be a compositional type-driven functional-application possible-worlds truth-theoretic semantics for L constructed in terms of the formal semanticist's proprietary language of the typed lambda calculus. This semantic model, like every other foundational semantic model, has been developed as an idealized model that prescinds from the vagueness of vague expressions, but some formal semanticists have attempted to show how at least certain kinds of vague expressions may be accommodated in a compositional semantics of the kind just mentioned. Arguably the most prominent such attempt is the semantics for vague gradable adjectives Christopher Kennedy offers in "Vagueness and grammar: The semantics of relative and absolute gradable adjectives."[15] On the semantics Kennedy offers, the meaning of 'tall' involves a contextually-determined measure function that maps individuals to precise degrees on a scale of measurement for height so that an utterance of 'John is tall' in context c is true if, for some precise degree of height δ onto which John is mapped by the measure function determined by c, John's degree of height $\geq \delta$, false otherwise. Since for Kennedy an utterance of 'John is tall' is either true or false, it's impossible for anyone to know the precise contextually-determined threshold John's height must exceed in order for the utterance to be true. So, as Kennedy recognizes, his account of vague gradable adjectives is an epistemic theory, according to which, although the context in which 'John is tall'

[15] Kennedy (2007).

is uttered determines a measure function which in turn determines a precise threshold that John's height must surpass in order for the sentence to be true relative to that context, it's impossible for anyone to know what that function or threshold is (for if one could know what the function was one could in principle know the threshold it determined, and if one could know that, then one could in principle determine whether the token of 'tall' produced in an utterance of 'John is tall' was true of any given individual, in which case nothing could be a borderline case of a thing of which the token was true, whereas the epistemicist, to account for what makes a predicate vague, must say that every utterance of a vague predicate admits the possibility of borderline cases). Even though it's impossible to know the measure function a context of utterance determines, we should still expect to be told *something* about how a context determines such a function. Yet all Kennedy tells us—by way of telling us how his theory provides a "contextualist" solution to the sorites paradox generated by a vague gradable adjective—is that if an utterance of 'tall' is being applied to a pair of individuals whose heights are very similar, then the context can't determine a measure function that determines a degree of height that only one of the pair of individuals meets or exceeds. In another much cited article,[16] Chris Barker suggests that a dynamic component be added to the semantics of vague gradable adjectives, but his account, like Kennedy's, requires a contextually-determined precise threshold, and, like Kennedy again, he offers no account of how that threshold is determined.[17]

If the token of 'here' Tom uttered definitely didn't refer to any precise or vague area, then, of course, the token of 'a little while ago' Tom uttered definitely didn't refer to any precise or vague period of time. I see no way in which Model *T* can accommodate, or be revised to accommodate, the vagueness of expressions like 'here' and 'a little while ago'. But who really cares about Model *T*? Well, the conclusions just reached about *T* seem to falsify a foundational assumption held by nearly every philosopher of language and semanticist—namely, that our language must have a compositional truth-theoretic semantics, where that is a finite theory of a language *L* that for each sentence of *L* generates a theorem that specifies what must be the case in order for an utterance of the sentence to be true (or false). We can see that this assumption is falsified by conclusions already reached, if those conclusions are correct. Given the description of Tom's utterance of 'A boy was here a little ago', we want to say that it was true—after all, we

[16] Barker (2002).

[17] Una Stojnić and Matthew Stone (forthcoming) claim to offer an account of how the Kennedy-required precise cutoffs are determined. They propose that the work, or the better part of it, can be done by mechanisms of discourse coherence, but while they might show how discourse-coherence considerations affect the determination of thresholds, I don't see that they even attempt to show how Kennedy-required precise thresholds are determined.

know that a human male of at most six years old was within four meters of Tom no more than five minutes before he spoke. If Tom's idiolect enjoyed a compositional truth theory of the kind now in question, then the truth-value of the token of 'A boy was here a little while ago' he uttered would be a function (in part) of the referential semantic values of the sentence token's constituent tokens of 'here' and 'a little while ago'. Those referential semantic values would evidently have to be things to which the tokens referred, or at least things to which the tokens indefinitely referred, and there seem not to be such things.

