

# *Truthier Than Thou: Truth, Supertruth and Probability of Truth*

NICHOLAS J.J. SMITH  
University of Sydney

## Abstract

Different formal tools are useful for different purposes. For example, when it comes to modelling degrees of belief, probability theory is a better tool than classical logic; when it comes to modelling the truth of mathematical claims, classical logic is a better tool than probability theory. In this paper I focus on a widely used formal tool and argue that it does not provide a good model of a phenomenon of which many think it does provide a good model: I shall argue that while supervaluationism may provide a model of probability of truth, or of assertability, it cannot provide a good model of truth—supertruth cannot be truth. The core of the argument is that an adequate model of truth must render certain connectives truth-functional (at least in certain circumstances)—and supervaluationism does not do so (in those circumstances).

## 1. Introduction

Different formal tools are useful for different purposes. For example, when it comes to modelling degrees of belief, probability theory is a better tool than classical logic; when it comes to modelling the truth of mathematical claims, classical logic is a better tool than probability theory. In this paper I focus on a widely used formal tool and argue that it does *not* provide a good model of a phenomenon of which many think it *does* provide a good model. I shall argue that supervaluationism cannot provide a good model of truth: supertruth cannot be truth. When I need specific examples I shall generally turn to the use of supervaluationism in connection with vagueness—but the argument against supervaluationism as a model of truth is quite general: it applies also to other uses of supervaluationism (as a model of truth), for example to model the truth conditions of statements about the future.<sup>1</sup> The argument applies to standard supervaluationism, to the degree-theoretic version and to the more recent conceptual spaces version; it does not apply to pluralvaluationism.<sup>2</sup> The theories just mentioned (pluralvaluationism, and varieties of supervaluationism) will be introduced in §2. In §3 I present the argument against supervaluationism as a model of truth. The core of the argument is that an adequate model of truth must make certain connectives truth-functional (at least in certain circumstances)—and supervaluationism does not do so (in those circumstances). In §4 I consider lessons to be drawn from the argument and make it clear that while supervaluationism

cannot be a good model of truth (or more precisely: truth cannot be supertruth), it may yet be a good model of some other significant property—for example, assertability.

## 2. Supervaluationisms and Plurivaluationism

In this section I sketch some formal theories.<sup>3</sup> Throughout we consider a standard first order language  $\mathcal{L}$  with individual constants  $a, b, c, \dots$  and predicates  $P, Q, R, \dots$  of each arity. Analogues of everything to be said about one-place predicates can be said about many-place predicates and analogues of everything to be said about disjunction and conjunction can be said about existential and universal quantification. Hence, for the sake of an uncluttered presentation in which the central points are not obscured by unnecessary detail, we shall only explicitly mention connectives and one-place predicates.

A *classical valuation*  $\mathfrak{V}_c$  of  $\mathcal{L}$  comprises:

- a set  $D$  (the *domain*)
- (where  $\mathcal{I}$  is the set of individual constants of  $\mathcal{L}$ ):  
a function from  $\mathcal{I}$  to  $D$  (assigning a *referent* to each individual constant)
- (where  $\mathcal{P}$  is the set of one-place predicates of  $\mathcal{L}$ ):  
a function from  $\mathcal{P}$  to  $\{0, 1\}^D$  (assigning an *extension* to each predicate: an extension is a total function from the domain to  $\{0, 1\}$ ; objects sent to 1 are in the extension and objects sent to 0 are not in the extension).

A classical valuation can be extended to a classical *model*  $\mathfrak{M}$  using the standard classical rules. In particular:

- the truth value of an atomic wff  $Pa$  is the value (1 or 0) to which the extension of  $P$  sends the referent of  $a$
- the truth values of negations, conjunctions, disjunctions and so on are determined by the classical truth tables.

A classical model assigns a truth value (1 or 0) to each closed wff of  $\mathcal{L}$ .

In *plurivaluationism*, a language  $\mathcal{L}$  is associated with a nonempty set of classical models. We call the members of this set the *acceptable models*.

A *partial valuation*  $\mathfrak{V}_p$  of  $\mathcal{L}$  is just like a classical valuation except that the extension of a predicate can be a *partial* function from the domain to  $\{0, 1\}$  (objects sent to 1 are in the extension; objects sent to 0 are outside the extension; objects sent nowhere are neither in nor out).

A *three-valued valuation* is just like a classical valuation except that the extension of a predicate is a total function from the domain to  $\{0, *, 1\}$ , where  $*$  is a third truth value in addition to 1 and 0. Everything that we say below in terms of partial valuations can be reformulated in terms of three-valued valuations. I shall use the term ‘tripartite’ when I wish to speak generally of three-valued and partial two-valued setups.

