

Expression-Meaning and Vagueness

(Revision of July 31, 2019)

Stephen Schiffer

New York University

Abstract

Brian Loar attempted to provide the Gricean program of intention-based semantics with an account of expression-meaning. But the theory he presented, like virtually every other foundational semantic or metasemantic theory, was an idealization that ignored vagueness. What would happen if we tried to devise theories that accommodated the vagueness of vague expressions? I offer arguments based on well-known features of vagueness that, if sound, show that neither Brian's nor any other extant theory could successfully make that adjustment, and this because, if sound, the arguments show not only that nothing can be the content of a vague expression, but also that standard assumptions about semantic compositionality are mistaken. This raises the question of what are the facts about a language whose explanation might seem to require semantic compositionality, and what, really, is needed to explain those facts.

In "Two Theories of Meaning," the most important of his papers in the philosophy of language, Brian Loar asked:

What facts about a sentence constitute its meaning—that is, its meaning in abstraction from any particular utterance, and in so far as it belongs to the language of a particular population?

The answer he proceeded to give was a theory of expression-meaning in support of the Gricean program of intention-based semantics (IBS). IBS is a two-stage program for reducing all questions about the intentionality of speech acts and linguistic expressions to questions about the intentionality of thought: in the first stage a certain notion of *speaker-meaning* is to be taken as foundational in the theory of meaning and defined, without recourse to any public-language semantic notions, in terms of acting with certain audience-directed intentions, and in the second

stage the semantic features of linguistic expressions are to be defined in terms of the now defined notion of speaker-meaning, together with certain ancillary notions, such as that of *convention*, also to be defined in terms of nonsemantic propositional attitudes. A widely-held view about the relation between language and thought is that it is only our intentional mental states—believing, intending, and the like—that have *original* intentionality, intentionality that doesn't have its source in something else's intentionality; the intentionality of speech acts and words is *derived* intentionality, intentionality inherited from that of associated mental states. But that derivation story is plausible only if there is a plausible explanation of *how* speech acts and words inherit their intentionality from the propositional attitudes implicated in linguistic behavior. IBS aims to provide that plausible explanation.¹ In the late sixties, when Brian began writing his D. Phil. thesis on *Sentence Meaning*, detailed revisions by Grice (1969) and myself (1972) of Grice's 1957 account of speaker-meaning were available to explain how propositional speech acts inherit their intentionality from the propositional attitudes in terms of which they, the speech acts, were defined, but there was nothing close to a comparably detailed account to show how expression-meaning was definable in terms of a Gricean account of speaker-meaning. The account of expression-meaning Brian advanced in his dissertation was intended to fill that gap, and that is the account a revised version of which finally appeared in print in "Two Theories of Meaning." That account, in my estimation, is probably as good an account of expression-meaning as IBS can hope to achieve. But it shares a certain feature with every other IBS theory, and, as we'll see, that feature proves to be the banana peel under every IBS theorist's heel. The feature to which I allude is that every theory constructed under the IBS banner ignores vagueness, as though it simply didn't exist, and thus makes no attempt to accommodate the vagueness of vague expressions or vague acts of speaker-meaning, notwithstanding that virtually every sentence we utter and virtually every act of speaker-meaning is vague.

Of course, proponents of IBS knew what they were doing. Part of the reason they ignored vagueness was that not to ignore it would require them to have a theory of vagueness, which they didn't have, and couldn't have had, since no IBS theorist had ever worked on vagueness. Yet that

¹ Much of the appeal of IBS for Brian lay in what he saw as its place in reducing *all* intentional notions to nonintentional notions. Once IBS succeeded in reducing semantic intentionality to mental intentionality, the task, as Brian conceived it, was then to reduce mental intentionality, and thereby all intentionality, to functional notions that enjoy wholly physical realizations. The aim of his book *Mind and Meaning* (1981) was to fulfill that physicalistic ambition.

doesn't explain why IBS theorists developed their semantic theories under the pretense that there were no vague expressions or vague speech acts, for they might have concluded that, since virtually every expression and speech act was vague, it would be reckless to devise theories of meaning, reference, and speech acts when one had no idea how the vagueness of an expression affected its semantics or how the vagueness of an act of speaker-meaning affected what was meant in its performance. Evidently, however, Grice, Brian, and I must have believed that our ignoring vagueness was nothing worse than a useful idealization, akin, perhaps, to Galileo's ignoring friction in his theories of bodies in motion. *Is that what we thought?* Actually, I'm not sure any of us ever considered the question; at least I can't *recall* a single time in the countless conversations I had about IBS issues with Grice and Brian in the twenty years between 1965 and 1985 that the question even came up. At the same time, I suspect that if the question had been put to us we would have come up with some version of the harmless-idealization line. If that is right, then we were wrong: vagueness, as I will here try to show, defeats the program of intention-based semantics. Brian's account of expression-meaning makes it possible to sketch the architecture of a theory that realizes the IBS program, and once we have that we'll be positioned to see why vagueness defeats IBS. Now, in idealizing vagueness out of their theories IBS theorists were in good company, for virtually every presentation of a foundational semantic or metasemantic theory ignores vagueness, and in seeing why vagueness defeats IBS we'll also see why it defeats pretty nearly every other extant theory of meaning.

I. The Architecture of IBS

That architecture has the following parts.

a. The most foundational IBS notion is that of *speaker-meaning*; this is where the semantic makes its first contact with the psychological. Speaker-meaning is the most general kind of illocutionary act, the genus of which all other kinds of illocutionary acts are species (see Schiffer 1972: Ch. IV). There are interrogative and imperatival acts of speaker-meaning, but present purposes will be adequately served if we identify speaker-meaning with *assertoric* speaker-meaning (henceforth simply speaker-meaning), i.e. meaning that such-and-such is the case, as when in uttering 'J'ai faim' Odile meant that she was hungry. IBS takes speaker-meaning to be a relation, *S meant p*, between a person *S* and a proposition *p* that she meant, where a proposition is any 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.² That relation, as already noted, is then defined wholly in terms of acting with certain audience-directed intentions. There is no official IBS definition of speaker-meaning, and IBS theorists may disagree among themselves about how exactly the details of that definition should go, but the following toy conditions will highlight the features of an IBS account of speaker-meaning most relevant to this essay:

For any person S , proposition p , and utterance x ,³ S meant p in uttering x iff for some person A and feature φ , S intended it to be mutual (or common) knowledge between A and S that x has φ and, on the basis primarily of that, mutual knowledge that S uttered x intending her utterance of x to result in A 's actively believing p .

Let's call the value of ' φ ' in an act of speaker-meaning its *inference-base feature* (its *IB-feature*, for short). Acts of speaker-meaning are typically performed by uttering sentences of a language common to the speaker and her audience, and, according to IBS, the IB-features of those sentences are their meanings. If you utter 'It's snowing' to communicate that it's snowing, the IB-feature of the sentence 'It's snowing' on which you rely is the meaning of that sentence in English. The reason IB-features are typically meaning properties is that meaning properties are *optimal* IB-features: if you want to tell your child that it's snowing, you would do much better to utter 'It's snowing' than to attempt to communicate that it's snowing by impersonating a snowflake. At the same time, a *sine qua non* of the IBS account of speaker-meaning is that the only intentional notions mentioned on its right-hand side are ordinary propositional-attitude notions, and, consequently, it's not a necessary condition for a person's meaning a proposition that what she utters have any sort of semantic property as its IB-feature. For example, during a lecture one might communicate to one's friend that one is bored by closing one's eyes and pretending to snore. The fact that the IB-feature needn't be a semantic feature is (supposed to be) what makes it in principle

² As we will later see (in §IV), 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 they can be assigned in a compositional truth-theoretic semantics, but for simplicity of exposition I will often 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.

³ Here, following Grice (1957), I use 'utterance' and its cognates as a technical term that applies both to linguistic and to non-linguistic items and behavior.

possible to define expression-meaning in terms of speaker-meaning. The IBS strategy for defining expression-meaning is to define it in terms of certain conventional regularities in communicative behavior that *explain* why a sentence's having the meaning it has in a certain population of speakers makes uttering that sentence an optimal way for a member of the population to communicate to another of its members a proposition that conforms to that meaning.

b. Having defined speaker-meaning, the next thing on the IBS agenda is defining all other agent-semantic notions in terms of the now defined notion of speaker-meaning. Here the most important notions are *illocutionary act* and *speaker-reference*. IBS dispatches illocutionary acts simply by identifying each kind of illocutionary act with a particular kind of act of speaker-meaning (see Schiffer 1972, Ch. IV). As for speaker-reference, there are a few interrelated notions of speaker-reference, but the two most important for present purposes are notions of what I will call *primary* (or first-order) speaker-reference (Schiffer 1981b; see also Neale 2016 and Schiffer 2016). I'll call the first of these notions *referring-in* (*referring_i*) and define it thus:

S referred_{*i*} to *o* in (the course of) uttering *x* iff, for some *o*-dependent proposition *p_o*, *S* meant *p_o* in uttering *x*.⁴

For example,

In nodding her head in response to the question “Is Henry coming to the party?”, Jane meant the Henry-dependent proposition *that Henry was coming to the party*, and thus referred_{*i*} to Henry in nodding her head.

In uttering ‘He’s coming to the party’ in response to the question “Is Henry coming to the party?”, Jane meant the Henry-dependent proposition *that Henry was coming to the party*, and thus referred_{*i*} to Henry in uttering that sentence.

