

# Plurivaluationism, supersententialism and the problem of the many languages

Rohan Sud<sup>1</sup> 

Received: 15 March 2017 / Accepted: 13 March 2018  
© Springer Science+Business Media B.V., part of Springer Nature 2018

**Abstract** According to the plurivaluationist, our vague discourse doesn't have a single meaning. Instead, it has many meanings, each of which is precise—and it is this plurality of meanings that is the source of vagueness. I believe plurivaluationist positions are underdeveloped and for this reason unpopular. This paper attempts to correct this situation by offering a particular development of plurivaluationism that I call supersententialism. The supersententialist leverages lessons from another area of research—the Problem of the Many—in service of the plurivaluationist position. The Problem reveals theoretical reasons to accept that there are many cats where we thought there was one; the supersententialist claims that we are in a similar situation with respect to languages, propositions and sentences. I argue that the parallel suggested by the supersententialist reveals unappreciated advantages and lines of defense for plurivaluationism.

**Keywords** Vagueness · Plurivaluationism · Supervaluationism · Problem of the many

The central claim of what Smith (2008) has dubbed *plurivaluationism* is, roughly:<sup>1</sup>

**Central Claim (Rough)** Our vague discourse doesn't have a single meaning. Instead, it has many meanings, each of which is precise.

---

<sup>1</sup> I am considering the view that Smith calls *classical* plurivaluationism. For ease of presentation, I drop the 'classical' qualification. Note, however, that as Smith uses the term, unqualified plurivaluationists are not committed to the claim that the many meanings are classical or precise. Nor are they committed to the claim that the plurality of meanings is the source of vagueness. (Smith ultimately defends a view he calls fuzzy plurivaluationism, which differs from classical plurivaluationism in denying both of these claims.)

---

✉ Rohan Sud  
rsud@umich.edu

<sup>1</sup> Department of Philosophy, Bates College, Lewiston, USA

According to the plurivaluationist, it is this plurality of precise meanings that explains vagueness. Plurivaluationist views have few, if any, adherents.<sup>2</sup> Instead, plurivaluationism is often dismissed in the literature. This has been a mistake: the position is simply under-developed.

What plurivaluationist views lack is a developed account of our talk about linguistic entities (e.g. talk of languages being spoken or sentences being true), including a view of the ontology of these linguistic objects. To give the problem a name, what the plurivaluationist strategy needs is a developed *talk-of-language account* to fit with the plurivaluationist's **Central Claim**. With the right talk-of-language account, I claim, the plurivaluationist strategy can be made compelling. In this paper, I will outline a talk-of-language account which, when coupled with the **Central Claim**, constitutes a particular plurivaluationist view that I call *supersententialism*.<sup>3</sup>

The supersententialist's preferred talk-of-language account draws lessons from another area of research—the Problem of the Many (PoM). The PoM reveals strong theoretical reasons to accept that there are many ordinary objects of a certain type (e.g. cats) where we pre-theoretically thought there was only one. The supersententialist claims that we are in a similar situation with respect to languages, propositions and sentences—we have strong theoretical reason to accept that there are many languages, propositions, and sentences where we pre-theoretically thought there was only one. The suggested parallel allows us to leverage research on the PoM in service of plurivaluationism. Our talk of linguistic objects (e.g. languages and sentences) is just an instance of our talk of other ordinary objects.

I claim that, with this talk-of-language account in hand, supersententialism is a compelling development of the plurivaluationist strategy. I will make my case indirectly, by arguing that supersententialism is preferable to the most popular approach to vagueness, known as *supervaluationism*. I make this case in two steps. In the first step, I will argue that the supersententialist development of plurivaluationism has a key advantage over supervaluationism. In particular, I will argue that unlike the supervaluationist, the supersententialist can vindicate the disquotational feature of the truth predicate – and can do so by drawing on lessons from the PoM – making the supersententialist's account of truth preferable. In the second step, I will further incorporate lessons from the PoM in order to defend supersententialism from various objections to the plurivaluationist strategy, including those raised by Williamson (1994a), Keefe (1998, 2000), and Schiffer (1998).

<sup>2</sup> Although Smith is the first to delineate (classical) plurivaluationism from supervaluationism, he does not endorse the view. Instead he endorses fuzzy plurivaluationism (see fn. 1). Lewis (1969, 1970, 1975) and Burns (1991) are read by Smith (2008) as endorsing plurivaluationism. (Although I disagree with this reading: see the end of Sect. 6 and fn. 27.) In the course of discussing a problem with speech reports and semantic plasticity, Dorr and Hawthorne (2014, §5 (esp. p. 333)) discuss a view of vagueness on which we are speaking many precise languages simultaneously but the view is not given an extended treatment. In his (2014, §3.4 (esp. fn. 13)), Dorr voices his endorsement of the view—although it is not defended. See also Varzi (2007) who delineates a view on which precisifications are treated as precise Lewisian languages. It's not clear, however, whether the view he outlines is meant to include the claims **Assertive Pluralism** or **Many Languages** below.

<sup>3</sup> Thanks to an anonymous referee for encouraging me to get clearer on the relationship between supersententialism and plurivaluationism.

## 1 Supervaluationism

In order to introduce the plurivaluationist strategy, it's helpful to begin with the more familiar, and more popular, supervaluationist approach to vagueness. Supervaluationism, at least as I use the term here, includes two claims. First: we speak a single vague language with many classical *precisifications*, where a precisification is a way of making the vague language precise which preserves the meanings that are settled in the vague language. Second:

**Supervaluationist Truth (ST)** A sentence of the vague language is true (false) when it is true (false) according to all of the precisifications of the language.<sup>4</sup>

The precisifications thus determine a unique set of truth- and falsity-conditions for each sentence of the single vague language we speak. In this sense, our discourse has a particular unique meaning according to the supervaluationist.

These two claims can be made formally rigorous as follows. We begin by defining an *uninterpreted language* where an uninterpreted language consists of a set of strings of symbols together with rules for classifying those strings into well-formed syntactic types (e.g. sentences, names, etc.). This uninterpreted language, which we can treat as first-order, is assigned meanings by a *supervaluationist model*, which we can take to be a set of classical models with a common domain.<sup>5</sup> The classical models contained by the supervaluationist model each assign precise meanings (e.g. intensions) to strings of symbols in the uninterpreted language in the standard way. These models thus represent the precisifications of the vague language. As is usual, a sentence of the uninterpreted language is true on one of the contained classical models when the intension assigned to that sentence by the classical model is true when evaluated at the actual world.

In order to define what it is for a sentence to be *true on the supervaluationist model*, supervaluationists define the technical notion of *supertruth* (*superfalsity*):  $s$  is supertrue (superfalse) on a supervaluationist model  $\mathcal{M}_S$  iff for all classical models  $\mathcal{M}_C$  in  $\mathcal{M}_S$ ,  $s$  is true on  $\mathcal{M}_C$ . In accordance with (ST), supervaluationists then identify truth (falsity) on a supervaluationist model with supertruth (superfalsity) on that model.

We have now defined the technical notion truth on a supervaluationist model. It remains, however, to link this technical notion with the ordinary notion of truth simpliciter—an unrelativized notion of truth. Supervaluationists, following semantic orthodoxy, link the notions via an *intended model*. A model is intended when, roughly,<sup>6</sup> the model assigns meanings to the parts of uninterpreted language that those parts in fact have. Supervaluationists claim that there is a unique intended model and that a sentence is true (false) just in case it is true (false) on that intended model.

<sup>4</sup> A sentence of the vague language is true on a precisification when, roughly, the sentence of the language that would result from that way of making the vague language precise is true when evaluated at the actual world.

<sup>5</sup> A more powerful model theory would more closely resemble the model theory of quantified modal logic (without a world designated as actual), thus allowing us to assign meanings to a richer uninterpreted language that contains a determinacy operator. The simpler model theory, however, suffices to illustrate the main points of this paper.

<sup>6</sup> cf. Smith (2008, pp. 31–32). I offer a slightly revised definition in fn. 36.

A further question one might ask is what ‘meta-semantic’ facts make a particular model the intended one. That’s a matter of substantial debate (plausible candidates include facts involving linguistic conventions and intentions, causal relations, and naturalness). What’s important for our purposes is that the supervaluationist believes that these meta-semantic facts succeed in singling out a particular supervaluationist model as intended and then use that model to link the ordinary notion of truth simpliciter with their model theory.

## 2 Supervaluationism and the truth objection

Supervaluationists accept (ST) in order to explain the unassertibility of borderline sentences and their negations. As an example consider the borderline sentence:

(1) Bob is bald.

where Bob is a borderline case of baldness. Borderline sentences like (1) and their negations are supposed to be true on some but not all precisifications. In virtue of accepting (ST), this implies that neither borderline sentences nor their negations are true. This result, combined with a maxim to assert only truths:

**Truth-Assertion (TA)** Don’t assert something unless it is true.

allows us to explain our hesitation in asserting either borderline sentences (like (1)) or their negations.

The model theory reflects this. On the intended model, borderline sentences are neither supertrue nor superfalse: they are true on some of the classical models and false on others. Given the definition of truth on a model, these borderline sentences are not true on the intended model and thus not true. Similarly with their negations.