What about 'boy'? Can the *T* theorist see a way to handle its vagueness? When the theorist was ignoring the fact that 'boy' was vague, and treating the word as though it wasn't vague, she supposed (relative to that idealization) that there was such a thing as *the* property of being a boy and therefore took the character of 'boy' to be a constant function that took tokens of 'boy' to that property. Now that she seeks to show how 'boy' should be treated in a compositional semantics that accommodates vagueness, she appreciates that in order to do that she must take into account the way the word is subject to penumbral shift. Suppose she is initially determined to achieve that accommodation without abandoning *T*'s tenet that every expression (that contributes a semantic value to the truth conditions of utterances of sentences containing it) is either a context-insensitive expression with a constant character or a context-sensitive expression with a non-constant character. Then I believe that the following are the options she will want to consider.

1. She might say that 'boy' has a non-constant character and that, consequently, penumbral shift was nothing more than tokens of a context-sensitive word taking different contents. Let's first suppose the theorist requires contents of tokens of 'boy' to be metaphysically-precise properties, where, you may recall, a property is metaphysically precise just in case, necessarily, everything either has the property or else fails to have it. Then she will need a non-constant character for 'boy' that is specifiable as an instantiation of the schema:

> A property φ is the content of a token τ of 'boy iff
> (1) (eligibility condition) …;
> (2) (selection condition) ….

The eligibility condition for 'boy' would specify what must be the case in order for a property to be a candidate for being the content of a token of the word. Such a property would have to entail the property of being a human male, but since 'male' and 'human' are vague, there is no such thing as *the* property of being a human male. But let's give the *T* theorist a head start by pretending that there is a precise property that is *the* property of being a human male and, to give her even more of a head start, let's also pretend that for each human male there was a precise moment when he came into existence. Then every precise property φ that is eligible to be the content of a token of 'boy' would be such that, for some positive real number n and for any x, x has φ iff x is a human male whose age in milliseconds $\leq n$. But now we can see a problem with getting an eligibility

condition for 'boy'. Although vague expressions are subject to penumbral shift, we know that the shifts that may occur must be within certain limits. For example, a nine-year-old human male must always count as a boy, and a thirty-nine-year-old human male must never count as a boy. So no property can be the content of a token of 'boy' if it's instantiated by a person thirty-nine years old or if it's not instantiated by a nine-year-old human male. But of course the same can be said for indefinitely many other pairs of ages. If there were a real number j that was the largest number such that every token of 'boy' had to be true of every human male whose age in milliseconds $\leq j$ and a real number j' that was the smallest number such that no token of 'boy' could be true of a human male whose age in milliseconds $\geq j'$, then a property φ would be eligible for being the content of a token of 'boy' just in case, for some real number $k \geq j \leq j'$ and for any x, x instantiates φ iff x's age in milliseconds $\leq k$. But of course there can be no such numbers as j and j'; there is no greatest lower bound and least upper bound to determine the class of precise properties eligible to be the content of a token of 'boy'; yet there being such bounds would seem to be what is required for there to be an eligibility condition for 'boy'.

There is an equally big problem in finding a selection condition, and we already know what it is. Suppose there were a correct eligibility condition for 'boy'. Then indefinitely many nearly identical precise properties—e.g. the property of being a human male whose age in milliseconds $\leq$ 531,066,240,000, the property of being a human male whose age in milliseconds $\leq$ 531,066,240,001, the property of being a human male whose age in milliseconds $\leq$ 531,066,240,002, etc.—would be eligible for being the content of a given token of 'boy'. The selection condition would have to say what must be the case for one of those properties to be the content of an arbitrary token of 'boy'. To do this the selection condition would have to specify a property of properties whose possession by just one of the eligible properties explained why it alone of the indefinitely many eligible properties was the token's content. I submit, once again, that it's very implausible that there should be such a magical property. One final point about there being a non-constant character that maps tokens of 'boy' onto precise properties. A theory that claimed that 'boy' had a character, whether non-constant or constant, that was a function from tokens of the word to precise properties would have to be an epistemic theory of the vagueness of 'boy', and therefore one who accepts that theory would have to claim that, if 'boy' has such a character, then it's impossible for anyone to know what it is, so a theorist couldn't justify her acceptance of the theory by *finding* the correct substitution instance of the above-displayed schema. So what might justify her in accepting that the word has such a character? This problem wouldn't arise if the non-constant character determined metaphysically-vague properties as the contents of tokens of 'boy', for then an individual's being a borderline case would be explained by his falling in the penumbra of the property that was the operative content. But there would still be the problem of getting an eligibility condition, for the eligibility condition would still need to

be defined in terms of a greatest lower bound and a least upper bound, and there can be no such bounds, and the problem with getting a selection condition would be the same as it was when contents were required to be precise properties: it would be impossible to justify the claim that just one of the indefinitely many eligible properties had the magical property whose possession made it the content of a given token of 'boy'. I conclude not only that the vagueness of 'boy' can't be accommodated in a Model *T* semantics by assigning it a non-constant character, but also that no compositional truth theory in the style of any model can accommodate the word's vagueness in that way.