Although in both examples Jane referred to Henry, there is an important difference between them: in the first example there was no part of Jane’s “utterance” (in this case, her nodding) *with which* she referred to him, but in the second example there was; there she referred to him *with the*

⁴ An *o*-dependent proposition is one that is individuated partly in terms of *o* and wouldn’t exist if *o* didn’t exist.

demonstrative pronoun ‘he’. *Referring-with* ($referring_w$) is the second of the two notions of primary reference to which I alluded, and it may be defined (at least to a first approximation) thus:

In uttering x , S referred_w to o with the i^{th} occurrence of e in x , e_{ix} , iff for some property φ and audience A , S uttered x intending it to be mutual knowledge between S and A that e_{ix} has φ and, on the basis primarily of that, mutual knowledge that S referred_i to o in uttering x .⁵

Primary speaker-reference, then, is *de re* speaker-meaning. To refer to a thing o in an act of primary speaker-reference is to mean some proposition *about* o . In $(n > 1)^{\text{th}}$ order speaker reference the speaker refers to a thing not in order to mean something about it, but in order to refer to a thing to which she is making an $(n - 1)^{\text{th}}$ order reference. For example, in uttering

That boy next to Mary is noisy

the speaker is making a primary reference to *that boy next to Mary* and a secondary reference to *Mary*; and in

That boy next to that man talking to Mary is noisy

the speaker is making a primary reference to *that boy next to that man talking to Mary*, a secondary reference to *that man talking to Mary*, and a tertiary reference to *Mary*. And so on.⁶

c. Having defined speaker-meaning and all other agent-semantic notions in terms of nonsemantic propositional attitudes, the one remaining thing on the IBS agenda is to define expression-meaning in terms of the already constructed IBS definitions. The strategy for doing this that best suited IBS, and the one Brian adopted, was inspired by David Lewis (1969). That strategy has two parts. In the first part, a language is taken to be a certain kind of abstract object that may or may not be used by anyone—namely, a finitely specifiable function from infinitely many finite strings of sounds (marks, gestures, neural states, whatever) to things of the kind meanings must be. Now, a language may be used in more than one way. Of special importance to an understanding of IBS is that a language may be used in one, or both, of the following two ways. First, it may be used by a population as its *public language*, the language members of the

⁵This borrows from Stephen Neale’s improvement of a definition I gave in (1981b) (see Neale (2016), p. 281).

⁶Schiffer (1981b) contains a recursive definition of speaker-reference devised by Brian and myself which accommodates speaker-reference at any order.

population use to communicate with one another. Second, if, as many suppose, we think—i.e. process information—in a neural system of mental representation, one’s “brain’s language of synaptic interconnections and neural spikes” (Lewis 1983, p. 346), then a language may be used by a person as her *language of thought*. So an expression ε can’t simply mean μ for x ; ε can only mean μ for x by meaning μ in a language x uses as a public language or in a language x uses as a language of thought (or in a language x uses in some other way). Since use as a public language and use as a language of thought are the only two uses that need for present purposes to come into play, it will be convenient to stipulate that:

ε means _{P} μ for x iff for some language L and population Q , (i) $L(\varepsilon) = \mu$,
(ii) L is the (or a) public language of Q , and (iii) x is a member of Q .

ε means _{T} μ for x iff for some language L , (i) $L(\varepsilon) = \mu$ and (ii) L is x ’s language of thought.

Let’s say that the *public-language relation* is that relation that must obtain between a language L and a population Q in order for L to be the public language of Q , and let’s say that the *language-of-thought relation* is that relation that must obtain between a language L and a person x in order for L to be x ’s language of thought. Then we see that to define expression-meaning _{P} is just to define the public-language relation, and that to define expression-meaning _{T} is just to define the language-of-thought relation. The expression-meaning that IBS aims to define in terms of the propositional attitudes that define speaker-meaning is expression-meaning _{P} . So the second part of the Lewis-inspired IBS strategy is to define the public-language relation, and for the IBS theorist this will require defining that relation in terms of the way members of P use L to perform acts of speaker-meaning.

d. An IBS account of expression-meaning _{P} requires more than defining expression-meaning _{P} in terms of the propositional attitudes that define speaker-meaning. This is because IBS aims not merely to *define* linguistic intentionality in terms of mental intentionality; it aims to *reduce* linguistic intentionality to mental intentionality. This reductionist aspiration can be achieved only if expression-meaning _{T} can be explicated without any appeal to the notion of expression-meaning _{P} ; that is to say, only if the language-of-thought relation can be defined in terms of notions whose explications require no recourse to the public-language relation. This may

be explained in the following way. IBS theorists have no trouble accepting the hypothesis that we think in a language of thought, provided it entails some refinement of the idea that:

Necessarily, for any proposition p and person x , x believes p just in case
a sentence which means _{T} p for x is tokened in x 's belief box.⁷

And when a proposition p 's being the content of a belief is defined in terms of a mentalese sentence's meaning _{T} , then the most major commitment of IBS may be put by saying that mentalese sentences are "the basic objects of interpretation; their content confers content upon thoughts; and thoughts transmit their content to outer speech" (McGinn 1982, p. 70). So IBS not only needs an account of the public-language relation, it also needs at least to show that a certain kind of account of the language-of-thought relation is possible. But while IBS took itself to have a strategy for defining the public-language relation, its having a strategy for defining the language-of-thought relation was left for a future occasion.

Let's return to the Lewis-inspired IBS strategy for explaining expression-meaning _{P} . The first part of that strategy required defining a language as a finitely specifiable function from finite sequences of sound, marks, whatever, onto things that might be their public-language meanings. In order to implement this part of his strategy, the IBS theorist must first do two things. So far I've characterized IBS as holding that speaker-meaning is a relation between speakers and propositions of some stripe or other. The first thing the IBS theorist must do is settle on the stripe of the propositions we mean; this is essential in order to know what word meanings should be taken to be. IBS theorists may disagree among themselves as to the nature of the propositions that comprise the range of the speaker-meaning relation, but for convenience of exposition I will initially suppose them to be Russellian, structured entities made up of the objects and properties our speech acts are about. I don't really think they can be Russellian propositions, for I don't think that what a speaker would mean in uttering 'J. K. Rowling is the author of the Harry Potter novels' is the same as what he would mean in uttering 'Robert Galbraith is the author of the Harry Potter novels', but once I reach the conclusions I need to reach relative to the Russellian assumption, it will be easy to show (in §IV) that those conclusions would also have been reached relative to any other assumption about the nature of the propositions we mean. The second thing the IBS theorist

⁷ Talk of a Mentalese sentence's being tokened in one's belief box is a metaphor for one's standing in that functional relation to the sentence that is constitutive of the sentence's being tokened as a belief (see my (1981a).

must do is say what expression-meanings must be relative to the assumption that the things we mean are Russellian propositions. Now IBS takes sentences to be devices for performing acts of speaker-meaning, and it takes their meanings to provide bases from which hearers may infer what speakers mean in producing unembedded utterances of those sentences. If declarative sentences were context *insensitive*, the IBS theorist could take their meanings to be the propositions that, all else being equal, speakers would mean in producing unembedded utterances of them, and it would take the meanings of words and other subsentential expressions to be the objects and properties that compose those propositions. But since hardly any of the sentences we utter are context insensitive, IBS will want to take expression-meanings to be things that *constrain* what speakers can mean in uttering those sentences. This may be spelled out in the following way.

Like virtually every other philosopher who has written about context-sensitive expressions, the IBS theorist recognizes two kinds of meanings, which I'll call *meanings* and *contents*. Meanings are properties of expression-types, contents of expression-tokens (\approx utterances of expressions⁸). The meaning of an expression-type ε determines, for any x and any token ε_τ of ε , what is required for x to be the content of ε_τ . Contents of declarative sentence tokens are (we're supposing) Russellian propositions, and if a proposition is the content of an unembedded token of a declarative sentence, then that proposition is what the speaker *said* in producing the token, if she said anything in producing it. The content of every expression token is its contribution to the content of the unembedded sentence token in which it occurs (an unembedded sentence token counts as occurring in itself). The meaning of a context-*insensitive* expression will assign the same content to every token of the expression which has a content, whereas the meaning of a context-sensitive expression will permit different tokens of the expression to have different contents. I'll call this way of understanding meanings and contents the *meaning/content model* (the M/C model, for short).

Given the M/C model it's useful to represent the meaning of an expression as a function from its tokens to their contents; to define such a function is to specify, for any possible content c and expression ε , the conditions that must be satisfied in order for c to be the content of a token

⁸ Sometimes, for stylistic reasons, I'll speak of the content of an utterance of an expression, rather than of a token of the expression.

of ε .⁹ So conceptualized, the meaning of a context-insensitive expression will be a constant function that maps every token of the expression onto the same content, and the meaning of a context-sensitive expression will be a nonconstant function whose values may differ from one token of the expression to the next. Since IBS takes a language to be a *finitely specifiable* function that maps its expressions onto their meanings, it is committed to taking the meaning of a complex expression to be a syntax-determined function of the meanings of its constituent words (or other morphemes); consequently, the only meanings that must be specified in defining a language are those given in the base-clauses for the language's morphemes. Specifying the meaning of a context-insensitive word merely requires saying what unique content each of its tokens must have. For example, the theorist might say that the meaning of 'Saul Kripke' is that function f such that,