Supervaluationists’ adherence to (ST), however, leads to the Truth Objection: some instances of the T-schema (‘ $p$ ’ is true iff  $p$ ) are not true and not assertible. From a borderline sentence like (1), there are two paths to arriving at the untruth of an instance of the T-schema.<sup>7</sup>

**PATH ONE** Instances of excluded middle are true on every classical model. They are therefore supertrue on every supervaluationist model, including the intended model, and therefore true simpliciter. In particular, the relevant instance:

(2) Bob is bald or it’s not the case that Bob is bald.

is true. Now suppose for *reductio* that the supervaluationist accepts the relevant instances of the T-schema:

(3) ‘Bob is bald’ is true iff Bob is bald.

(4) ‘It’s not the case that Bob is bald’ is true iff it’s not the case that Bob is bald.

Substituting the left-hand sides of (3) and (4) into (2) we get:

(5) ‘Bob is bald’ is true or ‘It’s not the case that Bob is bald’ is true.

<sup>7</sup> Williamson (1994b, p. 162) points out the first path and Keefe (2000, p. 214) points out the second.

By (ST), (5) is claiming that either a sentence or its negation is true on every precisification. But we concluded above that neither (1) nor its negation were true on every precisification. Short of giving up some classical rule of logic, supervaluationists are forced to admit that some instances of the T-schema—like (3) and (4)—are not true and, by (TA), not assertible.

PATH TWO Consider again the relevant instance of the T-schema:

(3) ‘Bob is bald’ is true iff Bob is bald.

From (ST) we learned that (1) is not true (because it’s true on some precisifications and false on others):

(6) It’s not the case that ‘Bob is bald’ is true.

Importantly, the claim made by (6) is not borderline:<sup>8</sup> if it were borderline, then it would not be true, but (6) is true. By (ST) it’s true on every precisification. Because (1) is not false on every precisification, there is some precisification  $p^*$  on which it is true. But then note that (3) is false on  $p^*$ : its left-hand side is false (because it’s false on every precisification) and its right-hand side is true on  $p^*$ . Supervaluationists are therefore forced to admit that some instances of the T-schema—like (3)—are not true and, by (TA), not assertible.

But plainly, the objection goes, all instances of the T-schema are true and assertible—they are not borderline sentences. Indeed, the T-schema features so centrally in our concept of truth that its instances appear to be trivialities. To deny the truth of some instance of the T-schema would be to give up on a central feature of truth and raises the suspicion that supervaluationists have failed to give an account of the everyday notion of truth.<sup>9</sup>

### 3 Plurivaluationism

With the theses of supervaluationism in clear view, we are in a position to outline the plurivaluationist strategy. That’s the task of this section. In further sections, I will suggest a particular development of the plurivaluationist strategy that avoids the Truth Objection.

The plurivaluationist strategy was first delineated from supervaluationism by Smith, who describes the distinguishing feature of the view as follows:

Plurivaluationism...countenances only classical models – but **instead of holding that a vague discourse is associated with a single intended model it holds that it is associated with multiple equally acceptable models**. These models are precisely the classical models that [represent the precisifications in] the (uniquely intended) tripartite base model in the supervaluationist picture. (Smith 2015, p. 1247, emphasis added)

<sup>8</sup> I am ignoring higher-order vagueness here. With higher-order vagueness, the point is simply that (6) does not adopt the first-order vagueness of (1).

<sup>9</sup> Keefe (2000) tries to lessen the incredulity of the supervaluationist position by pointing out that they can accept that the right and left hand side of instances of the T-schema are *mutually entailing*. This, however, is only the case on a impoverished view of supervaluationist entailment that is objectionable for other reasons. The problem re-arises for more sophisticated views (cf. Williams (2008)).

Following Smith, we can state the **Central Claim** of plurivaluationism more carefully as:

**Central Claim (Refined)** Our uninterpreted language has several intended classical models.

I take this claim to distinguish the plurivaluationist strategy. For the supervaluationist, there is one intended model that contains many classical models serving a purely instrumental function of determining the semantic value assigned by the supervaluationist model that contains them. For the plurivaluationist, there are just the classical models, several of which are intended.

One *prima facie* reason in support of the **Central Claim** is this. Whatever the metasemantic facts are that determine correctness for a model, it's incredible to claim that those facts are rich enough to pick out a single model as intended. (The epistemicist's commitment to this claim is one of the major reasons that the view is so implausible.) The plurivaluationist agrees and claims that those facts only delimit a *range* of models as correct.

The **Central Claim** has an immediate consequence. Recall: a model is intended iff it assigns the meanings to the parts of the uninterpreted language that they in fact have. If there are *multiple* intended models, then those parts of uninterpreted language have multiple correct meanings.<sup>10</sup> In typical contexts, we assert the meanings of the sentences we utter. So **Central Claim** seems to imply that, when we utter a sentence of uninterpreted language, we are asserting *many* precise contents—the various precise contents assigned by each of the many intended classical models.<sup>11</sup>

**Assertive Pluralism** When we utter an uninterpreted sentence, we assert each of the contents assigned to that sentence by any of the intended models.

**Assertive Pluralism** is not a bug of the plurivaluationist view: the plurality of asserted contents figures crucially in the plurivaluationist explanation of why we hesitate to utter either a borderline sentence (like (1)) or its negation. According to the plurivaluationist, if I were to utter (1), I would assert many precise contents, some of which are not true. In order to conform to (**TA**), I refrain from uttering (1). Similarly with the negation of (1). So I stay silent.

#### 4 Plurivaluationism and talk-of-language accounts

I take the **Central Claim** (and its consequence **Assertive Pluralism**) to define the class of plurivaluationist views and to constitute the core of any view in that class. But, a complete plurivaluationist view must couple that core with what I earlier called a talk-of-language account—an account of what linguistic objects there are and our ordinary talk about those linguistic entities. Most relevantly, that ordinary talk includes, importantly, predicating truth simpliciter to sentences. Recall that the supervaluationist

<sup>10</sup> The conclusion also follows by similar reasoning from the revised definition of an intended model I give in fn. 36.

<sup>11</sup> Smith agrees that plurivaluationism has this consequence: see his (2011, §2.6) and (2015, p. 1248).

was able to link the technical notion of (super)truth in a supervaluationist model with the ordinary notion of truth simpliciter via a unique intended model: a sentence is true simpliciter just in case it is true on the unique intended model. The plurivaluationist, on the other hand, denies that there is a unique intended model, so it's not obvious how they can link the technical notion of truth on a classical model with truth simpliciter. This is one of the main jobs of what I'm calling a talk-of-language account.

In this section, I hope to illustrate the importance of plurivaluationist talk-of-language accounts by discussing the talk-of-language account sketched by Smith. I'll raise some reasons his talk-of-language account is unsatisfactory. Most importantly, I will argue that his talk-of-language account succumbs to its own version of the Truth Objection. In Sects. 6–8, I will introduce and defend my own preferred talk-of-language account which avoids the Truth Objection and allows a full vindication of the T-schema.

#### 4.1 Smith's talk-of-language account

When delineating plurivaluationism from supervaluationism, Smith offers a brief discussion of his own talk-of-language account. Despite the brevity, we can extract three theses of Smith's account which together provide an illustrative prototype of a talk-of-language-account. Those three theses are as follows:

**Thesis 1 (Smith)** Sentences are, strictly and literally speaking, not true simpliciter.

As the plurivaluationist lacks the idea of a unique intended model of vague discourse, he cannot follow the route of identifying truth simpliciter for statements in that discourse with truth on the intended model. When it comes to truth, he can only say, of each acceptable model, whether a statement is true or false on that model....there can be no further overarching semantic fact, such as that the statement is 'true simpliciter'. (Smith 2015, p. 1248)

**Thesis 2 (Smith)** We can talk as if a sentence is true simpliciter iff the sentence is true on *all* of the intended models.

Suppose that a given sentence is true on *all* of [the intended models]. Then in an obvious sense, it *does not matter* that we have many acceptable interpretations and not just one: we can still say that our sentence is true....But if our sentence is true on some acceptable interpretations and false on others, then it *does* matter...: we can say *neither* that our sentence is true *simpliciter*, *nor* that it is false *simpliciter*. (Smith 2008, p. 99)

**Thesis 3 (Smith)** If all of the contents I assert share a truth-value, then we can talk as if there is only one content.

...if *all* the claims I make are true (or false)...then we can pretend (talk as if) I make only one claim, which is true (or false). (Smith 2011, p. 8)<sup>12</sup>

<sup>12</sup> I take Smith to be talking about propositions by 'claims' (rather than sentences) for he takes claims to be true or false, but **Thesis 1** denies that the notion of truth (simpliciter) applies to sentences.

## 4.2 Concerns with Smith's account

There are at least three reasons why we should not rest content with Smith's talk-of-language account.

First, because Smith's discussion doesn't go much beyond the quoted passages,<sup>13</sup> the details of his account are difficult to decipher. For example: it's not clear what Smith means in **Theses 2 and 3** by 'talking as if' there is one content or that there is a sentence that is true simpliciter. Some passages (e.g. the last quote) suggest we are *pretending* that there is a single intended interpretation. Other passages suggest that our 'talk as if' is *convenient abbreviation* (e.g. when he writes: "we are *merely* summing up the actual, long description of the semantic facts...in a short phrase." (2008, p. 109)). On this latter suggestion, the claim that a sentence is true simpliciter is shorthand for the claim that the sentence is true on all of the intended models. Perhaps a similar story can be told for **Thesis 3**.

Second and relatedly, a more developed account would explain *why* we talk as if as a sentence is true simpliciter or as if there is one content. Such an explanation might, for instance, subsume **Theses 2 and 3** under more general principles of when we 'talk as if' certain claims are true.