2. When she was treating 'boy' as though it weren't vague, the *T* theorist assumed that some property was *the* property of being a boy and said that 'boy' had a constant character that mapped tokens of the word onto that property. When she attempted to accommodate the vagueness of 'boy' she realized that penumbral shift prevented her from saying that, for penumbral shift shows that there may at the same time be two tokens of 'boy' one of which is definitely true of a person while the other of which is definitely false of him. The theorist might, however, attempt to account for penumbral shift by saying that at any given time 'boy' has a constant character in one's idiolect, but that which constant function was its character in one's idiolect was apt to change from one utterance to the next. Then what I have labelled 'penumbral shift' would simply be a manifestation of these shifts in meaning. If she takes this line, the theorist would owe an explanation of why we never know what property is the content of a token of 'boy' and why we never know what character the term has in our idiolect at any given time; and she would have to explain why there are limits to how much the word's character can change (every token of 'boy' must be true of a human male if he is nine years old and false of him if he is thirty-nine years old). If the theorist takes the contents of tokens of 'boy' to be metaphysically-precise properties, then she must do more than explain why we don't ever know what these contents are; in order to accommodate the fact that it's impossible for anyone to identify a precise cutoff in any sorites series, she would also have to explain why it's impossible for anyone to know what these contents are. Perhaps the theorist would offer to explain all these things in part by appeal to Timothy Williamson's views about margins for error and inexact knowledge.[18] The contemplated move would remain implausible for a reason that should by now be familiar: it would be the implausibility of supposing that there was a magical property of properties ψ such that (i) only one property can have ψ at any given time, (ii) whichever property had ψ at a time *t* was the meaning of 'boy' in one's idiolect at *t* (i.e. was the property that had to be the content at *t* of every token of 'boy' produced at *t*), and (iii) at any given moment just one of the indefinitely many nearly identical properties that were equally good candidates for being the meaning of 'boy' in one's idiolect at

[18] Williamson (1994: Chapter 8).

that time—e.g. the property of being a human male whose age in milliseconds ≤ 531,066,240,000, the property of being a human male whose age in milliseconds ≤ 531,066,240,001, the property of being a human male whose age in milliseconds ≤ 531,066,240,002, etc.—had ψ.

3. The theorist might try a supervaluationist move. The supervaluationist on vagueness says that a vague sentence is true just in case it's true under every admissible precisification of the language to which it belongs, false just in case it's false under every admissible precisification, and neither true nor false just in case it's true under some admissible precisification while false under another. A precisification is a bivalent model-theoretic interpretation of the language wherein the set assigned as extension to a vague term includes everything to which the term definitely applies, nothing to which it definitely doesn't apply, and may divide everything else between the term's extension and counterextension in any which way, provided it respects those analytical connections between terms that Kit Fine called penumbral connections.[19] A precisification is admissible just in case it respects all penumbral connections, so, for example, if a precisification assigns Ed to the extension of 'tall' and Ted is taller than Ed, then that precisification is admissible only if it also assigns Ted to the extension of 'tall'. What might the *T* theorist say about the *meaning* of 'boy' in order to get her supervaluationist account of the truth conditions of vague sentences? One thing she might say is that, while no proposition is definitely the meaning of e.g. 'Sasha is a boy', two or more are each such that it's indefinite whether it's the sentence's meaning. These propositions would then determine the sentence's admissible precisification, and the sentence would be true iff each of the propositions it indefinitely meant is true, false iff each is false, and neither true nor false iff at least one off those propositions is true and another one of them is false. But this won't do, for those propositions would be like the proposition *that Sasha is a human male whose age in milliseconds ≤ 531,066,240,003*, and we know that the sentence definitely doesn't mean that proposition. A better way to implement supervaluationsim at the level of meaning would be to take the meaning as well as the extension of a term to be relative to an interpretation that assigns meanings as well as extensions. An interpretation *I* would then assign a bivalent property φ to 'boy' as its meaning relative to *I* provided that φ entailed every property a thing must have in order for 'boy' to be definitely true of it and also respected all penumbral connections. Since φ has to entail every property a thing must have in order for 'boy' to be definitely true of it, it would *a fortiori* have to entail the complement of every property a thing must have in order for 'boy' to be definitely false of it. This revised version of model *T* would say that 'boy' has no complete meaning absolutely but only relative to a model-theoretic meaning interpretation of the language. The word would, however, have a partial meaning that determined which properties were eligible to be assigned as its meaning