⁹ On David Kaplan's (1989a, 1989b) version of the M/C model, meaning is called *character*, and is represented as a function from "contexts of utterance" to Russellian contents, where a "context of utterance" is a sequence of exactly the things needed to determine the content of *any* context-sensitive expression relative to such a "context." But (i) the representation of an expression's meaning as a mapping of the expression's tokens onto their contents is merely a notational variant of Kaplan's characters *if* there really are sequences of things capable of determining the contents of context-sensitive expressions, since in that case any function from tokens to contents would have to be defined in terms of the way Kaplan's sequences determine the contents expressions have relative to them; (ii) it can't be assumed that there are Kaplanian sequences to determine the content of any context-sensitive expression, so Kaplan's representation of expression meanings begs questions that aren't begged by the IBS construal of how the meaning/content model should represent meanings; and (iii) it's arguable that for all but a few indexical expressions—e.g. 'I' and 'yesterday' (but *not* 'here' and 'now')—most of the work in determining the reference of a speaker's utterance of a demonstrative or indexical is done by the speaker's referential intentions (e.g. it's *very* arguable that in a speaker's utterance of 'He doesn't love her', the referent of 'he' is the male to whom the speaker referred with 'he' and the referent of 'her' is the female to whom the speaker referred with 'her'), and this, in fact, is the view Kaplan came to accept in (1989b); but then there is no need for a technical conception of "contexts of utterance" as Kaplanian sequences if the references of context-sensitive terms are determined by speakers' referential intentions (see Schiffer 2003, §3.7). Kaplan does offer an argument (1989a: §XIII) to show that we won't be able to have a logic capable of accounting for the validity of arguments formulated with sentences containing demonstratives or indexicals if we represent those arguments as containing *utterances/tokens* of those expressions, but the soundness of Kaplan's argument is discussable. It's also worth remarking that the practice among formal semanticists of representing referential pronouns and demonstratives as free variables and then giving the truth conditions for the sentences containing them relative to "contextually-determined assignment functions" isn't incompatible with taking the meaning of an expression to be a function from its tokens to their contents.

for any token τ of ‘Saul Kripke’, $f(\tau) = \text{Saul Kripke}$. As for context-sensitive words it will prove useful to see specifications of the functions that are their meanings as completions of the schema

For any token ε_τ of ε and any x , $f(\varepsilon_\tau) = x$ iff

- (1) [*eligibility condition*] ...;
- (2) [*selection condition*] ...,

where the eligibility condition for an expression specifies the condition something must satisfy in order for it to be eligible to be the content of a token of the expression, and the selection condition selects which of the things eligible to be the content is the content. For example, one might specify the meaning of the demonstrative ‘she’ thus:

x is the content of a token τ of ‘she’ 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 kinds of 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 meaning 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. Perhaps no word 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.¹⁰

Having defined a language as a finitely specifiable mapping of sounds onto meanings, the IBS theorist sets out to complete his account of expression-meaning by defining the public-language relation, the relation that must obtain between a population P and a language L in order for L to be a public language of P . There were known to be problems that arose in attempting to define the public-language relation, perhaps the biggest being what David Lewis (1975) dubbed the *meaning-without-use problem*, and they were the problems Brian sought to solve in “Two Theories of Meaning.” Now, in order for a definition of the public-language relation to provide

¹⁰ 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.

an *IBS* account of expression-meaning its defining conditions must involve no ineliminable semantic notions, but only those propositional-attitude notions that, according to *IBS*, define speaker-meaning and convention, and elsewhere (2017) I have argued that that requirement can't be met, and that therefore there can be no correct *IBS* account of expression-meaning. Be that as it may however. I will now present a new argument to show why no *IBS* account of expression-meaning can be correct. This argument will have relevance well beyond *IBS*, however; for, if it's sound, it will show that no account of expression-meaning that presupposes the M/C model can be correct, and a corollary of this will be that natural languages don't have the kind of compositional semantics that nearly every philosopher of language and semanticist supposes they must have.

II. Penumbral Shift and Context Sensitivity

How might the M/C model accommodate vague terms when it's not ignoring their vagueness? Many vague expressions are obviously context sensitive, whether or not their vagueness is taken into account, and therefore can be accommodated in the model only if it can assign them meanings that are specifiable by instances of the schema:

For any token ε_τ of ε and any x , $f(\varepsilon_\tau) = x$ iff

(1) [eligibility condition] ...;

(2) [selection condition]

Such terms include, for example, 'here', 'fast', 'tall', and 'intelligent'. But other vague expressions will be treated by the model as context insensitive when their vagueness is ignored. Such terms include 'boy', 'bachelor', 'bald man', 'neurotic', 'logician', 'athlete', and 'triangular'. For example, when its vagueness is being ignored (and the things we mean are taken to be Russellian propositions), the *IBS* theorist would say that the meaning of 'boy' is a constant function that maps every token of the word onto the property of being a boy. Let's call such terms *predicates**. What can the model say about *predicates** when it's *not* ignoring their vagueness? Let's start by noticing that every *predicate** has a feature, consequent on its being vague, such that, owing to its having that feature, the *predicate** can be accommodated in the M/C model only if it can be assigned a context-sensitive meaning, i.e. a meaning that is specifiable as an instance of the just displayed schema. I call this feature *penumbral shift*, and it may be explained in the following way.

The notion of penumbral shift is defined in terms of the notion of a *penumbral profile*. Every token of every predicate* has a penumbral profile, where two predicate* tokens have the same penumbral profile just in case if either applies to a thing, then so does the other; if either is such that it's indeterminate whether it applies to a thing, then so is the other; if either is such that it's indeterminate whether it's indeterminate whether it applies to a thing, then so is the other; and so on.¹¹ We may think of a predicate* token's penumbral profile as determining the token's *penumbra*, where a vague predicate* token's penumbra is the range of things to which the token's application is anything other than unqualifiedly correct or unqualifiedly incorrect. Thus, a vague predicate* token's penumbra includes the token's borderline cases, borderline cases of its borderline cases, borderline cases of its borderline cases of its borderline cases, and so on.¹² *Penumbral shift* is then the fact that the penumbral profiles, and thus the penumbras, of a

¹¹ My use of 'vague', 'determinate', 'indeterminate', and 'borderline' throughout this essay is pretheoretic in that nothing I say will rely on any particular philosophical theory of vagueness or indeterminacy but will instead remain neutral on that score. So, for example, as far as this essay is concerned it's not assumed that 'It's indeterminate 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 indeterminate whether *x* is *F*'. (ii) 'It's indeterminate whether *x* is *F*' doesn't entail 'x is borderline *F*'. For example, if, as many suppose, it's indeterminate whether the continuum hypothesis is true, that wouldn't entail that it was borderline true or even vague. (iii) If it's indeterminate whether *x* is *F*, then it's impossible for anyone to know whether *x* is *F*. If, for example, it's indeterminate whether Harold is bald, then nothing can count as one's discovering that Harold is in fact bald or that he is in fact not bald. If one did come to know that Harold was bald, then one would thereby come to know that it wasn't indeterminate whether Harold was bald. (iv) One might wonder what the difference is supposed to be between *x*'s being *F* and *x*'s being determinately *F*. What, for example, is the difference between its being true that it's raining and its being determinately true that it's raining? If they are two distinct facts, in what can the difference between them consist? Are we to imagine that it's raining harder if we suppose that it's determinately true, as opposed to being merely true, that it's raining? Of course not. It's not that one might know that Sadie is lethargic and then wonder whether she's also determinately lethargic. To understand my vagueness-related use of the jargon expressions 'determinate' and 'indeterminate' it's enough to know that, while I assume that 'x is determinately *F*' entails 'x is *F*', I don't assume that 'x is *F*' entails 'x is determinately *F*'. For 'x is not determinately *F*' must be consistent with 'It's indeterminate whether *x* is *F*', and we don't want to assume that that entails 'x is not *F*'. So we can't deny that it's indeterminate whether *x* is *F* by claiming that *x* is *F* or that *x* isn't *F*, but we can deny it by claiming that *x* is determinately *F* or that *x* is determinately not *F*. The usefulness of 'x is determinately *F*/not *F*' resides in its incompatibility with 'It's indeterminate whether *x* is *F*'.

¹² Every token of a vague predicate has a penumbral profile and penumbra, and is also subject to a version of what I'm about to call *penumbral shift*. My focus on predicates* is for expository economy.

predicate*'s tokens may shift somewhat from one token to the next; that is to say, two tokens of a predicate* may have somewhat different penumbral profiles, and thus two different penumbras. The "somewhat" qualification is important. Typically, there will be numerous things of which every token of a predicate* must be true, and even more things of which every token of it must be false. 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 there to be three simultaneous tokens of any predicate*, one of which is true of the thing to which it's applied, another of which is false of the thing to which it's applied, and still another of which is such that it's indeterminate whether it's true of the thing to which it's applied. Here are three examples:

- ◆ At a party George is asked whether Henrietta came to the party with anyone, and he replies 'She came with some bald guy'. That utterance would most likely be accepted in the context in which it occurred as true, even if it transpires 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, someone 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 to say 'That's undetermined by the use of "bald" in everyday speech; such a man would be neither determinately bald nor determinately 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 among 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 determinately true when said of an unmarried eighteen-year-old male, and may even count as false.

- ◆ An utterance of ‘Mary is getting married to a boy from Boston’ would count as false if the male whom she is about to marry is fifty-two years old, but is apt not to count as false if he is twenty-seven years old. At the same time, if every one of the seven male professors in one’s department is over forty except Henry, who is twenty-seven, an utterance of ‘Six men and one boy are professors in my department’ would count as a misuse of ‘boy’.

Let’s take ‘boy’ as our exemplar of a predicate* and ask how a theorist might attempt to accommodate it in the M/C model when its vagueness is taken into account.

We may start by asking what the IBS theorist might say about the content of the token of ‘boy’ uttered by Jack in the following example.

Jack and Jill are on a first date, and at the moment are telling one another about their respective families. Jack has just told Jill that he has a sister who has a seven-year-old child named ‘Morgan’, and then, realizing that ‘Morgan’ is a unisex name, adds ‘Morgan is a boy’.