Third, as it stands Smith's talk-of-language account fails to provide an adequate account of our ordinary talk of sentences being true. In particular, the view seems to predict that instances of the T-schema are unassertible and is thus open to an analogue of the Truth Objection. Recall that the supervaluationist's commitment to **(ST)** led them to the Truth Objection. Instances of the T-schema such as:

(3) 'Bob is bald' is true iff Bob is bald.

were not true on every precisification, and thus were not true and not assertible. This indicated that the supervaluationist treatment of truth is mistaken because (3) is plainly assertible. However, there is a strong case to be made that Smith's talk-of-language account raises analogous problems.

Note first that, according to Smith's talk-of-language account, instances of the T-schema like

(7) 'Grass is green' is true iff grass is green.

are not assertible in strict and literal contexts. To see this recall that, according to the plurivaluationist, we can assert (7) only if it's true on every intended model (this was how they explained, via **Assertive Pluralism** and **(TA)**, why we do not assert borderline sentences or their negations). In strict and literal contexts, the right-hand side is true on all of the intended models. However, by **Thesis 1**, in such contexts, the left-hand side is false on every intended model. (According to **Thesis 1**, strictly and literally speaking, no sentence is true simpliciter—at best they are true-relative-to-a-model. So, in strict and literal contexts, we can say "Grass is green' is not true'. So, the previous sentence is true on every intended model. Because the models are classical, the sentence "Grass is green' is true' is false on every model.) Thus, the entire biconditional is false on every model and is thus unassertible. Indeed, the negation of (7) is assertible!

<sup>13</sup> This is not a criticism of Smith: his primary goal is to delineate supervaluationism and plurivaluationism.



A defender of Smith's talk-of-language account might accept this result, but lean on **Thesis 2** in order to claim that, in ordinary contexts, we can still talk as if the T-schema holds, thus alleviating the perceived implausibility of strict and literal denials of the T-schema. However, it's far from clear whether such a strategy will work when we consider those instances of the T-schema that caused trouble for the supervaluationist—instances like (3) that involve borderline sentences. On the first path to the Truth Objection to the supervaluationist, we noted that the T-schema immediately implies that either a sentence or its negation is true. Presumably, a denial of this immediate consequence of the T-schema is just as implausible as a denial of the T-schema itself. But consider again (5):

(5) 'Bob is bald' is true or 'It's not the case that Bob is bald' is true.

By **Thesis 2**, neither disjunct is assertible on their own. Depending on how we read Smith's account, this may be because we drop the pretense that there is a single intended model and switch into literal talk or because the truth simpliciter claim is convenient shorthand for the claim that the sentence is true on every intended model. If it's the former, we presumably also drop the pretense when evaluating the entire disjunction (5). In that case, we wouldn't assert the disjunction (when taken literally, the intended models all assign false to the disjuncts). If it's the latter, (5) is presumably shorthand for the claim that

(8) 'Bob is bald' is true on every intended model or 'It's not the case that Bob is bald' is true on every intended model.

But (8) is not the case: neither the sentence nor its negation is true on every intended model.

A similar argument follows an analogue of the second path to the Truth Objection. We noted that the supervaluationist cannot accept the right-to-left direction of (3) because the left-hand side is superfalse but the right-hand side is not. Consider the relevant conditional

(9) If Bob is bald, then 'Bob is bald' is true.

By **Thesis 2**, the consequent is unassertible on its own. Depending on how we read Smith, this may be because we drop the pretense that there is a single intended model and switch into literal talk or because the truth simpliciter claim is convenient shorthand for the claim that the sentence is true on every model. If it's the former, we presumably also drop the pretense when evaluating the entire conditional (9). In that case, we wouldn't assert the conditional because, by **Assertive Pluralism**, we would be asserting some false proposition (when taken literally, one of the intended models assigns true to the antecedent and false to the consequent). If it's the latter, (9) is presumably shorthand for the claim that

(10) If Bob is bald, then 'Bob is bald' is true on every intended model.

But this abbreviated claim is false on some intended model.

Thus, it's not clear that Smith's talk-of-language account conforms to our ordinary talk of sentences being true. I don't take this concern for his account to be insuperable. Perhaps Smith's talk-of-language account can be finessed—for instance by offering

a more nuanced story of how abbreviated descriptions and pretense work in disjunctions and conditionals than the simple story presumed above. This suggested strategy, however, is difficult to assess precisely for the first two reasons I gave for not resting content with Smith's talk-of-language account: that account is missing details and a unifying explanation.

\*\*\*

Given the difficulties with Smith's talk-of-language account, I propose we look for an alternative. In Sects. 5–8, I develop my preferred alternative. Before we turn to that alternative, however, it's instructive to briefly consider a different alternative to Smith's talk-of-language account, one that I will ultimately set aside. According to this tempting proposal, the predicate 'true' is a *vague* predicate that expresses differing properties on differing intended models. In particular, according to this proposal, on each intended model  $\mathcal{M}$ , 'true' expresses the property had by uninterpreted sentences of being true-on- $\mathcal{M}$ .<sup>14</sup> Similarly for the predicate 'false'. And, in a typical utterance, we token a single sentence that we can choose to predicate these vague terms 'true' and 'false' to. Unlike Smith's talk-of-language account, this proposal does *not* accept **Thesis 1**: we recover strict and literal talk of sentential truth. There is no trouble with accepting simple predications of truth like "Grass is green' is true", even in strict and literal contexts. And, most importantly, instances of the T-schema like (3) and (7) are true on every intended model, even in strict and literal contexts. On this proposal, borderline sentences are true or false, although it's indeterminate *which* truth-value they have.

The present proposal attempts to give up **Thesis 1** and retain talk of truth simpliciter. And, the proposal's treatment of truth does not conflict with the T-schema in the way Smith's talk-of-language account does. So, it is tempting to think that the plurivaluationist can rehabilitate our talk of truth by following this proposal. This, however, is not the case. The proposal merely moves the bump in the rug without addressing the plurivaluationist's real difficulty with the notion of truth simpliciter. That difficulty arises from the plurivaluationist's commitment to **Assertive Pluralism**.<sup>15</sup>

Recall that, according to **Assertive Pluralism**, when I utter the uninterpreted sentence 'Bob is bald', I am asserting multiple contents some of which are true and some of which are not. And, as we pointed out at the end of Sect. 3, **Assertive Pluralism** is essential to the plurivaluationist's explanation of vagueness: borderline sentences and their negations are unassertible because uttering borderline sentences or their negations requires asserting some contents that are not true, which is impermissible according to **(TA)**. But, **Assertive Pluralism** appears to conflict with the more general truth schema **(T\*)**: If  $S$  says that  $p$ , then  $S$  is true iff  $p$  is true (cf. Williamson (1994a)).<sup>16</sup> **(T\*)** captures the popular claim that a sentence inherits the truth-value of what the sentence says and implies the following principle of uniformity:<sup>17</sup> a sentence cannot say a true

<sup>14</sup> Thanks to an anonymous referee for pressing me to discuss this alternative.

<sup>15</sup> This point is broadly inspired from related discussions in Williamson (1994a, b), Andjelković and Williamson (2000).

<sup>16</sup> I am ignoring context sensitivity and disputes about the existence of propositions.

<sup>17</sup> Cf. Andjelković and Williamson (2000). Although I won't discuss it here, several authors have (i) suggested that the principle of uniformity indicates that **(T\*)** is not definitional of truth and (ii) have

and an untrue proposition.<sup>18</sup> On the present proposal, when I utter ‘Bob is bald’ I am tokening a sentence that is true or false, but the sentence says multiple propositions some of which are true others of which are not. Thus, the proposal violates (T\*) and the principle of uniformity.<sup>19</sup> As with Smith’s talk-of-language account, perhaps the present proposal on which ‘true’ is vague can be finessed to avoid this difficulty, but it is not obvious how to do so.

In short: The plurivaluationist explains why borderline sentences and their negations are unassertible by claiming that in uttering such sentences we are asserting a plurality of contents of differing truth-values. But, because sentences adopt the truth-value of propositions they express, this makes it mysterious which truth-value to assign to the borderline sentence. Reflecting on this concern makes clear the motivation behind the rejection of strict and literal talk of truth in Smith’s talk-of-language account, a motivation that is not clearly overcome by taking the truth predicate to be vague.

At this point we might wonder: Is there *any* talk-of-language account that retains the plurivaluationist’s claim of **Assertive Pluralism**, yet vindicates our ordinary talk—given by the T-schema and (T\*)—of sentences being true? For the rest of the paper, I will develop such an account. The simple idea behind my account is to pair the plurivaluationist’s plenitudinous ontology of propositions asserted with a plenitudinous ontology of sentences tokened, such that each sentence says exactly one proposition. This plurality of sentences allows us to fully vindicate talk of sentential truth in conformity with the T-schema and (T\*). And unlike the tempting proposal just considered, we need not reduce the monadic notion of sentential truth to a relativized notion of truth-on-this-or-that-model.<sup>20</sup> I call the resulting version of plurivaluationism ‘supersententialism’. The best way to motivate and defend the supersententialist’s plenitudinous ontology of sentences requires us to consider a different dialectic—the debate over the Problem of the Many. I turn to that debate in the next section.

---

Footnote 17 continued

explored whether to accept the principle of uniformity. See Dutilh Novaes (2008) for helpful discussion and references. Thanks to Catarina Dutilh Novaes for discussion.

<sup>18</sup> Proof: Suppose for *reductio* that *S* says propositions  $q_1$  and  $q_2$ , and  $q_1$  is a true proposition, but  $q_2$  is not a true proposition. Then one instance of (T\*) implies that *S* is true, and another instance of (T\*) implies that *S* is not true. Contradiction.