[19] Fine (1975).

in an admissible precisification. On this version of supervaluationist meaning, the fact that a property is definitely not the meaning of 'boy' is no barrier to its being the meaning assigned to the word in an admissible precisification.

Supervaluationsim also has some features that aren't so nice. Some of these are about its semantics for quantifiers and connectives, and thus about its solution to the sorites, another about its inability to explicate vagueness since it takes a notion of definite truth as primitive and defines the notion of a precisification in terms of it. But only three problems are directly relevant to the issues of this paper. One is that any proposal for a supervaluationist semantics for a language will have apparent counterexamples.[20] For example, consider a true utterance of

(*) Alice wouldn't find Bob attractive if he weren't bald.

Supervaluationism entails that the utterance of (*) is true only if it's true on every admissible precisification of it, but, although (*) is true, every admissible precisifications of it—e.g.

Alice wouldn't find Bob attractive if he didn't have fewer than 1,593.02 hairs on his scalp

—is definitely false. Another counterexample is this nice one from Crispin Wright:

No-one can knowledgeably identify a precise boundary between those who are tall and those who are not.[21]

(A few philosophers have argued that these examples aren't the counterexamples to supervaluationism they are alleged to be.[22] Their arguments deserve to be addressed, but it won't be possible for me to address them here.) A second problem begins to unfold with the observation that, as stated above, the supervaluationist treatment of 'boy' requires there to be real numbers j, j' such that 'boy' is definitely true of a human male iff his age in milliseconds $\leq j$ and definitely false of him iff his age in milliseconds $\geq j'$. But of course there are no such numbers. This is sometimes said to be an issue about higher-order vagueness that is due to the fact that 'admissible precisification' is vague in the metalanguage in which a supervaluationist semantics for an object language L is constructed, and that the vagueness of 'admissible precisification' merely shows that it's indefinite—or indefinite whether it's indefinite, or indefinite whether it's indefinite whether it's indefinite, or …which supervaluationist truth (or meaning) theory for L is correct. I don't think this can be right, however: each such semantics would definitely not be correct, for it would entail that were a pair of boundary-fixing real numbers of the kind just described, and it's not indefinite, or indefinite whether it's indefinite, etc. whether there are such numbers: we know that there definitely are not such numbers. A third problem is that as so far described supervaluationsim does nothing to accommodate the effect of penumbral shift on 'boy', but the most obvious way of

[20] I've offered various counterexamples in various publications, beginning with Schiffer (1998).
[21] Wright (2004).
[22] See e.g. Weatherson (2003), García-Carpintero (2010), and Keefe (2010).

doing of that would be to have an admissible precisification of the word assign it a non-constant function as its meaning, and, I submit, we've seen why that can't work.

The conclusion now to be drawn is not just about Model *T*: it's about every semantic model according to which the languages we speak must have compositional truth-theoretic semantics, which by my reckoning must be every model of which I'm aware. By a compositional truth-theoretic semantics for a language *L* I mean a finite theory that generates for each of the infinitely many truth-evaluable sentences of *L* a theorem that specifies, for each truth-value the theory recognizes, what must be the case in order for a token of the sentence to have that truth-value, and, of course, to have it as a function of semantic values of its constituent expression tokens. This will require that, for each expression *e* of *L* (whose tokens must have semantic values if expression tokens containing them are to have semantic values), the theory must determine, by axiom or theorem, a condition that specifies what must be the case for something to be the semantic value of a token of *e*. I submit that, if what I have said up to this point is correct (a pretty big if, to be sure), then we have good reason to doubt that our languages have compositional truth-theoretic semantics. For consider the perfectly ordinary general noun 'boy' as it occurs in Jane's English idiolect, *J*. If *J* enjoys a compositional truth-theoretic semantics, then it will contain an axiom which is equivalent to a completion of

(C) The character of 'boy' in *J* = that function *f* such that, for any token τ_b of 'boy' produced by Jane and for any property φ, $f(\tau_b) = \varphi$ if … τ_b … φ …, but is undefined otherwise.