IBS holds that the things we mean, and thus the contents of declarative sentence tokens, are propositions, and we are provisionally assuming that IBS takes those propositions to be Russellian propositions. If an IBS theorist were ignoring vagueness, she would say that what Jack meant in uttering ‘Morgan is a boy’, and the content of the token of that sentence he uttered, is the singular proposition $\langle \text{Morgan, the property of being a boy} \rangle$; ¹³ for if the theorist is ignoring vagueness, then she would say that ‘boy’ is context insensitive, and that its meaning is that function f such that for any token τ of ‘boy’ and any x , $f(\tau) = x$ iff $x =$ the property of being a boy, and that therefore the content of the token of ‘boy’ Jack uttered is the same as every other token of ‘boy’, viz. the property of being a boy. But now we are supposing the theorist not to be ignoring vagueness, and ‘boy’ is vague. It’s vague because even if the token of ‘Morgan is a boy’ Jack uttered is determinately true or determinately false, it might have been neither: Morgan might have been a borderline case of a

¹³ For convenience of exposition I will ignore syntax-determined temporal reference when nothing turns on it.

thing to which ‘boy’ applies, either because Morgan is a borderline case of a male, or of a human being, or because he’s of an age such that it’s indeterminate whether ‘boy’, as uttered by Jack, applies to a human male of that age. So what can the theorist say about the content of the token of ‘boy’ Jack uttered? All that we can say at this point is that its content can’t be the property of being a boy, and this because there can be no such property as *the* property of being a boy. For if there were such a property, then it would have to be the content of every token of ‘boy’, and penumbral shift secures that no property has to be the content of every token of the word. Because of penumbral shift, the M/C model must treat ‘boy’ as a context-sensitive term; that is to say, if the model is to account for the word’s meaning in a way that takes account of its vagueness, then it can do so only by finding a true completion of the schema:

For any token τ of ‘boy’ and any property φ , $f(\tau) = \varphi$ iff

- (1) [eligibility condition] ...;
- (2) [selection condition]

What options does the M/C model have for achieving such a completion? In order to answer this, we would first have to know what options are available to the model for saying what the contents of tokens of ‘boy’ are, given the provisional assumption that the contents of declarative sentence tokens are Russellian propositions. We can make this question more tractable by returning to the example of Jack’s utterance of ‘Morgan is a boy’ (above p. ...) and asking what the M/C model can say about the content of the token of ‘boy’ Jack uttered. However, since properties, like propositions, are mind- and language-independent entities, the M/C theorist can’t say what the content of a token of ‘boy’ might be without first taking a stand on the issue of ontic vagueness, or vagueness in the world, at least as regards properties.

III. Ontic Vagueness and the Meaning of ‘Boy’

The issue of ontic vagueness is a contest between a view we might call *no-vagueness-in-the-world* and one we might call *vagueness-in-the-world*. What exactly is at issue in this contest doesn’t escape vagueness, but I think the following gloss should be good enough for present purposes.

No-vagueness-in-the-world holds that vagueness is exclusively a feature of language and thought and that, therefore, nothing outside of language and thought can be vague in its own right, independently of the vagueness of language and thought, while vagueness-in-the-world holds that vagueness isn’t exclusively a feature of language and thought and that, therefore, properties or

things outside of language and thought may be vague in their own right. I'll say that things that aren't vague in their own right are *metaphysically precise*, and that things that are vague in their own right are *metaphysically vague*. A Russellian proposition is metaphysically precise just in case each of its constituent properties or objects is metaphysically precise, metaphysically vague just in case one or more of its constituent properties or objects is metaphysically vague. Since our concern now is to characterize the kind of property that might function as the content of the token of 'boy' Jack uttered when he uttered 'Marlon is a boy', we need only to engage the issue of ontic vagueness as regards properties. So, what would it be for a property to be metaphysically precise or vague? What, to begin, would it be for a property to be metaphysically vague? Some philosophers doubt whether we can make coherent sense of vagueness-in-the-world, but in order to give the views against which I'll be arguing the best run for their money, I will assume that vagueness-in-the-world is a view to be reckoned with. Consequently, when I ask what it would be for a property to be metaphysically vague, I'm really asking what is the best answer to that question that a proponent of vagueness-in-the-world might give. How would such a theorist say what a metaphysically-vague property is?

One thing the theorist will want to say is that φ is a metaphysically-vague property *only if* it's possible for there to be a metaphysically-precise object o such that it's metaphysically indeterminate whether o has φ , where that means that there is no fact of the matter as to whether o has φ .¹⁴ But the theorist won't be able to say that φ is metaphysically vague *if* o is metaphysically precise and it's metaphysically indeterminate whether o has φ . This is because no-*vagueness*-in-the-world doesn't entail no-*indeterminacy*-in-the-world and we can't rule out that there might be a precise object x and a precise property φ such that, while x can't be a borderline instance of φ , it's nevertheless metaphysically indeterminate whether x has φ —for example, it may be metaphysically indeterminate that a given alpha particle will decay within a certain period, where that metaphysical indeterminacy can't be explained by attributing vagueness to any object or property. What must a vague property be like so that its being metaphysically indeterminate whether a thing has that property can be explained in terms of the property's vagueness? In explaining penumbral shift in §II, I said that every token of a predicate* has (what I called) a

¹⁴ We need the qualification about a metaphysically-precise *object* because if an object is metaphysically vague (whatever that might mean) it may be that it's entirely owing to its vagueness that it's a borderline case of a thing that instantiates a certain property.

penumbral profile which determines the token's *penumbra*, which is to say, the range of things to which the token's application is anything other than unqualifiedly correct or unqualifiedly incorrect. I believe that what the proponent of metaphysically-vague properties needs to say is that a vague property is a property that has a penumbral profile which determines a penumbra for the property, where that is the range of things to which the property's application is anything other than unqualifiedly correct or unqualifiedly incorrect. If φ is a metaphysically-vague property of a kind that enters into normal discourse, then we should expect there to be numerous things to which φ unqualifiedly applies, and even more things to which φ unqualifiedly doesn't apply. But *the property of being a thing to which φ unqualifiedly applies/doesn't apply must* evidently also be vague. This means that not only can there be no precise cutoff as regards the things which have φ , there can also be no precise cutoff as regards the things such that it's metaphysically determinate that they have φ . A metaphysically-vague property will not have a precise condition of application at any order of vagueness. By contrast, a metaphysically-precise property will have an absolute, unconditionally precise condition of application, even if it's possible for there to be things for which it's metaphysically indeterminate whether they satisfy that condition. Now, there are reasons to question whether the foregoing characterization of metaphysically-vague properties is coherent, but my concern hasn't been to show that there are, or even that it's metaphysically possible for there to be, metaphysically-vague properties; my concern has been merely to say what I think the proponent of vagueness-in-the-world needs to say about what it is for a property to be metaphysically vague.

In any case, we are now positioned to say that if the vagueness of 'boy' can be accommodated in the M/C model relative to the assumption that Russellian propositions are the contents of declarative sentence tokens, then the following options (in which β is the token of 'boy' Jack uttered and we're back to a pretheoretic use of 'determinate' and its cognates) are the ones worth taking seriously:

- (A) Some metaphysically-precise property is the content of β .
- (B) No metaphysically-precise property is such that it's determinately the content of β , but myriad metaphysically-precise properties are each such that it's indeterminate whether it's the content of β .
- (C) Some metaphysically-vague property is the content of β .

- (D) No metaphysically-vague property is such that it's determinately the content of β , but myriad metaphysically-vague properties are each such that it's indeterminate whether it's the content of β .

Although (A)-(D) are the only answers we need to 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 indeterminate whether X , so that, for example, in addition to (A), logical space also contains the answers:

- (A2) It's indeterminate whether (A).
 (A3) It's indeterminate whether (A2).
 (A4) It's indeterminate whether (A3)
 ⋮

But once our critical examination of (A)-(D) is concluded it will be clear why we have no need to consider permutations of them induced by higher-order vagueness.¹⁵

Now that we know the options available to the M/C model as regards the content of the token of 'boy' Jack uttered, we can discuss, relative to each of those options, the prospect of achieving a correct completion of the schema:

- For any token τ of 'boy' and any property φ , $f(\tau) = \varphi$ iff
- (1) [eligibility condition] ...;
 - (2) [selection condition]

Re (A) [Some metaphysically-precise property is the content of β (the token of 'boy' Jack uttered)]. The token of 'Morgan is a boy' Jack uttered is vague. It's vague because even if it's determinately true or determinately false, it might have been neither; it might have been borderline true/false. The vagueness of the sentence token is entailed by the vagueness of β , the token of 'boy' contained in it, and, let's suppose, that is the only source of the sentence's vagueness. Given that and the assumption that Russellian propositions are the contents of declarative sentence tokens, the content of the sentence token Jack uttered is the metaphysically-precise proposition $\langle \text{Morgan}, \varphi \rangle$, if (A) is correct and φ is the metaphysically-precise property that is the content of

¹⁵ Another option in logical space not worth taking seriously is that 'boy' has a constant meaning that changes from moment to moment.

β . But how can the token of ‘Morgan is a boy’ be vague if its content is a metaphysically-precise proposition? That is possible if, and, I should think, only if, the epistemic theory of vagueness is correct. For that theory holds that a borderline sentence is true or else false, and we cannot know which. The fact that (A) presupposes the epistemic theory is already enough to make (A) problematic (see e.g. Schiffer 2003, pp.183-4), but we can put that polemical remark aside, because, as we’re about to see, we can show that (A) is false without having independently to show that the epistemic theory is false.