A classical valuation  $\mathfrak{V}_c$  *extends* a partial valuation  $\mathfrak{V}_p$  if:

- $\mathfrak{V}_p$  and  $\mathfrak{V}_c$  have the same domain
- $\mathfrak{V}_p$  and  $\mathfrak{V}_c$  assign the same referents to names
- for all predicates  $P$  and for all  $x$  in the domain (and where  $P_p$  is the extension of  $P$  on  $\mathfrak{V}_p$  and  $P_c$  is the extension of  $P$  on  $\mathfrak{V}_c$ ):

-- if  $P_p(x) = 1$  then  $P_c(x) = 1$

-- if  $P_p(x) = 0$  then  $P_c(x) = 0$

(i.e. the extension of  $P$  on  $\mathfrak{V}_c$  just closes the gaps in the extension of  $P$  on  $\mathfrak{V}_p$ : it does not move anything from in (1) to out (0) or vice versa).

We call a classical model an *extension* of  $\mathfrak{V}_p$  if it is determined (using the standard classical rules) by a classical valuation  $\mathfrak{V}_c$  that extends  $\mathfrak{V}_p$ .

There are various ways of extending a partial valuation to a model, which assigns a truth value or a gap (i.e. no value) to each closed wff of  $\mathcal{L}$ . One way—call it the *recursive* way—mirrors the classical story as closely as possible:

- the truth value of an atomic wff  $Pa$  is the value (1, 0 or no value) to which the extension of  $P$  sends the referent of  $a$
- the classical truth tables are extended to cover the cases where one or more component proposition has no value. Perhaps the best known option here is Kleene's strong tables—but there are many other options.<sup>4</sup>

Another way is the supervaluationist route. Given a partial valuation  $\mathfrak{V}_p$ , the associated *supervaluation* is a partial function  $v$  from closed wffs  $\alpha$  to  $\{0, 1\}$  defined as follows:

$$v(\alpha) = \begin{cases} 1 & \text{if } \alpha \text{ has the value 1 on every extension of } \mathfrak{V}_p \\ 0 & \text{if } \alpha \text{ has the value 0 on every extension of } \mathfrak{V}_p \\ \text{undefined} & \text{otherwise} \end{cases}$$

Note that 'otherwise' here means that  $\alpha$  has the value 1 on some extension of  $\mathfrak{V}_p$  and has the value 0 on some (other) extension of  $\mathfrak{V}_p$ . The property of having the value 1 (true) on every extension is often referred to as 'supertruth' and so because the supervaluation assigns the value 1 (true) to sentences that are supertrue, supervaluationism is often summed up in the slogan: truth is supertruth.

Consider atomic wffs such as  $Pa$  for a moment. They can be assigned a value by the supervaluation in the way just specified—or alternatively (as in the recursive way of extending a partial valuation to a model) they can be assigned a value with reference only to the partial valuation (not its extensions): the truth value of an atomic wff  $Pa$  is the value (1, 0 or no value) to which the extension of  $P$  sends the referent of  $a$ . Both options yield the same values for atomic wffs.

One can refine the supervaluationist view by considering not *all* extensions of  $\mathfrak{V}_p$  but only those satisfying certain conditions. These extensions are then called the *admissible extensions* and in the definition of the supervaluation  $v$ , 'extension' is replaced by 'admissible extension'. We shall use  $E(\mathfrak{V}_p)$  to denote the set of all

admissible extensions of  $\mathfrak{A}_p$ . The unrefined version of supervaluationism presented above is the special case of the refined version where every extension is admissible.

An important variant of the supervaluationist framework is the degree-theoretic form of supervaluationism. Here we suppose there to be a normalised measure function  $\mu$  on the powerset of  $E(\mathfrak{A}_p)$ . Where  $A$  is the set of admissible extensions on which  $\alpha$  has the value 1, the supervaluation  $v$  is then defined thus:

$$v(\alpha) = \mu(A)$$

It may be possible to define  $\mu$  only on some  $\sigma$ -field of subsets of  $E(\mathfrak{A}_p)$ , not on the full powerset. In that case  $v$  is defined only for  $\alpha$  such that  $A$  is a measurable set. Decock and Douven have recently proposed a particular implementation of the degree-theoretic form of supervaluationism in the framework of conceptual spaces.<sup>5</sup> They show that the values assigned by the supervaluation  $v$  in their account behave in the same ways as the *verities* of Edgington (1997) and are formally probabilities.

### 3. Truth and Truth-Functionality

I shall now argue that the formal framework of supervaluations cannot provide a good model of truth (or more precisely: truth cannot be supertruth).<sup>6</sup> Of course in order to argue this I shall have to assume something about truth—something that a good model of truth must capture. I shall assume the classic, orthodox view that *truth is saying it how it is*. A claim is true if things are the way it claims them to be; it is false if things are not as it claims them to be. This idea goes back at least as far as Plato and Aristotle:

SOCRATES: But how about truth, then? You would acknowledge that there is in words a true and a false?

HERMOGENES: Certainly.

SOCRATES: And there are true and false propositions?

HERMOGENES: To be sure.

SOCRATES: And a true proposition says that which is, and a false proposition says that which is not?

HERMOGENES: Yes, what other answer is possible? (Plato, c.360 BC.)