It will then use (C) to generate for, say, the sentence 'Sasha is a boy' in Jane's idiolect the theorem:

(T) A token τ_s of 'Sasha is a boy' produced by Jane is *true/false* if, for some φ, token τ_b of 'boy', and function *f*, τ_b is the token of 'boy' in τ_s, *f* is the character of 'boy', $f(\tau_b) = \varphi$, and Sasha has/doesn't have φ; otherwise no truth-value is defined for τ_s.

But if the conclusions I reached about 'boy' are correct, then there can be no such axiom as (C) and therefore no such theorem as (T). And if that is correct, then the same goes for every other vague predicate, and nearly every predicate is vague.

III. Vague Expression-Meaning: What Might It Be?

Many philosophers of language and semanticists seem to take it to be a datum that every spoken language has a compositional truth-theoretic semantics (a compositional semantics, for short). But it's not a datum. To be justified in believing that every spoken language must have a compositional semantics, one must be justified in believing that there is a feature shared by all

spoken languages that can't be *explained* unless each spoken language has a compositional semantics. It's not easy to say what a compositional semantics is needed to explain.

It might be suggested that:

> Speakers of a language *L* are able to know the truth conditions of indefinitely many sentences of *L*, sentences constructed from the same finite stock of morphemes and syntactic structures. This can be explained only on the assumption that *L* has a compositional semantics.

But we need a more compelling answer than that. It's not clear in what sense we're supposed to "know" the truth conditions of novel sentences, nor what "truth conditions" are supposed to be such that we know them; and although it's true, and thus a fact to be explained, that in some sense of 'understand' speakers of a language are able to understand indefinitely many of its novel sentences, it's far from clear in what that understanding consists, and even less clear that there is any sense in which it requires knowing the truth conditions of those sentences. But the reason the displayed claim doesn't tell us what a compositional semantics is needed to explain is that the claim's talk of knowing the truth conditions of, or understanding, novel sentences is a red herring. At most it gives a reason to suppose that knowing a language requires having "tacit" knowledge of a compositional semantics for that language, but if the claim implies that the understood language has a compositional semantics, it's only because it implies that infinitely many sentences of the language *have* truth conditions that are determined by the words and syntactic structures that compose those sentences, and that *that* is what a compositional semantics is needed to explain. To this it might be replied that:

> Infinitely many sentences provide true substitution instances of the schema
>
> Necessarily, '*S*' is true iff *S*.
>
> One famous example is:
>
> Necessarily, 'snow is white' is true iff snow is white.
>
> This can be explained only on the assumption that English has a compositional semantics.

Of course the mere fact that infinitely many true sentences are instances of the displayed *T*-schema doesn't *per se* imply that English has a compositional semantics: there are infinitely many true instances of the schema

Necessarily, '*n* is a prime number' is true iff *n* is a prime number,

but while that plausibly requires a certain tiny fragment of English to have a compositional semantics, it doesn't show that *English* needs a compositional semantics. The issue would be discussable, however, if, notwithstanding the vagueness of 'snow is white', the instance of the *T*-schema

(*S*) Necessarily, 'Snow is white' is true iff snow is white

were true. Many theorists do believe that (*S*) is true, notwithstanding the vagueness of 'snow is white'. For, they would say, the vagueness of the sentence used on the right-hand side cancels out the vagueness of the sentence referred to on the left-hand side, in that the former matches the latter "umbra for umbra and penumbra for penumbra."[23] These theorists might even go on to say that the same canceling-out goes on in

Necessarily, 'snow' refers to snow.