Penumbral shift entails that there is no such property as *the* property of being a boy. If there were such a property, it would be the content of every token of ‘boy’, and that is precisely what penumbral shift precludes. Given penumbral shift, if (A) is true, then different tokens of ‘boy’ may have different metaphysically-precise properties as their contents. In fact, the content of a particular token of ‘boy’ would have to be one of uncountably many metaphysically-precise properties that are equally eligible to be the content of a token of ‘boy’. Let’s ignore the fact that ‘male’ and ‘human being’ are vague and pretend that they express metaphysically-precise properties. Let’s also pretend that for each person there is a precise instant of time at which he or she came into existence. Then a good candidate for being the content of β if (A) is correct would be, say, *the property of being a male human being whose age in seconds $\leq 531,066,240$* . But if that is a good candidate, then so is *the property of being a male human being whose age in seconds $\leq 531,066,240.1$* , *the property of being a male human being whose age in seconds $\leq 531,066,240.01$* , *the property of being a male human being whose age in seconds $\leq 531,066,240.001$* , *the property of being a male human being whose age in seconds $\leq 531,066,239$* , *the property of being a male human being whose age in seconds $\leq 531,066,239.05733$* , *the property of being a male human being whose age in seconds $\leq 536,000,674$* , *the property of being a male human being whose age in seconds $\leq 527,000,674.7009$* , ...—well, you get the point. Let’s suppose (for *reductio*) that (A) is true and that among the uncountably many equally good candidates for being the content of β the one that happens to be its content is *the property of being a male human being whose age in seconds $\leq 530,802,239.057$* . Call this property φ^* . It can’t be a brute inexplicable fact that φ^* is β ’s content; there must be some feature that φ^* alone among the countlessly many equally good candidates has which makes *it* the content of β . What might

that feature be? It can't be that Jack *intended* to express φ^* with β .¹⁶ It can't be any *causal relation* that relates β to φ^* but to no other equally good candidate. And it's preposterous to suppose it's somehow miraculously determined by the use of 'boy' among speakers of English. I submit that there is *nothing* that could make one alone of these countless many candidate properties the content of β , and that, therefore, (A) is false.

There is another, independent, reason why no metaphysically-precise property can be the content of β . If (A) were true, then, owing to penumbral shift, the meaning of 'boy' would have to be a function that was specifiable as an instance of the schema:

For any token τ of 'boy' and any metaphysically-precise property $\varphi, f(\tau)$
 = φ iff
 (1) [eligibility condition] ...;
 (2) [selection condition]

The considerations rehearsed in the preceding paragraph show that there can be no correct selection condition. But we can also see that there can be no correct eligibility condition. For let's continue to pretend that 'male' and 'human being' express metaphysically-precise properties and that for each person there is a precise instant of time at which he or she came into existence. Then in order for there to be a correct eligibility condition for 'boy' there would have to be real numbers j, k such that (a) j is the largest number such that, necessarily, *every* token of 'boy' is determinately true of every human male whose age in seconds $\leq j$ and (b) k is the smallest number such that, necessarily, *every* token of 'boy' is determinately false of every human male whose age in seconds $\geq k$. If there were such numbers, then we could specify the eligibility condition by saying that:

For any token τ of 'boy' and any metaphysically-precise property $\varphi, f(\tau)$
 = φ *only if*, for some real number $j' \geq j, < k$, a male human being has φ
 iff his age in seconds $\leq j'$.

But there being such a j' is precluded by the fact that there can be no determinate cutoffs at any order of vagueness: owing to high-order vagueness, there can be no greatest lower bound or least upper bound to determine the metaphysically-precise properties that are eligible to be the content of a token of 'boy'.

¹⁶ For Jack to have such an intention he would have to have some way of thinking of the property under which he had that intention, but he clearly has no such way of thinking of it. See the discussion of *de re* intentions in Schiffer (2019).

(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 prescind 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. The most prominent such attempt is the semantics for vague gradable adjectives Christopher Kennedy (2007) offers. On that semantics, the meaning of 'tall' involves a contextually-determined measure function which determines a scale of measurement for height and a precise degree of height on that scale so that an utterance of 'John is tall' in context c is true if John's height surpasses the degree-of-height threshold determined by c , 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 surpass 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, even though the context in which 'John is tall' 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, Chris Barker (2002) 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's again, it offers no account of how that threshold is determined.¹⁷)

Re (B) [No metaphysically-precise property is such that it's determinately the content of β , but myriad metaphysically-precise properties are each such that it's indeterminate whether it's the content of β .] Let D be the determinately operator, so that DP means that *it's determinately the case that P* . As already remarked (see n. 11), we don't want P to entail DP because we don't want $\neg DP$ to entail $\neg P$. Nevertheless, one can't be *warranted* in *asserting P* unless one would be equally warranted in asserting that it was determinately the case that P —whence the infelicity of 'Morgan is a boy but it's indeterminate whether or not he's a boy'. Similarly, if one has the concept of a proposition's being determinately true and one knows that one knows that p is true, then one knows that p is determinately true, and if one is justified in believing that one is justified in believing that p is true, then one is justified in believing that p is determinately true. In the just concluded discussion of (A) I offered two arguments to show that every metaphysically-precise property is such that it couldn't be the content of β . I could not coherently have reached that conclusion and then added "But while the foregoing argument shows that every metaphysically-precise property is such that it couldn't be the content of β , it's still an open question whether or not there is a metaphysically-precise property such that it's indeterminate whether it's the content of β ". This is because to whatever extent those arguments justify me in concluding that no metaphysically-precise property is the content of β , to that same extent they justify me in concluding that it's determinately the case that no metaphysically-precise property is the content of β . Now, if it's determinately the case that no metaphysically-precise property is the content of β , then every metaphysically-precise property is such that it's determinately not the content of β ; and if that is so, then no precise property is such that it's indeterminate whether it's the content of β . In short, the considerations that show (A) to be false also show (B) to be false.

¹⁷ Una Stojnić and Matthew Stone (manuscript) 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.

It's also important to notice that the very same eligibility-condition problem that arises for (A) also arises for (B). This is because in order for a property to be such that it's indeterminate whether it's the content of a token of β , 'boy' must have a meaning that determines what must be the case in order for a property to be the content of a token of 'boy', and such a meaning, as we've seen, requires a correct eligibility condition for 'boy'. So the impossibility of such a condition falsifies (B) as well as (A).

Re (C) [Some metaphysically-vague property is the content of β]. If there are metaphysically-vague properties they are every bit as finely individuated as metaphysically-precise properties. So if a metaphysically-vague property φ is a good candidate for being the content of β , then so are countlessly many metaphysically-vague properties that are virtually indistinguishable from φ . For example, suppose that metaphysically-vague properties φ and ψ are identical except that, whereas it is indeterminate whether φ applies to a human male whose age in seconds is 522,000,000, it is merely indeterminate whether it is indeterminate whether ψ applies to him. Then we should expect that if either property is a good candidate for being the content of β , then so is the other. (A) was shown to be false by the fact that there was nothing to explain what would make any precise property such that it alone among the uncountably many equally eligible precise properties was the content of β . In the same way, (C) is shown to be false by the fact that there is nothing to explain what would make any metaphysically-vague property such that it alone among the uncountably many equally eligible metaphysically-vague properties was the content of β .

Also like the falsity of (A), the falsity of (C) is overdetermined. If (C) were correct, there would be a correct completion of the M/C schema

For any token τ of 'boy' and any metaphysically-vague property $\varphi, f(\tau)$
 = φ iff
 (1) [eligibility condition] ...;
 (2) [selection condition] ...,

but the argument of the preceding paragraph shows that there is no correct selection condition for this schema, just as the corresponding argument for metaphysically-precise properties showed that there could be no correct selection conditions for those properties. The discussion of (A) also contained an argument to show that there could be no correct eligibility condition for

metaphysically-precise properties, and the same argument is available, *mutatis mutandis*, to show that there also can't be a correct eligibility condition for metaphysically-vague properties. If a metaphysically-vague property is the content of β and Morgan has that property, then Jack's utterance of 'Morgan is a boy' is true, and if Morgan doesn't have that property, then the utterance is false, and if Morgan is a borderline instance of the property, then it's indeterminate whether the utterance is true. What matters as regards an eligibility condition are the conditions a property imposes on its determinate application. So, if we continue to pretend that 'male' and 'human being' express metaphysically-precise properties and that for each person there is a precise instant of time at which he or she came into existence, then, as before, in order for there to be a correct eligibility condition for the just-displayed schema there would have to be real numbers j, k such that (a) j is the largest number such that, necessarily, every token of 'boy' is determinately true of every human male whose age in seconds $\leq j$ and (b) k is the smallest number such that, necessarily, every token of 'boy' is determinately false of every human male whose age in seconds $\geq k$. And if there were such numbers, then we could specify the eligibility condition by saying that:

For any token τ of 'boy' and any metaphysically-vague property $\varphi, f(\tau)$
 = φ only if, for some real number $j' \geq j, < k$, a human male has φ iff his
 age in seconds $\leq j'$.

But there being such a j' is precluded by the fact that there can be no determinate cutoffs at any order of vagueness: owing to high-order vagueness there can be no greatest lower bound or least upper bound to determine even the metaphysically-vague properties that are eligible to be the content of a token of 'boy'. Therefore, every metaphysically-vague property, like every metaphysically-precise property, is such that it's not the content of β .