<sup>19</sup> Faced with this difficulty, might we continue to hold that the truth predicate is vague but simply give up **Assertive Pluralism**? We might instead claim that in a given utterance, we assert just one precise proposition, but it’s vague which precise proposition is the one we assert. Because this position gives up **Assertive Pluralism**, it is no longer a plurivaluationist position (according to the envisaged theory one model is intended, but it’s vague which model that is) and cannot help itself to the plurivaluationist’s explanation for the unassertibility of borderline sentences. Instead the position becomes a form of what’s know as *non-standard supervaluationism* (cf. the position in Field (1994)) according to which our bivalent sentences have a single precise meaning but it’s indeterminate which meaning it has. Although I won’t discuss non-standard supervaluationism here, see Williamson (1994b, p. 164) for a criticism (He argues that the non-standard supervaluationist fails to give a non-circular account of the ‘philosophical significance’ of indeterminacy.)

<sup>20</sup> Thanks to an anonymous referee for emphasizing this advantage of my proposal.

## 5 The lessons of the many

Consider an ordinary cat, Tibbles, on a mat. Tibbles is shedding and some of her hairs are in the process of leaving her body. Consider one such hair, *h*, such that, of *h*, it's indeterminate whether it is part of Tibbles. Call the mereological sum of Tibbles and *h* 'Tibbles+' and the mereological difference of Tibbles and *h* 'Tibbles-'.<sup>21</sup>

Now consider the following claims:

- (11) a. There is exactly one cat on the mat.  
 b. # Tibbles is Tibbles+.  
 c. # Tibbles is Tibbles-.

Both Tibbles+ and Tibbles- seem to have all of the requisite features of catiness (both meow, have whiskers, etc.) So, both seem to be cats. Yet, the first claim is plainly assertible. Moreover, Tibbles+ and Tibbles- appear to be the *only* candidates for being Tibbles. So, it seems that Tibbles is identical to Tibbles+ or she is identical to Tibbles-. Yet, the last two claims are unassertible: we hesitate to identify Tibbles with Tibbles+ or to identify her with Tibbles-. This is one case of the infamous Problem of the Many: we must explain the assertibility of the first claim and the unassertibility of the second two.

One strategy for solving the problem rejects the claim that Tibbles+ and Tibbles- are the only candidates for being Tibbles and posits a sort of dualism between Tibbles and the objects Tibbles+ and Tibbles-. According to one development of this strategy, there is a metaphysically vague object that 'Tibbles' determinately refers to (van Inwagen 1990; Tye 1996).<sup>21</sup> Lewis (1993) claims, and I agree, that dualism is a solution of last resort. This follows from a hesitation to posit vague objects based solely on non-scientific, philosophical observations. This is the first lesson to draw from the Problem of the Many:

**No Dualism** Don't posit a dualism between a 'vague object' and a 'precise object'; there are no vague objects.

Instead, Lewis advocates solutions that locate the source of the Problem in our language, while retaining the anti-dualist claim that Tibbles is Tibbles+ or Tibbles-. He proposes two such solutions. On the first solution, Lewis rejects the claim that *both* Tibbles+ and Tibbles- are cats (perhaps because a requisite feature of being a cat is not overlapping with another cat), and deploys the supervaluationist machinery to develop the claim that it's vague which is a cat. On this proposal, the predicate 'cat' is vague: on one precisification, Tibbles+ is in the extension of 'cat' and on another precisification Tibbles- is in the extension. To explain the assertibility of (11a), in no precisification are both Tibbles+ and Tibbles- within the extension of 'cat'. In order to explain the unassertibility of (11b) and (11c), the supervaluationist claims that 'Tibbles' refers to Tibbles+ on some precisifications and Tibbles- on others. And, in order to explain the assertibility of a claim like 'Tibbles is a cat', they posit a 'penumbral

<sup>21</sup> On another development, Tibbles is not identical to any material composite, but is rather indeterminately constituted by one of the material composites (Lowe 1995) or constituted by multiple material composites (Jones 2015).

connection’ (a connection between meanings that is preserved in all precisifications) between the meanings of the name ‘Tibbles’ and the predicate ‘cat’ such that, for any precisification, the referent of ‘Tibbles’ according to that precisification is in the extension of ‘cat’ according to that precisification. Call this the *Solution by Vagueness*.

On Lewis’s second solution, which we can call the *Solution by Plenitude*, both Tibbles+ and Tibbles– are cats. In order to explain the assertibility of (11a), Lewis claims that ordinary notions of counting are by ‘almost-identity’ rather than identity, where two objects are almost-identical just in case they massively overlap. The Solution by Plenitude admits that, strictly speaking, there are many cats on the mat. The plenitudinous-philosopher, therefore, rejects the vagueness-philosopher’s insistence that the predicate ‘cat’ is vague.<sup>22</sup> However, the plenitudinous-philosopher follows the vagueness-philosopher in treating the singular terms ‘Tibbles’ and ‘the cat on the mat’ as vague. That explains our hesitation in asserting (11b) or (11c) and the assertibility of claims like ‘Tibbles is a cat’.

There has been an ongoing debate over which is the right solution to the Problem of the Many.<sup>23</sup> However, I believe the research underlying the Solution by Plenitude can be fruitfully applied in defense of plurivaluationism. In addition to **No Dualism**, there are at least three more lessons worth distilling from the Solution by Plenitude:

**Plenitudinous Ontology** Strictly speaking, there are many cats on the mat.

**Almost-Identity** We count by almost-identity rather than strict identity.

**Vague Singular Terms** Our singular terms vaguely refer to one of the many cats on the mat.

With these four lessons from the Problem distilled, I am in a position to present, over the next two sections, my preferred talk-of-language account to pair with the **Central Claim** of plurivaluationism.

## 6 Many languages

When we reflected on the nature of cats, we noticed pressure to admit that there are many cats where we ordinarily thought there was one. As we did with cats, we can ask about the ontology of *languages*. What is a language?

In Sect. 1, we said what an *uninterpreted* language is: it’s a set of strings of symbols together with a grammar—rules for classifying those strings into well-formed syntactic types. But *languages*—the sorts of things we name with terms like ‘English’ or ‘Hindi’—are more than syntactic structures. One does not become fluent in Spanish simply by knowing the grammar of the language. We can emphasize the distinction

<sup>22</sup> At least vague in the same way that the vagueness-philosopher claims—the plenitudinous-philosopher might admit vagueness as to which plurality of cat-candidates count as the many cats on the mat.

<sup>23</sup> Lewis, for instance, claimed that both solutions are adequate, and that each requires elements of the other. See Williams (2006) for one powerful argument in favor of the plenitude solution. See Weatherson (2003) for an excellent defense of the Solution by Vagueness.

between languages and uninterpreted languages by calling the former ‘interpreted languages’.

So, what is an (interpreted) language? Lewis (1975) answers this question by treating languages as functions from sentences of an uninterpreted language to sentential semantic values. We can extend Lewis’s definition by also adding assignments from subsentential expressions in the uninterpreted language to subsentential semantic values.<sup>24</sup> Each language is precise in that it associates each uninterpreted sentence with a precise content—a content that is true or false when evaluated at any possible world. On this conception of languages as functions, there are a massive number of languages. Languages can thus be associated with classical models with the corresponding assignments of meanings. Just as we talked of a sentence of uninterpreted language being true on a model, we can talk of a sentence of uninterpreted language being *true on a language*.

With languages thus understood, we can ask a related question: what does it mean for a community to speak (or use) a particular language? A minimal answer is, roughly,<sup>25</sup> a community speaks a language when the parts of the uninterpreted language that the community utters have the meanings assigned to them by that language. We can put this in terms of intended models: a community speaks a language when the model associated with that language is intended.

A more informative answer would also explain in virtue of which meta-semantic facts we speak the languages that we do. Indeed, Lewis’s early works are aimed at giving such an explanation in terms of prevailing conventions in the linguistic population. We noted earlier that, regardless of how this more informative answer is spelled out, it is incredible to think that the meta-semantic facts are rich enough to single out a unique language as spoken. In other words: several of the precise languages will be on par with respect to these meta-semantic facts. Our position with respect to languages being spoken by a community is thus analogous to the position we are in with respect to cats on a mat upon reflecting on the Problem of the Many.

Consider that position. There are many mereological sums that are equally good candidates for being (what we presumed to be) the single cat on a mat. So: which is the cat? One answer posits a sort of dualism according to which there is a ‘vague’ cat over-and-above the ‘precise’ mereological sums. Another solution, the Solution by Vagueness, claimed that our term ‘cat’ is vague but determinately includes only one mereological sum. However, a third solution – the Solution by Plenitude – accepts that there are many cats where we ordinarily thought there was one.

There are many languages (as Lewis understands them) that are equally good candidates for being (what we presumed to be) the single language we speak. So: which is the language we speak? The supervaluationist posits a sort of dualism according to which there is a ‘vague’ language over-and-above the ‘precise’ Lewisian languages. Another answer, following the Solution by Vagueness, would claim that it’s vague

<sup>24</sup> We might further require that the assignments conform to compositionality requirements so that the semantic value of a composite expression in the uninterpreted language is composed of the semantic values assigned to the parts of that expression. Languages are thus the functions generated by what Lewis calls ‘grammars’ on p. 175 of his (1975).