Necessarily, 'white' refers to whiteness.

and that this shows that it's easy to derive (*S*) in a compositional semantics. *Penumbral shift*, however, shows this argument to be mistaken. The vague sentence 'snow is white' is subject to penumbral shift, and the context-sensitivity this induces entails that every token of (*S*) is false: it's false in something like the way that

The sentence type 'It's a prime number' is true iff 2 is a prime number

is false. A context-sensitive sentence doesn't have truth conditions; only its tokens do, and they may have different truth conditions. A sentence-type has truth conditions only if each of its tokens have those truth conditions. So consider the just-displayed token of (*S*), which I'll call S_τ. Since the token of '"snow is white"' on the left-hand side of S_τ refers to the sentence *type* 'snow is white', S_τ is true just in case every token of the sentence type 'snow is white' must have the same truth conditions as the token of 'snow is white' on the right-hand side of S_τ, but that is what penumbral shift shows not to be the case. For, owing to penumbral shift, there may be two utterances of 'snow is white' by the same person in different contexts of utterance, one of which is true, the other false, even though snow underwent no change of color between the time the first utterance occurred and the time the second one occurred. Likewise, the sentence token (call it *W*)

Necessarily, 'White' is true of a thing iff it's white

is false. For *W* is true only if the word type 'white' referred to on its left-hand side is such that a token of it is true of a thing iff the token of 'white' on *W*'s right-hand side is true of it; but, owing to penumbral shift, one token of 'white' might be true of a thing while at the same time a second token of it wasn't true of it.

In order to determine whether a compositional semantics for English is needed to explain how English sentences have the truth conditions they have—or, more to the point, in order to see whether that might be explained even though those truth conditions can't be compositionally generated—we need first to determine the sense in which they do have truth conditions. Let's see what sort of truth conditions can be ascribed to the sentence 'Sasha is a boy', wherein 'Sasha'

[23] Quine (1960: 37). Quoted in Lepore and Ludwig (2005: 139).

uniquely refers to a particular person, when we suppose that the sentence belongs to a language that is the idiolect of each of us.

We know that there is no condition *C* such that it's both (i) knowable whether *C* obtains and (ii) knowable that

Necessarily, a token of 'Sasha is a boy' is true *if, and only if, C*(Sasha).

(For 'boy' wouldn't be vague—wouldn't feature in a sorites paradox—if (i) and (ii) were true.) But we do know that there are some people of whom every token of 'boy' must be definitely true (e.g. your nine-year-old son Charlie), and some people of whom every token of 'boy' must be definitely false (e.g. your daughter Matilda), so it might seem reasonable to suppose that there are conditions *C*, *C′* semantically associated with 'boy' such that it's definitely true that

Necessarily, a token of 'Sasha is a boy' is true *if C*(Sasha)

and that

Necessarily, a token of 'Sasha is a boy' is true *only if C′*(Sasha).

And if it's reasonable to think there are such conditions, then shouldn't it also be reasonable to think that a compositional theory could tell us what they are and explained what made them the sufficient or necessary conditions they are, in part by ascribing some sort of semantic property to the word 'boy'?

Yes, that would be *very* reasonable, *but*, first, there are no such conditions, and second, if there were we couldn't know what they were, and thus couldn't construct a compositional semantics in which they featured. The reason we couldn't know what these conditions were is that they would have to be metaphysically-precise or metaphysically-vague properties, and, as we (ahem) know from Part I, we're not able to think of such properties in a way that would enable us to specify or intend to refer to them. The more serious problem is that there are no such conditions. For while every token of e.g.

Necessarily, a token of 'Sasha is a boy' is true *if* Sasha is a six-year-old human male

and of e.g.

Necessarily, a token of 'Sasha is a boy' is true *only if* Sasha is not female.

is definitely true, there is no condition *C* expressed or determined by any token of 'is a six-year-old human male' such that

Necessarily, a token of 'Sasha is a boy' is true *if C*(Sasha)

and there is no condition *C′* expressed or determined by any token of 'is not female' such that

Necessarily, a token of 'Sasha is a boy' is true *only if C′*(Sasha).

This is because, being vague, 'boy', 'six-year-old human male' and 'not female' are subject to penumbral shift, and while no token of either expression denotes any condition (= property) at all,

there are limits to how much the penumbral profiles of those expressions can change from one token to the next, so that, for example, every token of 'boy' must be true of a person if any token of 'six-year-old human male' is true of him, and so that no token of 'boy' can be true of a person unless every token of 'not female' is true of him. Similarly, 'Midtown', in its use that purports to be the name of a part of Manhattan, is such that, while no token of 'Midtown' even indefinitely refers to a part of Manhattan, every token of 'I'm in Midtown' is definitely true when uttered by a person in Times Square. That is how a token of 'Sasha of a boy' can be true, but not as a syntax-determined function of any semantic values of the sentence token's parts.