Re (D) [No metaphysically-vague property is such that it's determinately the content of β , but myriad metaphysically-vague properties are each such that it's indeterminate whether it's the content of β]. (D) stands to (C) as (B) stands to (A). I claim that the discussion of (C) warrants concluding that (C) is false, and to whatever extent it warrants holding that no metaphysically-vague property is the content of β , to that same extent it warrants holding that it's determinately the case that no metaphysically-vague property is the content of β , and that therefore each metaphysically-vague property is such that it's determinately not the content of β , and therefore

no metaphysically-vague property is such that it's indeterminate whether it's the content of β . Just as the case against (A) is simultaneously a case against (B), so the case against (C) is simultaneously a case against (D).

IV. The Argument Extended

The options (A)-(D) don't presuppose IBS. They are, I submit, the options available to *any* theorist who thinks the M/C model can accommodate vagueness and who takes Russellian propositions to be the contents of tokens of vague sentences (or to sentences relative to "contexts of utterance"). Would (A)-(D) fair better on another conception of propositions? When I provisionally supposed we were addressing a view on which the contents of sentence tokens were taken to be Russellian propositions I said that that was merely an expository convenience, since the arguments to be deployed to show that Russellian propositions couldn't be the contents of tokens of vague sentences could be redeployed to show that no propositions of any other kind could be their contents either. The idea is this. The arguments against (A)-(D) are effectually arguments against any theory that takes contents of tokens of vague predicates like 'boy' to be metaphysically-precise or metaphysically-vague properties. A theorist who takes the contents of sentence tokens to be propositions construed as functions from possible worlds to truth-values will construe properties as functions from possible worlds to sets of things in those worlds (e.g. she would represent the property of being a tiger as the smallest function that maps every possible world onto the set of things in that world that are tigers). It's not clear to me how the possible-worlds theorist might represent the notion of a metaphysically-vague property, but if we continue to pretend that 'male' and 'human being' express metaphysically-precise properties and that for each person there was a precise instant of time at which he or she came into existence, then the theorist would say that an example of a function eligible to be the content of β (the token of 'boy' Jack uttered) might be, say, that function that maps each possible world onto the set of human males in that world whose age in seconds $\leq 530,156,794$. But if such a function is eligible to be the content of β , then so is the function that maps each possible world onto the set of human males in that world whose age in seconds $\leq 530,156,795$, the function that maps each possible world onto the set of human males in that world whose age in seconds $\leq 530,156,794.0057$, and so on for countless many other functions that differ only in some irrelevant way from them. But then the sort of considerations which show that there is nothing that could make one alone of the countless many equally eligible

properties the content of β would also show that there was nothing that could make one alone of the countless many equally eligible functions the content of β , and it should be obvious that the considerations pertaining to higher-order vagueness that make it impossible for there to be a correct eligibility condition for properties would make it equally impossible for there to be one for the functions that are the possible-world theorist's surrogates for properties. If anything is the content of β —or if anything is such that it's indeterminate (or indeterminate whether it's indeterminate, or ...) whether it's the content of β —and it's not a property or property-surrogate, then it would have to be something that *determines* a property in the way that a Fregean “mode of presentation”/“way of thinking” of a property is supposed to do. But then there would be countless many equally eligible such things, and it should by now be obvious that they, too, would fail for the same sort of reasons that the other candidate contents fail.

From this I conclude that, if there are such *things* as the contents of tokens of vague sentences, then they aren't propositions of any kind. At the same time, it would surely have been correct to report that, in uttering ‘Morgan is a boy’, Jack told Jill, and thus meant, that his sister's child Morgan was a boy, and from this we may further conclude that speaker-meaning isn't a relation, *S meant p*, between a speaker *S* and a proposition *p* that she meant;¹⁸ for if it were, then propositions *would* be the contents of sentence tokens. And, I should think, we may also conclude that there are no such *things* as the contents of vague sentence tokens or the things we mean; for if there were such things they would evidently have to be language- and mind-independent things that have truth conditions, and that just is the generic sense of ‘proposition’ (see above p. 000).

Davidsonians, however, claim that semantics has no need for such *things* as “contents,” and that the only semantics a public language needs is an extensional compositional truth theory, where, roughly speaking, an extensional compositional truth theory for a language *L* is a finitely storable theory of *L* that generates for each truth-evaluable sentence of *L* a theorem of the form ‘An utterance of σ is true iff ...’, where ‘iff’ is the material biconditional. Could such a semantics accommodate the vagueness of ‘boy’? Some theorists would say the answer is yes, and they would support that answer in the following way:

The words ‘snow’ and ‘white’ are vague, and therefore so is the sentence
‘Snow is white’. Yet that doesn't affect the truth of the “T-sentence”

¹⁸ There is a freestanding argument for this conclusion in Schiffer (2019).

(T) ‘Snow is white’ is true iff snow is white,
and this because the vagueness of (T)’s right-hand side cancels out the
vagueness of its left-hand side, in that each side matches the other “umbra
for umbra and penumbra for penumbra” (Quine 1960, p. 37; quoted in
Lepore and Ludwig 2005, p. 139).

But this is wrong; it’s falsified by the context sensitivity induced by penumbral shift. Not only is
(T) not true, it’s also the case that every assertive utterance of it is false. For suppose the displayed
token of (T) was assertively uttered. The left-hand side of that token ascribes truth to the sentence
type ‘snow is white’. If a sentence type is true, then every assertively uttered token of it is true,
and whatever is required for the truth of the sentence type is also required for the truth of every
assertively uttered token of it. In this way, we could have more precisely expressed what is
supposed to be expressed by (T) by:

(T’) For every (assertively uttered) token τ of ‘snow is white’, τ is
true iff snow is white.

The problem with (T’) is that *every* assertive utterance of *it*—i.e. of (T’) itself—is *false*. For
suppose the displayed token of (T’) was assertively uttered. Let’s call that token ‘(T’)’ $_{\tau}$. Then
every assertively uttered token of ‘snow is white’ would have exactly the same truth and falsity
conditions as the token of ‘snow is white’ on the right-hand side of (T’)’ $_{\tau}$. But since ‘snow is white’
is vague, it’s subject to penumbral shift, and a consequence of that is that there may be assertively
uttered tokens of ‘snow is white’ with different truth-values, in which case the penumbral profile
of the token of ‘snow is white’ on the right-hand side of (T’)’ $_{\tau}$ wouldn’t determine the truth
conditions of every assertively uttered token of ‘snow is white’, and from that it follows that (T’)’ $_{\tau}$
is false, and the same goes for every assertively uttered token of (T’). In effect, an assertive
utterance of (T’) would be in the same sort of pickle as an assertive utterance of

For every (assertively uttered) token τ of ‘she is French’, τ is true iff she
is French

in which the token of ‘she’ on the right-hand side referred to a particular female Sabine. Just as
what would be needed to accommodate ‘she’ in a compositional truth theory would be a finitely
statable meaning for ‘she’ that specified what must be the case in order for a token of ‘she’ to refer
to a particular thing, so what would be needed to accommodate the vague sentence ‘snow is white’
in a compositional truth theory would be a finitely statable meaning for the sentence that specified

what must be the case in order for any token of ‘snow is white’ to be true; but if we could achieve such a specification, then we could show how the M/C model is able to accommodate vague terms like ‘snow’, ‘white’, and ‘boy’. As regards ‘boy’, for example, what kind of base axiom should we expect a compositional truth-theoretic semantics to have for the word in order for it to capture the word’s penumbral-shift-induced context sensitivity? Well, let’s continue to pretend that ‘male’ and ‘human being’ have precise extensions and that for each person there was a precise instant of time at which he or she came into existence. And let’s also pretend that there are real numbers j, k such that j is the largest number such that every token of ‘boy’ is determinately true of a male human being whose age in seconds $\leq j$ and k is the smallest number such that every token of ‘boy’ is determinately false of every male human being whose age in seconds $\geq k$. Then we could devise a base axiom for ‘boy’ in a compositional truth theory *provided* we could find a condition C such that:

For any x , token τ of ‘boy’, and time t , τ is true of x at t iff at t x is a male human being whose age in seconds = some real number n such that:

- (1) (eligibility condition) $j \geq n < k$;
- (2) (selection condition) n alone of the countlessly many eligible numbers—i.e. the uncountably many between j and k —satisfies condition C .

But of course we already know that no such base axiom is to be found, for there can be no upper or lower bounds to serve as witnesses for ‘ j ’ and ‘ k ’ in the displayed conditions, and even if there were there would be no condition to do the job required for there to be a correct selection condition.

V. What Now?

A *compositional meaning theory for a language L* would be a finitely storable theory of L that generates for each expression ε of L a theorem that assigns to ε its meaning in L , where, for any possible content c and any token ε_τ of ε , that meaning determines what must be the case for c to be the content of ε_τ . A *compositional truth theory for a language L* would be a finitely storable theory of L that generates for each truth-evaluable sentence of L a theorem that specifies what must be the case for a token of that sentence to be true. If everything I’ve so far said in this paper is correct (a big if, to be sure), then neither kind of compositional semantics can accommodate the vagueness of vague expressions, and since virtually every sentence of every natural language is vague, this means that no natural language can have either kind of semantics. Might natural languages have some other kind of semantics? In order to address this question we should have on board two further, but so far unmentioned, features of

vagueness—namely, what I will call *penumbral ignorance* and *penumbral indifference*.