<sup>25</sup> cf. fn. 36.

which of the many candidate-languages delimited by the meta-semantic facts is the one we speak.<sup>26</sup> A third solution—adopting **Plenitudinous Ontology**—accepts that:

**Many Languages** Strictly speaking, we are speaking many languages simultaneously.

I call the position that accepts **Many Languages** (together with additional theses to be spelled out below) *supersententialism*. Given the Lewisian understanding of languages, **Many Languages** is a translation of the **Central Claim** of plurivaluationism. If we speak those languages with associated models that are intended, and there are many intended models, then we speak many languages. I thus take supersententialism to be a form of plurivaluationism—one that combines the plurivaluationist view with a talk-of-language account inspired by the Problem of the Many and the Lewisian view of the ontology of languages.

(For bookkeeping purposes, supersententialism can also be helpfully classified as a particular enactment of a more general strategy sometimes called *pragmatism*. Pragmatic theories of vagueness treat languages as precise and locate the source of vagueness in our relationship with these precise languages. As Keefe (1998, 2000) notes, there are various ways one might specify this relationship, generating different enactments of the pragmatic strategy. One version—sometimes endorsed by Lewis<sup>27</sup>—suggests speakers switch from one precise language to another depending on highly transient features of context, and is a forerunner to contemporary contextualist theories of vagueness (cf. Fara (2000)). The analogy with the Problem shows us two alternative developments corresponding to the Solutions by Vagueness and Plenitude. Supersententialism, draws inspiration from the Solution by Plenitude, and combines the pragmatic strategy with the plurivaluationist’s **Central Claim** with the upshot of **Many Languages**.<sup>28</sup>)

At this point, we might wonder: what benefit does this way of construing languages confer to the plurivaluationist? The supersententialist’s commitment to **Many Languages** motivates a plenitudinous ontology of other sorts of linguistic objects—

<sup>26</sup> This would be akin to a sort of *non-standard supervaluationism*, according to which the meaning of our language is given by one classical model but it is indeterminate which that is. As Williamson (1994b, p. 164) correctly points out, this suggestion does not give a non-circular account of the ‘philosophical significance’ of indeterminacy.

<sup>27</sup> Cf. Keefe (1998). Lewis writes “we are free to settle these indeterminacies however we like” (1975, p. 188) and “the different languages of the cluster...may be differently suited for individual opinions, tastes, and conversational purposes. If everyone can pick from the cluster, incompatible preferences among languages may all be satisfied” (1969, p. 202). (In other places, it should be noted, Lewis suggests a non-pragmatic treatment of vagueness more akin to supervaluationism.) Linda Burns (1991, p. 182), in developing Lewis’s pragmatism, makes similar statements such as “where there is vagueness speakers must be represented as alternating between members of a range of such languages” (182) and “speakers may adopt different languages from one another and shift from one language to another at different times” (186). Neither Lewis nor Burns make the claim of **Assertive Pluralism** or use it to explain why we do not assert borderline sentences.

<sup>28</sup> Smith (2008) writes “Pragmatism, then, is simply a stylistic variant of plurivaluationism”. As should be clear by now, I disagree on two fronts. First, pragmatism can be developed in ways that are not plurivaluationist. Second, the pragmatic treatment of vagueness makes claims about the ontology of languages that are not part of the **Central Claim** of plurivaluationism. Part of the work of this paper is to show that those claims can form part of a talk-of-language account that can be leveraged to great effect.

asserted propositions and interpreted sentences—for which we’re also facing instances of the Problem. Consider the case of propositions. Because we are speaking many precise languages with precise propositions as contents, in a particular utterance, we assert each of these many precise propositions (**Assertive Pluralism**). And, most importantly, **Many Languages** allows the supersentialist to posit a plenitude of sentences. This is because sentences are individuated by the languages they are of.

To see this, consider a typical utterance of ‘there is a torch’ by a speaker of American-English (in which ‘torch’ refers to objects with open flames) and a separate utterance of ‘there is a torch’ by a speaker of British-English (in which ‘torch’ refers to flashlights).<sup>29</sup> Have these speakers tokened the *very same* sentence? In a sense they have, in a sense they have not. It’s uncontroversial that they have tokened the same string of symbols. That is, they have tokened the same sentence *of the uninterpreted language*. Similarly it’s uncontroversial that this single sentence of uninterpreted language occurred *in* British-English and American-English. But, I claim, they have *also* tokened different sentences. On my proposal, we token sentences *of* the languages we speak. One speaker tokens a sentence *of* American-English and one tokens a sentence *of* British-English. And those sentences are not identical: the sentence of British-English may be true while the sentence of American-English false. These sentences are thus individuated by the languages they are of, so if we speak many languages, then we token many sentences—one for each of the languages we speak. On at least one disambiguation of ‘sentence’, there are sentences individuated by languages. To mark the difference, we can call sentences of uninterpreted language *uninterpreted sentences* and sentences of interpreted language *interpreted sentences*. If we wish, we can treat the latter as ordered pairs of uninterpreted sentences and Lewisian languages.

An alternative ontology might deny that there are distinct sentences *of* British-English and *of* American-English. Instead, on this proposal, there is *only* the sentence of uninterpreted language—and that single sentence occurs *in* American-English and British-English.<sup>30</sup> While we can plainly say of this lonely sentence that it is true on this-or-that language, difficulties arise for understanding attributions of truth simpliciter—attributions that do not mention a language (e.g. ‘all of the sentences on the board are true’ or ‘Bob is bald’ is true if Bob is bald’). If an uninterpreted sentence only occurred in a single language that we use, we could say that an uninterpreted sentence is true simpliciter when it is true on the language it occurs in. But if the uninterpreted sentence occurs in many languages that we speak – as the supersentialist proposes—it’s not clear how to define truth simpliciter. We might try something akin to the strategy Smith used when facing the analogue puzzle with intended models, but I’ve already suggested reasons not to rest content with Smith’s proposal. While I don’t wish to deny that a view with a sparse sentence ontology can be made to work,<sup>31</sup> I do not

<sup>29</sup> Thanks to Chip Sebens for suggesting this example.

<sup>30</sup> This appears to be Lewis’s view. See his (1975, p. 163). I suspect this is the more popular view, and may have been Smith’s view when he gave his talk-of-language account.

<sup>31</sup> For instance, I could imagine someone saying that truth simpliciter is predicated of a plenitude of *occurrences* rather than a plenitude of sentences, where a given uninterpreted sentence can figure in many occurrences by occurring in different languages. Much of what I say below regarding a plenitude of sentences might be parroted for a plenitude of occurrences.



want to pursue that strategy here. Instead, I posit an ontology according to which interpreted sentences are individuated by the languages they are of. Thus, according to the supersentialist:

**Many Sentences** In a typical utterance, we token many interpreted sentences – one of each of the languages that we speak – in addition to the uninterpreted sentence.

In borderline cases, I token several interpreted sentences, some of which are true, and some of which are false. Because we refrain from asserting falsehoods, we hesitate to utter these sentences. Thus, truth of all of the interpreted sentences is a necessary condition for assertibility.

## 7 Return to the truth objection

**Many Sentences** induces an important divergence between the talk-of-language account that Smith sketches and the supersentialist's account. According to the talk-of-language account sketched by Smith, strictly and literally speaking, sentences are not true simpliciter (**Thesis 1**). As we noted above at the end of Sect. 4, he was pushed towards this claim because it was unclear which truth-value to assign a borderline sentence that expresses multiple propositions with different truth-values. The supersentialist's plenitudinous ontology of sentences and languages opens up space for an alternative: there are a plenitude of tokened sentences—one of each language. And each of these sentences expresses a single proposition and adopts the truth-value of the proposition it expresses (in accordance with the uniformity principle of (T\*)). Moreover, pace Smith's **Thesis 1**, we can thus predicate truth simpliciter of each of these many sentences (even in strict and literal contexts). In addition, we can link the model theory as follows. A sentence of an interpreted language is true simpliciter when the corresponding uninterpreted sentence is true on the model that is associated with the language the interpreted sentence is of. And because we speak a language when the model associated with that language is intended, the notion of an intended model also fixes which interpreted sentences we token: for each of the many intended models, we token a distinct sentence of the language associated with that intended model.

This plenitude of sentences also provides the supersentialist with the tools to rebut the Truth Objection. Consider again (3):

(3) 'Bob is bald' is true iff Bob is bald.

We must show how each of the many sentences tokened by an utterance of (3) (each spoken in one of the many languages we simultaneously speak) are true. The right-hand side of the biconditional is true when tokened in some languages being spoken and false when tokened in others; our challenge is to explain how the left-hand side of the biconditional is true when tokened in all and only those languages in which tokens of the right-hand side are true.

By positing a plenitude of sentences being uttered in a single utterance the supersentialist can apply the lesson **Vague Singular Terms** from the Problem of the

Many: we can posit vagueness in the sentence being named by the quotation marks.<sup>32</sup> The supersentialist can also posit a penumbral connection between the quotation “Bob is bald” and the right-hand side ‘Bob is bald’ such that the sentence named in a language by the quotation is true just in case the sentence ‘Bob is bald’ is true on that same language. More carefully: for any language  $L$  being spoken, the interpreted sentence being named in  $L$  by the quotation in an utterance of (3) has the same truth conditions as the sentence of  $L$  ‘Bob is bald’. One way this might happen is if, for any language  $L$  being spoken, the interpreted sentence being named in  $L$  by the quotation in an utterance of (3) is a sentence of  $L$  itself!