Very well, but if we can't have a compositional theory that answers the question posed in the subtitle of Part III, how might we answer it? This is a very difficult question to answer, but it will help if we first get rid of the weasel word 'meaning' in 'vague expression-meaning'. Better—or at least a little clearer—to use the label 'semantic properties of vague expressions'. A theorist who believes that every spoken language must have a compositional truth-theoretic semantics, knows what *she* has to do—viz. figure out the properties that vague expressions of a language must have in order for them to be accommodated in a correct compositional semantics for that language. But I've suggested that we've yet to be given a good reason to suppose that a spoken language must have a compositional semantics, and that we have been given reason to doubt that compositionality hypothesis. So, now that we're—I mean *I*'m—unmoored from the constraints imposed by the compositionality hypothesis, we can see some of the difficult questions a meaning theory for a vague language—that is to say, for any spoken language—must answer when it's precluded from being, or requiring, a compositional truth-theoretic semantics for the language. For example, to mention just a few:

(1) The token of a 'A boy was here a little while ago' Tom uttered was definitely true, and that is in large part due to its contained tokens of 'boy', 'here', and 'a little while ago', but the token of 'boy' referred to/expressed no property, the token of 'here' referred to no area, and the token of 'a little while ago' referred to no period of time. What do the vague expression-tokens in the sentence token contribute to the sentence token such that those contributions are essential to the token's being definitely true?

(2) If I say 'being a six-year-old male child is sufficient for being a boy', I say something that's definitely true. But the uttered token of 'being a six-year-old male child' refers to no property, and therefore states no sufficient condition for being a boy. How, then, can the displayed sentence token be true?

(3) If Jill says 'Sasha is a boy' and you say

Jill said that Sasha was a boy

or

Jill knows that Sasha is a boy

Then you give a true report of what it is that Jill said and of what it is that she knows. This evidently shows that neither

Jill said/believes/knows something

Jill said/believes/knows that Sasha is a boy

entails

$\exists x$(Jill said/believes/knows x).

How can that be? What does it imply about the logical form of saying reports and propositional-attitude reports—or is it that sentences don't have "logical forms" if the languages to which they belong don't have compositional truth-theoretic semantics? But then how are we to explain the validity or invalidity of arguments constructed from those sentences?

(4) Penumbral shift is puzzling for at least the following reason. The penumbral profile of a token of 'boy' determines the token's conditions of application. Penumbral shift occurs when different tokens of 'boy' have different penumbral profiles when that isn't explained by a change in the meaning of 'boy'. It's natural to suppose that when the meaning of a general term is consistent with tokens of the term having different penumbral profiles, the word's meaning is a non-constant character that determines its tokens to have different penumbral profiles under specified conditions. What makes penumbral shift interesting, and puzzling, is that a term like 'boy' seems not to have a non-constant character. Yet it's because the meaning of e.g. 'boy' is instrumental in determining the penumbral profiles of the word's tokens that every token of the word must be true of every six-year-old human male and not true of any sixty-year-old human male. The big question is what determines the penumbral profile of a token of 'boy' after the word's meaning has done its work? Another big question is whether a systematic account of those determinants is possible. That's a big question because if a systematic account of it isn't possible, then it's hard to see how any sort of systematic semantics for 'boy' is possible. Sometimes it's clear that conversational features, such as the topic under discussion, are instrumental in causing penumbral shift. David Lewis's

> "rule of accommodation"[24] offers an explanation of one way penumbral shift may occur. But sometimes penumbral shift seems to be due to a speaker's whims, mood swings, or other kinds of conceptual noise. To the extent that that is so, a systematic account of the semantic features of vague expressions may be mostly out of reach. Still, we should want to know how penumbral profiles are caused, even if it's mostly by a random process. The task of getting that knowledge, however, would appear to be outside the philosopher's purview. In any case, there is an *a priori* reason for being skeptical of finding a systematic account of what determines the penumbral profiles of a vague expression's tokens—namely, that it's impossible to know what a token's penumbral profile is. For a token's penumbral profile includes the token's penumbra, the area, itself vague, within which the term's application is anything other than definitely correct or definitely incorrect. Thus, a token's penumbra includes the token's borderline cases, borderline cases of its borderline cases, borderline cases of its borderline cases of borderline cases, and so on. Because there can be no definite set of these iterations—any attempt to specify such a set will itself admit borderline cases.