Penumbra *ignorance* is the fact that speakers tend to have very little, if any, awareness of the penumbral profiles of the vague expression tokens they utter. Many adults are probably unaware that most of the words they use are vague (in the sense that they admit of possible borderline cases), and it's *possible* for someone to count as a competent speaker of a language even if he doesn't even have the concept of a vague expression. My guess is that the concept of a borderline case isn't one a child acquires until fairly late in her linguistic development, well after she is regarded as a competent speaker.¹⁹ That is the most extreme form of penumbral ignorance. But fully competent speakers of a language also manifest penumbral ignorance. One typical manifestation is illustrated in the following example. Jane, an intelligent native speaker of English, is asked if she has any children, and she replies, 'Yes. I have a 7-year-old boy'. Although different tokens of 'boy' may have different penumbral profiles, Jane's child is such that every token of 'boy' must be true of him. Now suppose that, after explaining to her what the question means, we ask Jane, "What must be true of an individual in order for it to be a borderline case of a thing to which 'boy', as you just now used it, applies?" That is not a question Jane considered in making her utterance, nor did she have any reason to consider it: she knew that what she said was most definitely true, so she had no reason to wonder what sort of state of affairs would have made her utterance merely borderline true. What Jane *ought* to say in response to the question is that, never having considered what would count as a borderline boy, she is in no position to answer it. If Jane did venture an answer, it would have no special authority, either about the truth of her answer or about what she would say in relevant counterfactual circumstances.²⁰ I hazard, first, that most utterances of vague sentences are like Jane's in that, all else being equal, when a rational speaker applies a vague predicate to a thing, she is confident that the predicate is definitely true of the thing, and, second, that if in such cases the hearer doesn't believe what the speaker said it won't be because the predicate she uttered is vague. These speakers have no reason to consider what something would

¹⁹ Susan Carey (p.c.) believes that no one has studied this question.

²⁰ 'Midtown' is the name of a vaguely defined area of Manhattan: if you're in Times Square, you're definitely in Midtown and if from there you move in a straight line far enough in any direction you will eventually definitely not be in Midtown, but at no point will you have crossed an invisible line on one side of which you're definitely in Midtown and on the other side of which you're definitely not in Midtown. Nevertheless, I've found that when I ask savvy Manhattanites what the boundaries of Manhattan are, they don't hesitate to tell me—although no two of them tell me the same thing. A lawyer once told me the southern boundary was 34th St., but I got her to retract her statement when I asked, "You mean Madison Square Garden isn't in Midtown?"

have to be like in order for it to be a borderline case of a thing to which the token of the predicate she uttered applies. Penumbral ignorance is a consequence of this. But even when a speaker applies a vague predicate to a thing and her interlocutors deem that application to be merely borderline correct, they won't have considered what something would have to be like in order for it to be a borderline case of a borderline case of a thing to which the predicate, as uttered by the speaker, applies. Penumbral ignorance is a consequence of this too.

Penumbral indifference is the fact that we don't regard it as a barrier to knowing what someone said in uttering a vague sentence that we know very little about the utterance's penumbral profile and that, consequently, typically when a person x says something using a vague expression ε and another person y uses ε in reporting what x said, y pays no attention to whether the penumbral profile of her utterance of ε matches that of x 's utterance of ε . For example, if A assertively utters 'Jane's child is precocious' and B , reporting on the speech act A thereby performed, utters ' A told me that Jane's child is precocious', B 's utterance may count as true even though the penumbral profile of her utterance of 'precocious' is different from that of A 's. Related to this is the fact that when, as we would say, we "believe what so-and-so said," it's unlikely that the belief we acquire will have exactly the same penumbral profile as the one to which the speaker gave voice in her utterance.²¹ Penumbral indifference has limits, of course. For example, if I know that Leona believes that 'home' can't be true of apartments, but only of houses, then I would be very reluctant to report her utterance of 'Morris is homeless' by uttering 'Leona said that Morris was homeless'. If there was any information that I thought was transmitted by her utterance it would be information I would express with the sentence 'Morris doesn't live in a house'. But what are these limits? If a speaker assertively utters a sentence in which a vague predicate φ is predicated of a certain thing, when will one find it unproblematic to use φ in reporting what she said or the information acquired from her utterance? I think two already noticed facts about our use of vague terms will help to answer that question. The first is that however much the penumbral profiles of two tokens of a vague predicate may differ, there will normally be myriad things to which every token of the predicate must apply and even more things to which no token of it can apply. Let's call things of the first kind the vague predicate's *positive anchors* and things of the second kind its *negative anchors* (penumbral shift occurs in the area between the positive and negative anchors), while

²¹ If we think in a *lingua mentis*, the penumbral profile of a belief is the penumbral profile of the mentalese sentence whose tokening in one's belief box realizes the belief.

recognizing that there can be no determinate set of positive or negative anchors. The second fact about our use of vague terms is that, all else being equal, a rational speaker won't apply a vague predicate to a thing unless she is confident that the predicate is definitely true of the thing and would be recognized as such by others apprised of the facts on which her assertion was grounded, and she expects that if her hearer doesn't accept her assertion, it won't be because of its vagueness. I think this explains why, when a speaker predicates a vague predicate of a thing, we are normally comfortable in using that same predicate to report either what the speaker said or the knowledge we acquired from her utterance.

The point about positive and negative anchors might seem to suggest that, even if vagueness precludes our language from having a correct compositional truth theory that, for *every* truth-evaluable sentence σ , issues in a theorem of the form

For any token σ_τ of σ , σ_τ is true iff ...,

the fact that most single-word vague predicates have positive and negative anchors would make it possible for our language to enjoy a correct compositional truth theory that, for *most* truth-evaluable sentences σ that will ever be uttered, issues in theorems of the forms

For any token σ_τ of σ , σ_τ is true if ...

For any token σ_τ of σ , σ_τ is false if ...

The possibility of this sort of compositional truth theory is suggested by the following line of reasoning. What stands in the way of the more ambitious compositional semantics is the impossibility of getting correct base axioms for single-word vague predicates. For example, such a base axiom for 'boy' would have to be a completion of the form

For any x , token τ of 'boy', and time t , τ is true of x at t iff ...,

and if what I have argued is correct, no such completion is available. At the same time, 'boy' has positive and negative anchors; that is to say, there are numerous individuals—e.g. your 9-year-old son Charlie—of whom every token of 'boy' must be unqualifiedly true and numerous individuals—e.g. Ruth Bader Ginsberg—of whom every token of 'boy' must be unqualifiedly false. This evidently entails that there are numerous conditions Ω such that:

(1) For any x , token τ of 'boy', and time t , τ is true of x at t if x satisfies Ω at t

and numerous conditions Ω' such that:

(2) For any x , token τ of 'boy', and time t , τ is false of x at t if x satisfies Ω' at t .

Further, given that there are such conditions, then it's apt to seem that there should be

specifications of them, such as, for example:

- (3) For any x , token τ of ‘boy’, and time t , τ is true of x at t if at t x is a 9-year-old male human being.
- (4) For any x , token τ of ‘boy’, and time t , τ is false of x at t if at t x is an 86-year-old female human being.

And if there are such specifications of sufficient conditions for the application or non-application of ‘boy’, then there would be analogous sufficient conditions for nearly every other single-word vague predicate, and there would evidently be nothing standing in the way of achieving a limited compositional truth theory of the kind described. The hope would then be that this limited compositional semantics could be put to work to explain some of the things compositional truth-theoretic semantics is thought to be needed to explain.

That line of reasoning, and the hope it aims to sustain, is problematic in at least two ways. One problem is that even if most single-word vague predicates have positive and negative anchors that enabled us to fashion base axioms for those predicates of the kind in question, the limited compositional semantics which that made possible wouldn’t go far enough in explaining what makes utterances of vague sentences true; for, as we know from the discussion of penumbral shift, utterances of vague predicates can be true of things, and known to be true of them, even though the things of which they are true are not among the predicate’s positive anchors (for example, owing to penumbral shift, there could be two simultaneous predications of ‘bald’ to a certain man, one of which was true of him while the other which was false of him). A second problem is more serious. It may be illustrated in terms of our running example of the vague word ‘boy’. The problem comes to light when we notice that, while it is plausible that there are numerous conditions Ω such that:

- (1) For any x , token τ of ‘boy’, and time t , τ is true of x at t if x satisfies Ω at t

the sentence

- (3) For any x , token τ of ‘boy’, and time t , τ is true of x at t if at t x is a 9-year-old male human being.

is not the true substitution instance of the open sentence (1) that it appears to be, and when we see why this is so we’ll see why no sentence would provide a substitution instance. In order for (3) to be a substitution instance of (1) that is true, ‘is a 9-year-old human male’ in (3) must express a

condition that satisfies the open sentence (1). But owing to its penumbral-shift-induced context sensitivity, the expression *type* ‘is a 9-year-old human male’ in (3) doesn’t express any condition. What about *tokens* of that expression? Although different tokens of it may express somewhat different conditions, won’t it nevertheless be the case that every token of ‘boy’ must be true of every token of every individual of whom *any* token of ‘is a 9-year-old human male’ is true? No. There *could* be such predicates; for example, I should think that every token of ‘boy’ must be false of every individual of whom any token of ‘is an 86-year-old female human being’ is true. But it’s not the case that every token of ‘boy’ must be true of every individual of whom any token of ‘is a 9-year-old human male’ is true. One reason this is so is that ‘male’ is vague and there may be two tokens of the term one of which is true of a certain individual while the other is false of that individual, and in that case there might easily be a token of ‘boy’ that is false of someone of whom a particular token of ‘is a 9-year-old human male’ is true. What all this suggests is that, while there are indeed numerous conditions Ω such that, for any person x , token τ of ‘boy’, and time t , τ is true of x at t if x satisfies Ω at t , it may not be possible to articulate any of these conditions! One way this could happen would be if no-vagueness-in-the-world were true. For then we should expect there to be metaphysically-precise properties whose possession by a person was sufficient for every literal utterance of ‘boy’ to be true of him. Then, even though we couldn’t know what any such metaphysically-precise property was, we might still be able to know when someone had such a property. After all, you already know that every literally uttered token of ‘boy’ is true of your 9-year-old son Charlie, even though you might not be able to specify any property of Charlie that makes that the case. In any event, if sufficient conditions for the application of a vague predicate token are inarticulable, then it won’t be possible for anyone to *articulate* even a limited compositional theory with a base axiom for ‘boy’ that stated a sufficient condition for any token of ‘boy’ being true of every individual who satisfied that condition. This is an important point because many theorists hold that if a compositional theory is to explain language understanding, then a representation of it must be implicated in the information processing language understanding requires; but it won’t be possible for there to be a *representation* of such a theory if the conditions the theory needs to assign to vague predicates in the theory’s base axioms are inarticulable.