This move is available to the supersentialist because she posits several languages being spoken simultaneously, and with it, a plenitudinous ontology of interpreted sentences. Without this ontology, there would be no plenitude of candidates to which the quotation-name can vaguely refer. The failure to appreciate this significant advantage is perhaps what has led many supervaluationists to be dismissive of proposals like supersentialism.

The supersentialist’s talk-of-language account has another benefit: it explains why, in simple attributions of truth, Smith’s **Thesis 2** holds. When an uninterpreted sentence ‘Grass is green’ is true on all intended models, then all of the interpreted sentences that are tokened are true simpliciter. We talk as if ‘Grass is green’ is true because such talk *is in fact true* on each of the languages we speak!

## 8 Objections considered

In this section, I reply to various objections to the view outlined above. One theme running throughout these replies is an exploitation of research done on the Problem of the Many in order to further develop a talk-of-language account which can rebut the present objections.

### 8.1 We speak one language

*Objection* I do not speak many languages. I speak exactly one language: English!

*Reply* Recall our lesson **Almost-Identity** from Lewis’s Solution by Plenitude to the Problem of the Many. There are in fact many cats on the mat but we count only one cat on the mat because we count by almost-identity rather than strict identity, where two objects are almost-identical if they massively overlap. Together with the claim that the name ‘Tibbles’ indeterminately refers to one of the many cats, we predict the relevant data.

The same strategy can be exploited by the supersentialist. There are in fact many languages that we speak, but we count only one language because we count languages by almost\*-identity, where two languages are almost\*-identical when the intensions

<sup>32</sup> Compare the move made at the end of Weatherson (2003) in which Weatherson takes the name of a predicate to be vague in order to rebut an argument by Merricks (2001) for metaphysical indeterminacy.

for the terms massively overlap.<sup>33</sup> Together with the claim that the name ‘English’ indeterminately refers to one of the many languages, we predict the relevant data. (Speaking loosely) we do speak exactly one language: English! The same solution can be extended to sentences tokened or propositions asserted, in the obvious way, to generate the result that (speaking loosely) I token only one sentence and assert only one proposition in a single utterance, thus offering a unifying explanation of Smith’s **Thesis 3**.<sup>34</sup>

Indeed, there is a good case to be made, independently of our discussion of vagueness, that we already have a grip on the notion of almost\*-identity and count languages by it. Compare again American-English and British-English. By Leibniz’s law, these languages are not strictly identical, because ‘torch’ means something different in each. Yet no American-English speaker claims to be bilingual as a result of speaking British-English!

## 8.2 What about the unsharp sentence?

*Objection* You’ve shown that the T-schema holds for each of the precise interpreted sentences that we speak. But that doesn’t solve the original problem. We must show that the T-schema holds for the *imprecise sentence*, not merely the precise interpreted sentences! Williamson (1994a) puts the objection best in the course of dismissing a proposal similar to supersententialism:

The utterance ‘*a* is a heap’ is by ordinary standards the utterance of a single sentence. If it is held that in reality many sentences were uttered, a special way of individuating sentences must be in play...Call the many nonstandardly individuated co-uttered sentences *sharpened*, and the original standardly individuated sentence *unsharpened*. One can still ask for the truth-value of unsharpened sentences...the disquotational principle must fail for unsharpened sentences...According to the original intuition, if *a* is a heap then the unsharpened ‘*a* is a heap’ is true...(180)

*Reply* If by ‘unsharpened sentence’ Williamson means an unsharpened *interpreted* sentence—a sentence of a *vague language*—then he appears to have a different ontology of sentences in mind than my own. According to my view, *there is no such thing* as an unsharpened interpreted sentence, because there are no such things as unsharp or vague languages. The only interpreted sentences are the plurality of ‘sharp’ interpreted sentences. That follows from the lesson **No Dualism** from the Problem of the Many: don’t adopt dualism unnecessarily. To posit an additional object—the unsharpened interpreted sentence—over and above the many sharp interpreted sentences is to

<sup>33</sup> Of course the relevant notion of overlap used to define almost\*-identity is not the mereological notion Lewis used to define almost-identity, but I take it that the notion is obvious enough. For example, take some suitable measure  $m$ . Two intentions  $i_1$  and  $i_2$  massively overlap if for all  $w$  in the domain of both  $i_1$  and  $i_2$ ,  $m(i_1(w) \cap i_2(w))$  is very close to  $m(i_1(w))$  and very close to  $m(i_2(w))$ .

<sup>34</sup> For more on the parallel between counting the many cats and the many languages, including how we react when the differences between the cats and languages are made salient or are relevant to our interests, see *Reply 3* in Sect. 8.3.

adopt a sort of ‘dualism’ about sentences that is no part of my theory. Just as I don’t posit an unsharp cat over and above the many sharp cats when solving the Problem of the Many, I do not posit an unsharp interpreted sentence over and above the many sharp interpreted sentences.

Williamson might suggest that without an unsharp interpreted sentence we cannot accommodate the ‘original intuition’ in favor of the T-schema, which appeared not to be directed towards these precise interpreted sentences. But again, the analogy with the PoM is useful here. If there is a sense in which the original intuition behind predications of truth applied to an unsharp interpreted sentence, then it’s the same sense in which our original intuition behind predications of the cat predicate applied to an unsharp cat. I’ve successfully accounted for both the unsharp-cat intuition and the unsharp-sentence intuition by positing vagueness in the relevant singular terms ‘Tibbles’, ‘the cat on the mat’, and “Bob is bald”.

*Objection (cont’d)* Denying the existence of such an unsharpened interpreted sentence is unmotivated because it’s unmotivated to deny the existence of a vague language. Consider a language which assigns to sentences the truth conditions suggested by the supervaluationist, complete with truth-gaps. We speak this ‘vague’ language, and thus token an interpreted sentence of this language. Here’s an argument from Keefe (2000) given when considering various ways of developing pragmatic treatments of vagueness (à la Lewis and Burns) and arguing that each collapses into supervaluationism:

[The supersentialist might insist] that talk of ‘the vague language’ is not acceptable...[but she] has no grounds on which to maintain this insistence. For she endorses Lewis’s account of the role of beliefs and intentions in fixing the language spoken and hence the meanings of sentences, and if we are seeking to be truthful and trusting in the whole cluster of languages then our beliefs and intentions are related in exactly the appropriate way to the corresponding language whose truth-conditions are given by the supervaluationist account. (144)

The supersentialist intends for her view to be compatible with Lewis’s account that we speak a language iff we are seeking to be trusting and truthful in that language. But we are being trusting and truthful in the vague language posited by the supervaluationist. So we should admit that we speak this vague language.

*Reply* The supersentialist need not insist that talk of ‘the vague language’ is unacceptable. She can accept that a language like the one posited by the supervaluationist exists. What she does not accept is that *we are using this vague language*. According to the supersentialist, Keefe is wrong that “our beliefs and intentions are related in exactly the appropriate way to the corresponding language whose truth-conditions are given by the supervaluationist account”.

On Lewis’s account of which languages we speak, we don’t speak a language  $L^*$  that assigns ‘Tibbles is a cat’ to the proposition that Tibbles is a table. To see why first note that it’s not reasonable to believe that Tibbles is a table. Presumably, we are not radically unreasonable in our beliefs about Tibbles. So, we don’t believe that Tibbles is a table. Yet we assert ‘Tibbles is a cat’. Thus, we are not seeking to be truthful in

any language like  $L^*$  that assigns ‘Tibbles is a cat’ to the proposition that Tibbles is a table. So, on Lewis’s theory we don’t speak  $L^*$ ; if we spoke  $L^*$  we would be saying things that are obviously not true.

Similarly, on Lewis’s account of which languages we speak, we don’t speak the ‘vague’ supervaluationist language that assigns “‘Harry is bald’ is true iff Harry is bald” to the proposition that ‘Harry is bald’ is true on every precisification iff Harry is bald. To see why first note that it’s not reasonable to believe that ‘Harry is bald’ is true on every precisification iff Harry is bald—when Harry is borderline bald, the biconditional is easily verified to be untrue (cf. Sect. 2). Presumably, we are not radically unreasonable in our beliefs about truth under precisifications or the biconditional. So, we don’t believe that ‘Harry is bald’ is true on every precisifications iff Harry is bald. Yet we assert “‘Harry is bald’ is true iff Harry is bald”. Thus, we are not seeking to be truthful in any language like the vague supervaluationist language that assigns “‘Harry is bald’ is true iff Harry is bald” to the proposition that ‘Harry is bald’ is true on every precisification iff Harry is bald. So, on Lewis’s theory we don’t speak the ‘vague’ supervaluationist language: if we spoke the vague language we would be saying things that are obviously not true.

*Objection (cont’d)* But plainly there is *something* that is an ‘unsharpened sentence’, even if it isn’t an interpreted sentence (a sentence of an interpreted language). Consider the individual’s *utterance* as distinct from the many sentences uttered. The T-schema fails for utterances. Or, consider the *uninterpreted sentence* uttered. The T-schema fails for uninterpreted sentences.

*Reply 1* It’s not clear that utterances are bearers of truth. Note the awkwardness of predicating truth to utterances when we are explicit about referring to the utterance rather than the sentence:<sup>35</sup>

- (12) a. # That speech act was true.  
b. # What he did in asserting that sentence was true.