If a "semantics" for a language must compositionally determine semantic values for complex expressions on the basis of semantic values determined for its parts, then no spoken language has a "semantics." What we might be able to have instead of a "semantics" for a language is a theory that, though it doesn't compositionally generate semantic values from semantic values, nevertheless is able to explain how the semantic properties of a complex expressions are determined by contributions its constituent expressions make to its having those properties, even though what the parts contribute aren't semantic values of their own. What sort of theory might accomplish that? I would answer this in refined detail were it not for the fact that it would make this paper too long for the Mind and Language seminar. But if you like cryptic hints, here's one. Pretend we think in (a neural version of) the language we speak. Then the theory will explain how the conceptual roles of complex expressions are a function of the conceptual roles of its parts, and how these conceptual roles are related to objects, properties, and states of affairs composed of them in ways that explain what makes a definitely true token of a sentence such as 'A boy was here a little while ago' definitely true.

---

[24] Lewis (1979).

Works Cited

Barker, C. (2002). "The Dynamics of Vagueness," *Linguistics and Philosophy*, 25(1): 1-36.

Barnett, D. (2011). "Does Vagueness Exclude Knowledge?" *Philosophy and Phenomenological Research* 82 (1): 22-45.

Beal, J. (ed.). (2004). *Liars and Heaps* (Oxford University Press).

Buchanan, R. (2010). "A Puzzle about Meaning and Communication," *Noûs*, 44(2): 340-71.

Buchanan, R. and Ostertag, G. (2005). "Has the Problem of Incompleteness Rested on a Mistake?" *Mind*, 114(456): 889-913.

Davidson, D. (1968). "On Saying That." *Synthese* 19: 130-146.

Dietz, R. and Moruzzi, S. (eds.). (2010). *Cuts and Clouds: Vagueness, Its Nature, and Its Logic* (Oxford University University Press).

Dorr, C. (2003). "Vagueness without Ignorance," *Philosophical Perspectives* 17 (1): 83-113.

Frege, G. (1892). "On Sense and Reference," *Zeitschrift für Philosophie und Philosophische Kritik*, 100: 25-50.

García-Carpentero, M. (2010). "Supervaluationism and the Report of Vague Contents." In Dietz and Moruzzi (2010): 345-359.

Heck, R. (2002). "Do Demonstratives Have Senses?" *Philosopher's Imprint* 2.

Keefe, R. (2010). "Supervaluationism, Indirect Speech Reports, and Demonstratives." In Dietz and Moruzzi (2010): 360-382.

Keil, G. and Poscher, R. (eds.). (2016). *Vagueness and Law: Philosophical and Legal Perspectives* (Oxford University Press).

Lepore, E. and Ludwig, K. (2005). *Donald Davidson: Meaning, Truth, Language, and Reality* (Oxford University Press).

Lewis, D. (1979). "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8 (1): 339—359.

McDowell, J. (1998a). "De Re Senses," in McDowell (1998b):. 214-27

_____ (1998b). *Meaning, Knowledge, and Reality* (Harvard University Press).

Miller, A. (ed.) (forthcoming). *Logic, Language and Mathematics: Essays for Crispin Wright, Volume II* (Oxford University Press.

Quine, W. (1960). *Word and Object* (MIT Press).

Schiffer, S. (1998). "Two Issues of Vagueness," *The Monist* 88: 193-214

_____ (2010). "Vague Properties," in Dietz and Moruzzi (2010): 109-130.

_____ (2016). "Philosophical and Jurisprudential Issues of Vagueness," in Keil and Poscher (2016): 23-48.

_____ (forthcoming). “Quandary and Intuitionism: Crispin Wright on Vagueness,” in Miller (forthcoming). Miller (ed.),
Stojnić, U. and Stone, M. (forthcoming). “Vagueness in Context.”
Weatherson, B. (2003). “Many many problems.” *Philosophical Quarterly* 53.
Williamson, T. (1994). *Vagueness* (Routledge).
_____ (1997). “Précis of *Vagueness*,” Philosophy and Phenomenological Research 62(4): 921-8.
Wright, C. (2004). “Vagueness: A fifth column approach,” in Beal (2004): 84-105.
_____ (2010). “The Illusion of Higher-Order Vagueness,” in Dietz and Moruzzi (2010).