Where does all this leave us? From my vantage point, the next question to ask is: If what I have said in this paper is correct, what facts about a language that seemed to implicate the need

for the language to have a compositional semantics remain to be explained, and what kind of theories are needed to explain them? This is not an easy question to answer. For example, one thing that remains to be explained is the fact that we are able to understand indefinitely many sentences of our language upon encountering them for the first time. What must be true of the sentences and of us for this to be so? What makes this a difficult question to answer is that in order to make progress on it we first need to know in what “understanding” a sentence consists. When vagueness is being ignored a ready answer is that to understand a sentence is to know the kind of speech act that would be performed in a literal and unembedded utterance of the sentence, and the kind of content that speech act must have. That may seem to be a reasonable thing to say, at least as a starting point, even when vagueness isn’t being ignored, only now we have to tackle the question of what talk of the content of an utterance can come to when there are no such *things* as “contents” (speaker-meaning can no longer be thought of a relation between a speaker and a proposition she means), especially given penumbral ignorance and indifference.

In any case, the question of understanding novel sentences may not be the most pressing question to be tackled. Suppose we think—i.e. process information—in a language-like neural system of mental representation, Mentalese. Then our spoken language will have a compositional semantics if Mentalese does, but we won’t need a compositional semantics to explain our “understanding” of Mentalese because we don’t need to know anything about the “brain’s language of synaptic interconnections and neural spikes” (Lewis 1983, p. 346) in order to think in it. So why might one think there is something about our thinking in Mentalese whose explanation requires it to have a compositional semantics? If there is nothing about Mentalese whose explanation requires the neural language to have a compositional semantics, then we won’t need a compositional semantics for our spoken language to explain our ability to understand its sentences, since sentence understanding may be understood as a wholly nonsemantic process whereby spoken sentences are mapped onto their Mentalese correlates.²² If Mentalese needs a compositional semantics, it will be to explain how our beliefs are realized by tokenings of Mentalese sentences. Fodor argues that Mentalese will need a compositional semantics precisely to explain that. But the way he formulates the argument suffers from his ignoring vagueness:

²² For the full story on this, see my (1987, Ch. 7, “Compositional Semantics and Language Understanding) and Jerry Fodor’s defense of my argument in his (1990a).

Mentalese must ... be productive. The reasoning is ... familiar: On the one hand, propositional attitudes derive their semantical properties (their 'intentional contents') from the semantics of the Mentalese expressions And, on the other hand, there are infinitely many psychologically possible propositional attitudes What needs to be explained is that (synonymy aside) each of the infinitely many syntactically distinct expressions of [Mentalese] has its distinctive truth conditions (1990a, p. 421).

In other words, if Harold has a belief that is true iff snow is white, then that is because a Mentalese sentence that is true iff snow is white is tokened in his belief box. Since there are infinitely beliefs Harold might have, each with its distinctive truth condition, there must also be infinitely many Mentalese sentences, each with its distinctive truth condition. So let's suppose Harold thinks in English. Then his Mentalese will need a compositional semantics that generates for the Mentalese sentence 'Snow is white' the T-sentence

'Snow is white' is true iff snow is white,

and similarly for each of the infinitely many other sentences of his Mentalese. But, if what I argued in §IV is correct, not only is that T-sentence not true, every token of it is false.

Notwithstanding that, something still remains to be explained, and the challenge for me is to show how to explain it in a way that is consistent with the arguments I've given. If we continue with the Mentalese story, the fact standing in need of explanation may be put like this:

Suppose Harold believes that Jane's car is white. Suppose, too, that that belief is realized by the Mentalese sentence 'Jane's car is white' being tokened in Harold's belief box. Then the fact that the sentence's tokening results in his having that belief must be explained at least partly in terms of properties of the syntax and words that compose the sentence. What remains to be explained is what those properties are and how it is that, owing at least in part to them, the sentence's tokening in Harold's belief box realizes his believing that Jane's car is white.

So far, so good; but, you'll recall, the challenge at this point is to assume that what I have argued in this paper is correct and then to determine (i) what remains to be explained given that assumption and (ii) whether what remains to be explained can be explained in a way that is compatible with

the assumption. One can't begin to speculate about how the tokening in Harold's belief box of 'Jane's car is white' realizes his belief that Jane's car is white before one has some hypothesis about how to understand the claim that *Harold believes that Jane's car is white*. This is no easy task on the assumption that what I have argued is correct, for, on that assumption, we can't understand Harold's believing *that Jane's car is white* as his standing in the belief relation to the proposition that Jane's car is white nor, indeed, to anything else, and we must recognize that the penumbral profile of one's uttered token of 'Jane's car is white' in one's uttered token of 'Harold believes that Jane's car is white' is very unlikely to be the same as the penumbral profile of the token of 'Jane's car is white' tokened in Harold's belief box. The starting point in these investigations, I believe, should be with the conceptual roles of Mentalese expressions and syntactic structures, where these are nonsemantic properties that determine the conceptual roles of the sentences those words and structures compose, where the conceptual role of a Mentalese sentence details how that sentence is counterfactually related to sensory stimulations, to other Mentalese sentences, and to bodily movements. One feature of these conceptual roles may be that whenever a sentence σ is tokened in one's belief box, so too is the sentence 'I believe that σ '. How one goes on from there is, alas, an open question.²³

²³ I'm grateful to Crispin Wright for helpful discussions on the topics of this paper. Some of this material was presented in the NYU Mind and Language Seminar that Cian Dorr and Jim Pryor gave in spring 2018. I benefited from their comments and from the comments made by Chris Barker and others during the general discussion.

Works Cited

- Almog, J., Perry, J, and Wettstein, H. (eds). (1989). *Themes from Kaplan* (Oxford University Press).
- Barker, C. (2002). "The Dynamics of Vagueness," *Linguistics and Philosophy*, 25 (1): 1-36.
- Capone, A. (ed). (2019). *Further Advances in Pragmatics and Philosophy, Part II* (Springer).
- Davidson, D. (2001). "Theories of Meaning and Learnable Languages," in *Inquiries into Truth and Interpretation* (Oxford University Press): 3-15.
- Evans, G. and McDowell, J. (eds.) (1976). *Truth and Meaning: Essays in Semantics* (Oxford University Press).
- Fodor, J. A. (1990a). "Stephen Schiffer's Dark Night of the Soul: A Review of *Remnants of Meaning*." In Fodor (1990b): 177-93.
- _____ (1990b). *A Theory of Content and Other Essays* (MIT Press).
- Grice, H. P. (1957). "Meaning," *Philosophical Review*, 66: 377-88.
- _____ (1969). "Utterer's Meaning and Intentions," *Philosophical Review*, 78: 147-77.
- Gunderson, K. (ed.). (1975). *Minnesota Studies in the Philosophy of Science* (University of Minnesota Press).
- Hale, R., Wright, C., and Miller, A. (eds.). (2017). *Blackwell Companion to the Philosophy of Language, Volume 1, Second Edition* (Wiley-Blackwell).
- Kaplan, D. (1989a). "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals," in Almog et. al. (1989).
- _____ (1989b). "Afterthoughts," in Almog et. al. (1989).
- Kennedy, C. (2007). "Vagueness and grammar: The semantics of relative and absolute gradable adjectives," *Linguistics and Philosophy* 30: 1-45.
- Lepore, E. and Ludwig, K. (2005). *Donald Davidson: Meaning, Truth, Language, and Reality* (Oxford University Press).
- Lewis, D. (1969). *Convention: A Philosophical Study* (Harvard University Press).
- _____ (1975). "Languages and Language," in Gunderson (1975): 3-35.
- _____ (1983). "New Work for a Theory of Universals," *Australasian Journal of Philosophy*,

61: 343-377.

Loar, B. (1975). "Two Theories of Meaning" in Evans and McDowell (1975).

_____ (1981). *Mind and Meaning* (Cambridge University Press).

McGinn, C. (1982). *The Character of Mind* (Oxford University Press).

Neale, S. (2016). "Silent Reference," in Ostertag (2016): 229-342.

Ostertag, G. (ed.) (2016). *Meanings and Other Things: Themes from the Work of Stephen Schiffer* (Oxford University Press).

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

Parret, H. and Bouveresse, J. (eds.) (1981). *Meaning and Understanding* (de Gruyter).

Schiffer, S. (1972). *Meaning* (Oxford University Press).

_____ (1981a). "Truth and the Theory of Content." In Parret and Boveresse (1981).

_____ (1981b). "Indexicals and the Theory of Reference," *Synthese*, 49: 43-100.

_____ (1987). *Remnants of Meaning* (MIT Press).

_____ (2003). *The Things We Mean* (Oxford University Press).

_____ (2016). "Reply to Neale," in Ostertag (2016): 495-501.

_____ (2017). "Intention and Convention in the Theory of Meaning," in Hale et. al. (2017): 49-72.

_____ (2019). "Vague Speaker-Meaning" in Capone (2019).

Stojnić, U. and Stone, M. (manuscript). "Vague Utterances in Context."