There is even less reason to treat uninterpreted sentences as bearers of truth—at least as bearers of truth *simpliciter*, which I take to be what is predicated in instances of the T-schema. Uninterpreted sentences don’t have meaning—at best they have meanings in a language or meanings assigned by a model. But meaningless objects cannot be true simpliciter. To see this, recall that the T-schema is the result of a more general schema ( $T^*$ ): If  $S$  says that  $p$ , then  $S$  is true iff  $p$  is true. An interpreted sentence ‘Bob is bald’ says that Bob is bald—thus, by ( $T^*$ ), the interpreted sentence ‘Bob is bald’ is true iff Bob is bald. But the uninterpreted sentence doesn’t say anything. So ( $T^*$ ) does *not* entail that the uninterpreted sentence ‘Bob is bald’ is true iff Bob is bald.<sup>36</sup>

<sup>35</sup> This point is from MacFarlane (2014, §3.1.2)

<sup>36</sup> Denying that expressions in an uninterpreted language have meanings requires a revision of the definitions of *intended model* and *speaking a language* that I gave in Sects. 1 and 6. A more careful definition of intended model is: a model for an uninterpreted language  $\mathcal{L}$  is intended for a community  $C$  when the model assigns to the parts of the uninterpreted language  $\mathcal{L}$  those meanings had by those parts *on one of the languages* spoken by  $C$ . We can similarly revise the minimal claim of what it means to say a community uses a language: a community uses a language  $\mathcal{L}$  when the language assigns to each expression of the uninterpreted language a meaning had by an expression that is co-tokened with the expression of the uninterpreted language.

*Reply 2* Even if utterances and uninterpreted sentences are truth-apt, our practice of quotation-naming would be ambiguous between naming utterances, uninterpreted sentences or interpreted sentences. I've shown that the T-schema holds for the interpreted sentences and thus for one disambiguation of the truth predicate and quotation-name. That is sufficient to explain the truth and assertibility of instances of the T-schema.

### 8.3 Did she really say that?

*Objection.* Mary observes Hsiao utter (1):

(1) Bob is bald.

According to your view, Hsiao asserted many propositions. For example, she asserts the proposition that Bob has less than 10,000 hairs on his head. But plainly (cf. Schiffer (1998)) Mary cannot say any of the following:

- (13) a. # Hsiao said that Bob has less than 10,000 hairs on his head.  
 b. # Hsiao said that Bob has less than 10,001 hairs on his head.  
 c. # Hsiao said that Bob has less than 10,002 hairs on his head.

*Reply 1 Pragmatic Implicature* We can accept the truth of (13), while explaining the awkwardness of the speech report pragmatically. Consider scalar implicatures. Imagine a scenario in which I've seen exactly how many chairs are in the room. If I utter:

(14) There are three chairs in the room.

I implicate that there are *exactly* three chairs in the room. For if there were more than three chairs in the room, say exactly five chairs in the room, I would have made a strictly stronger assertion like:

(15) There are five chairs in the room.

This implicature follows from the Gricean maxim of maximal informativeness. Indeed, the implicature is so strong, meta-linguistic negation is sometimes warranted. If my friend walks into the room and observes exactly five chairs, it may be appropriate for him to utter:

(16) There aren't three chairs in the room – there are five.

In the case of Hsiao's assertion, the maxim of maximal informativeness is also violated. If Mary had uttered

(17) Hsiao said that Bob is bald.

She would have asserted strictly more propositions than if she had uttered (13a). In particular, she would have also asserted the propositions expressed by (13b) and (13c) (among others). The fact that she choose to perform the less-informative speech act implicates the false claim that Hsiao *only* said that Bob has less than 10,000 hairs on his head. This explains the unassertibility of (13a)–(13c). And, as with the chair example, in some cases the implicature may be strong enough to warrant meta-linguistic negation.

Indeed, if this pragmatic explanation were true, we would expect the unassertibility of (13a) to wane as the conversational purpose of the assertion (or the ‘question under discussion’) is made specific to particular precise propositions, so that there is less reason to be maximally informative. And this prediction seems to be satisfied. For instance, suppose that we’re in a context in which the term ‘tall’ definitely applies to a child if she is taller than 5’0” and definitely doesn’t apply if the child is shorter than 4’7”. (So, if a child is between 4’7” and 5’0”, it is indeterminate whether she is tall, and inappropriate to assert ‘she is tall’.) A parent is planning a trip to an amusement park for which, let’s say, one cannot enter unless they are at least 4’9”. Arvind, speaking to the parent, says:

(18) Regina is tall.

When asked whether Regina will be tall enough to ride the rides, the parent reports:

(19) Don’t worry: Arvind said that Regina is taller than 4’9”.

(19) Strikes me as perfectly assertible. It is certainly much *less* awkward than if the conversational purpose of the assertion weren’t so specific, which is exactly what the pragmatic story would predict. Indeed, the parent could follow up the report by canceling the implicature:

(20) I don’t mean to say that’s all Arvind said. He said more than just that: he said that she is tall.

*Reply 2 Neo-Fregean Propositions* A similar objection has been raised against supervaluationism by Schiffer (1998). He notes that, if there are only precise propositions (propositions like those picked out in the that-clauses of (13)), it must be vague which precise proposition is being picked out by the that-clause in (17). If (17) is true for the supervaluationist, it must be true on all ways of precisifying the that-clause and so (13a)–(13c) must each be true, which he takes to be absurd.

One reply given on behalf of the supervaluationist to Schiffer’s complaint comes from García-Carpintero (2000, 2009). His reply is detailed and subtle and I won’t rehash it here.<sup>37</sup> But its basic inspiration is the following: on at least one disambiguation of ‘proposition’ and ‘what is said’, propositions are partially individuated by Fregean senses or modes of presentation. And, because what was said by Hsiao when she uttered (1) has a different Fregean sense or mode of presentation than the propositions picked out by the that-clauses in (13), those indirect speech reports are false. ‘Bald’ has a vague Fregean sense.

By following the supervaluationist in adopting a neo-Fregean conception of propositions on which they are partially individuated by something like Fregean senses,

<sup>37</sup> The dialectic is more complicated than I am reproducing here. Schiffer also thinks that positing vague propositions cannot explain *de re* indirect speech reports in which singular terms scope over the that-clause. Much of García-Carpintero (2000, 2009) is dedicated to rebutting this claim. An alternative solution, given by Weatherson (2003) and Keefe (2009), posits a penumbral connection between the sentence (1) uttered by Hsiao and the that-clause in (17) such that, although determinately there is a unique precise proposition that is both expressed by Hsiao’s utterance of (1) and subsequently picked out by the that-clause in the indirect speech report, it is indeterminate which. That approach (i) cannot explain the assertibility of the negations of (13a)–(13c) and (ii) is effectively criticized by Rohrs (2017) for conflicting with the truism that a sentence adopts the truth-value of the proposition it expresses.

supersentialists can respond to the present objection. The propositions expressed by an utterance of “Bob is bald” are all distinct from the proposition that Bob has less than 10,000 hairs on his head: in every language being spoken ‘bald’ and ‘has less than 10,000 hairs on his head’ have different semantic contributions despite having the same intension, because they contribute different Fregean senses. The supersentialist only needs to claim that an utterance of the word ‘bald’ tokens many predicates, one of which shares an intension with a predicate like ‘has less than 10,000 hairs on his head’, another of which shares an intension with a predicate like ‘has less than 10,001 hairs on his head’, etc. But, the supersentialist need not claim that the predicates tokened by ‘bald’ have the same Fregean sense as the predicates tokened by ‘has less than 10,000 hairs on his head’.

*Reply 3 Psychological Theory of Counting* My third reply continues the theme of drawing lessons from the Solution by Plenitude from the Problem of the Many, and explains our reluctance to accept claims like (13) using a psychological theory of counting. Recall the lesson **Almost-Identity**: when we say ‘there is exactly one cat on the mat’, we are counting by almost-identity rather than strict identity. Counting by almost-identity comes with a linguistic dimension and a psychological dimension.

According to the linguistic dimension, my utterance

(21) There is exactly one cat on the mat.

can be adequately paraphrased as (something like)

(22) Every cat on the mat massively overlaps each other.

either because (21) is true in virtue of expressing a proposition with the truth-conditions of the proposition expressed by (22), or (21) is strictly speaking false but the availability of the nearby proposition expressed by (22) blunts the incredibility of the falsehood.

There is also a psychological dimension to counting by almost-identity. Plainly, when the folk look at Tibbles on a mat, they are not simultaneously representing the many cats on the mat. Rather, they are only representing a single cat on the mat and are ignoring the distinctions between the many cats. That’s why, when the differences between the many cats become relevant, we are initially struck with puzzlement and deny that there are many cats on the mat. When the differences between the cats are made salient, the folk form a more fine-grained representation of the situation. Representing the situation fine-grainly, we realize that we’ve been ignoring the many cats, and have been counting by almost-identity all along, despite thinking that we were counting by strict identity. And, that fine-grained representation is short-lived—outside of philosophical contexts, we switch back to the coarse-grained representation of a single cat on the mat.

Now consider how this psychological story applies to the supersentialist. When Mary says “Hsiao said that Bob has less than 10,000 hairs on his head” it naturally strikes us as very odd. That’s because we’ve been operating under a coarse-grained representation of the many propositions expressed, treating them as one. Because we’ve been representing all of these propositions as one, when the difference between the propositions are made salient by Mary’s utterance, we are struck with puzzlement. It’s only after philosophical reflection that we realize that we’ve been counting them by almost-identity all along. This provides us with a sort of ‘error theory’ for our hesitation to accept (13).



## 8.4 Higher order vagueness

*Objection* You've presented your view while ignoring higher-order vagueness. But the phenomenon of higher-order vagueness is real, and your view is incompatible with it. According to your view, when I make an utterance, I am in fact uttering many interpreted sentences. And, an utterance is permissible just in case it tokens only true interpreted sentences. A borderline sentence is one that is simultaneously uttered with other interpreted sentences of differing truth-values. But then consider the set of interpreted sentences  $S$  which are tokened by a given utterance. Either (1) all of the sentences are true, (2) all of the sentences are false, or (3) some are true and some are false. In the first and second cases, the sentences are not borderline. In the third case, the sentences are borderline. But then there is no room for higher-order vagueness: no scenario corresponds with *borderline* borderline sentences.

*Reply* I will not attempt to directly show how my account can give an adequate account of higher-order vagueness. I've been arguing that supersententialism is preferable to supervaluationism. So, I will be content to show that higher-order vagueness poses no *distinctive* problem for my account, as compared to the supervaluationist. There is a parallel complaint that can be leveled against the supervaluationist, and the same replies she gives to her analogous objection is available to me.

Compare the following analogous complaint, directed against the supervaluationist, with the complaint above:

Consider a sentence  $S$  and the set of admissible precisifications. Either (1)  $S$  is true on each of the admissible precisifications, (2)  $S$  is false on each of the admissible precisifications, or (3)  $S$  is true on some precisifications but false on other precisifications. In the first and second cases, the sentence is not borderline. In the third case, the sentence is borderline. But there is no room for higher-order vagueness: no scenario corresponds with *borderline* borderline sentences.

There are two replies that are given by the supervaluationist. The first is to deny the phenomenon of higher-order vagueness; the supersententialist can, of course, also deny higher-order vagueness.

The second response is to argue that the meta-language term 'admissible precisifications' is itself indeterminate.<sup>38</sup> Although there is one and exactly one set of admissible precisifications, it's indeterminate which set that is. In other words, which supervaluationist model is the intended one is indeterminate. Thus, the supervaluationist can explain cases of borderline borderline sentences as cases in which the sentence is borderline on one model  $\mathcal{M}_1$  and not borderline in  $\mathcal{M}_2$ , where  $\mathcal{M}_1$  and  $\mathcal{M}_2$  are not determinately not the intended model. A parallel reply is open to the supersententialist. The supersententialist can claim that the theorist's term 'languages being spoken' is itself indeterminate. Although there is one and exactly one set of languages being spoken, it's indeterminate which set that is.

Note the strategy for reply here is quite general and applies to a range of *indiscriminate* objections: objections to supersententialism that have analogues that apply to the

<sup>38</sup> See Keefe (2000, Chapter 8) for this response.

supervaluationist. For instance, according to the supersententialist, on every language being spoken the following sentence is true: ‘There is a precise number of hairs that separates the bald from the not bald’. But the objection does not discriminate—the supervaluationist will admit that the sentence is true on every precisification and is therefore true in English. Because I am attempting to establish that supersententialism is preferable to supervaluationism, objections like the foregoing that poses no distinctive problem for my account, as compared to the supervaluationist, do not undermine my main thesis.<sup>39</sup>

## 9 Conclusion

The goal of this paper was to show that the plurivaluationist strategy for explaining vagueness is preferable to the supervaluationist approach. I did this by offering a new talk-of-language account which, when paired with the plurivaluationist’s **Central Claim**, constituted a form of plurivaluationism that I call supersententialism. The supersententialist’s talk-of-language account was inspired by applying the Solution by Plenitude from the Problem of the Many to the ontology of linguistic objects; that account was then leveraged in order to unearth new advantages and defenses from objections for the supersententialist. The supersententialist’s main advantage over the supervaluationist is that, according to the supersententialist, all instances of the T-schema are true and assertible. According to the supervaluationist, some instances of the T-schema are not true and are therefore unassertible. And, although the supersententialist makes some seemingly incredible claims, I attempted to draw from the Problem in order to show that these claims are not as incredible as they may have seemed. And with these claims, I hope to have moved towards a more plausible theory of vagueness.

**Acknowledgements** Thanks to Mercedes Maria Corredor, Daniel Drucker, Jim Joyce, Jeremy Lent, Chip Sebens, Ted Sider, Eric Swanson, three anonymous referees and, especially, David Manley and Brian Weatherston for helpful conversations and comments. Thanks also to audiences at the University of Michigan, the University of San Diego, Colgate University, the 10th Annual Mark L. Shapiro Graduate Philosophy Conference (and my commentators Geoffrey Grossman and Yongming Han), and the 2018 Eastern APA (and my commentator Catarina Dutilh Novaes).

## References

- Andjelković, M., & Williamson, T. (2000). Truth, falsity, and borderline cases. *Philosophical Topics*, 28(1), 211–244.
- Burns, L. C. (1991). *Vagueness: An investigation into natural language and the sorites paradox*. Berlin: Springer.
- Dorr, C. (2014). Quantifier variance and the collapse theorems. *The Monist*, 97(4), 503–570.
- Dorr, C., & Hawthorne, J. (2014). Semantic plasticity and attitude reports. *The Philosophical Review*, 123(3), 281–338.
- Dutilh Novaes, C. (2008). Tarski’s hidden theory of meaning: Sentences say exactly one thing. In S. Rahman, T. Tulenheimo, & E. Genot (Eds.), *Unity, truth and the liar*, Chapter 3 (pp. 41–63). Berlin: Springer.

<sup>39</sup> This includes the main objection that is raised by Smith (2008) against classical plurivaluationism—that it violates the thesis he calls *Closeness*. As he argues, this complaint also applies to supervaluationism.

- Fara, D. G. (2000). Shifting sands: An interest relative theory of vagueness. *Philosophical Topics*, 28(1), 45–81.
- Field, H. (1994). Disquotational truth and factually defective discourse. *The Philosophical Review*, 103, 405–452. Reprinted in Field, H. (2001). *Truth and the absence of fact*. Oxford: Oxford University Press, 222–258.
- García-Carpintero, M. (2000). Vagueness and indirect discourse. *Philosophical Issues*, 10, 258–270.
- García-Carpintero, M. (2009). Supervaluationism and the report of vague contents. In R. Dietz & S. Moruzzi (Eds.), *Cuts and clouds*. Oxford: Oxford University Press.
- Jones, N. (2015). Multiple constitution. *Oxford Studies in Metaphysics*, 9, 217–261.
- Keefe, R. (1998). Vagueness and language clusters. *Australasian Journal of Philosophy*, 76(4), 611–620.
- Keefe, R. (2000). *Theories of vagueness*. Cambridge: Cambridge University Press.
- Keefe, R. (2009). Supervaluationism, indirect speech reports, and demonstratives. In R. Dietz & S. Moruzzi (Eds.), *Cuts and clouds*. Oxford: Oxford University Press.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge: Harvard University Press.
- Lewis, D. (1970). General semantics. *Synthese*, 22, 18–67. Reprinted with Postscript in Lewis, D. (1983). *Philosophical papers* (Vol. 1). Oxford: Oxford University Press, 189–232.
- Lewis, D. (1975). Languages and language. *Minnesota Studies in the Philosophy of Science*, 7, 3–35. Reprinted in Lewis, D. (1983). *Philosophical papers* (Vol. 1). Oxford: Oxford University Press, 163–188.
- Lewis, D. (1993). Many but almost one. In J. B. Keith Campbell & L. Reinhardt (Eds.), *Ontology, causality, and mind: Essays on the philosophy of D.M. Armstrong*. Cambridge: Cambridge University Press. Reprinted in Lewis, D. (1999). *Papers in metaphysics and epistemology*. Cambridge: Cambridge University Press, 164–182.
- Lowe, E. J. (1995). The problem of the many and the vagueness of constitution. *Analysis*, 55, 179–182.
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.
- Merricks, T. (2001). Varieties of vagueness. *Philosophy and Phenomenological Research*, 62(1), 145–157.
- Rohrs, B. (2017). Supervaluational propositional content. *Synthese*, 194(6), 2185–2201.
- Schiffer, S. (1998). Two issues of vagueness. *The Monist*, 81, 193–214.
- Smith, N. J. J. (2008). *Vagueness and degrees of truth*. Oxford: Oxford University Press.
- Smith, N. J. J. (2011). Fuzzy logic and higher-order vagueness. In P. Cintula, C. Fermüller, L. Godo, & P. Hájek (Eds.), *Understanding vagueness: Logical, philosophical and linguistic perspectives*. London: College Publications.
- Smith, N. J. J. (2015). Fuzzy logics in theories of vagueness. In P. Cintula, C. Fermüller, & C. Noguera (Eds.), *Handbook of mathematical fuzzy logic*. London: College Publications.
- Tye, M. (1996). Fuzzy realism and the problem of the many. *Philosophical Studies*, 81, 215–225.
- van Inwagen, P. (1990). *Material beings*. Ithaca: Cornell University Press.
- Varzi, A. C. (2007). Supervaluationism and its logics. *Mind*, 116(463), 633–675.
- Weatherston, B. (2003). Many many problems. *Philosophical Quarterly*, 53(213), 481–501.
- Williams, J. R. G. (2006). An argument for the many. *Proceedings of the Aristotelian Society*, 106, 411–419.
- Williams, J. R. G. (2008). Supervaluationism and logical revisionism. *The Journal of Philosophy*, 105(4), 192–212.
- Williamson, T. (1994a). Definiteness and knowability. *The Southern Journal of Philosophy*, 33, 171–191.
- Williamson, T. (1994b). *Vagueness*. London: Routledge.