... we define what the true and the false are. To say of what is that it is not, or of what is not that it is, is false, while to say of what is that it is, and of what is not that it is not, is true (Aristotle, c.350 BC., Book IV ( $\Gamma$ ) §7)

This is just a rough guiding idea about truth—and that is all I am assuming. The rough idea has been used to motivate more precise, detailed theories of truth—some of which (such as certain versions of the correspondence theory of truth) are quite contentious. I am not assuming any precise, detailed theory of truth—I am assuming only the basic, guiding idea that truth is saying it how it is. Indeed, the whole point will be to assess precise models of truth (such as supervaluationism) according to whether or not they can accommodate this guiding idea. (The lack

of precision here makes my argument more forceful: the conclusion is not that supervenience conflicts with some very specific definition of truth—which one could then simply reject; it is that supervenience conflicts with the very idea of truth as saying it how it is.)

Implicit in this basic view of truth is the idea that each claim is *about* something—it has a *subject matter*. This subject matter is what must be the way the claim says it is in order for the claim to be true. Again, this basic idea has been spelled out more precisely in various different ways. For example, one branch of this line of thought leads to truthmaker theories; another branch leads to theories of semantic groundedness and grounded truth.<sup>7</sup> As far as the basic idea is concerned, it could be that all claims have the same subject matter: ‘the world’. However one feature of many of the more precise developments of the basic idea is the view that at least some truth-apt sentences have a *minimal* subject matter, such that how things are with this subject matter fixes whether the sentence is true *and vice versa*. Thus, if we know the sentence is true, we know *all about* how things are with the (minimal) subject matter of the sentence. To put it differently, the minimal subject matter of a sentence comprises just what *must* be fixed in order to settle the truth value of the sentence. So, for example, in a sense ‘John is tall’ is about John. But John is not the minimal subject matter of the sentence. When we know the claim is true, we do not know all about John: we do not know his weight, his occupation, his age, his marital status, etc. These things *need not* be fixed in order to settle the truth value of ‘John is tall’. The *minimal* subject matter of the claim that John is tall might instead be something like ‘how John stands with respect to the property of being tall’ or ‘John’s tallness or lack thereof’.

As anyone familiar with the literature on truthmaking will quickly be able to see, it is far from clear that *every* truth-apt sentence has a minimal subject matter. For example, just consider a disjunction such as ‘John is tall or Bill is bald’ or an existential claim such as ‘There are at least two places to get good coffee in Wellington’. For the argument to be given below to go through, it is not required that every proposition has a minimal subject matter. What will be assumed is that every proposition has a subject matter (remember that this is implicit in the basic guiding idea about truth) and that some propositions have minimal subject matters. In particular, I think it is extremely plausible that atomic propositions formed from a one-place predicate and a name (for example, ‘John is tall’) have minimal subject matters.

Now for a definition. Let us say that a connective  $*$  is *conservative* if the subject matter of  $\alpha * \beta$  is exhausted by the subject matters of  $\alpha$  and  $\beta$ .<sup>8</sup> In other words, what  $\alpha * \beta$  is about does not extend beyond the combination of what  $\alpha$  is about and what  $\beta$  is about. Of course this definition cannot be precise, because we are not working with any precise account of what subject matters are and of whether, for example, they combine set-theoretically or mereologically or in some other way. I trust however that at the level of basic, guiding ideas at which we are working, the idea is clear. Here’s another way of putting it: what you would have to investigate to determine whether  $\alpha * \beta$  is true does not extend beyond what you would have to investigate to determine whether  $\alpha$  is true together with what you would have to investigate to determine whether  $\beta$  is true.

Let's say that a connective  $*$  is *minimally conservative* if it is conservative in cases where the component propositions have minimal subject matters: that is, when  $\alpha$  and  $\beta$  have *minimal* subject matters, the subject matter of  $\alpha * \beta$  is exhausted by the subject matters of  $\alpha$  and  $\beta$ . Of course a conservative connective will be minimally conservative, but the converse need not hold.

Where  $\alpha$  is a claim or truth-apt sentence, let  $\langle \alpha \rangle$  denote  $\alpha$ 's subject matter and let  $[\alpha]$  denote  $\alpha$ 's truth value. We can now show that, where  $*$  is a minimally conservative connective and  $\alpha$  and  $\beta$  have minimal subject matters,<sup>9</sup> fixing  $[\alpha]$  and  $[\beta]$  fixes  $[\alpha * \beta]$ —and hence minimally conservative connectives must behave truth-functionally when the component propositions have minimal subject matters:

1. If you fix  $[\alpha]$  then (by the basic idea about truth, combined with the fact that  $\alpha$ 's subject matter is minimal) you fix how things are with  $\langle \alpha \rangle$  (i.e.  $\langle \alpha \rangle$  is, or is not, *the way  $\alpha$  says it is*;  $\langle \alpha \rangle$  is exactly the way it must be in order for  $[\alpha]$  to be what it is). Likewise if you fix  $[\beta]$  then you fix how things are with  $\langle \beta \rangle$ .
2. If you fix how things are with  $\langle \alpha \rangle$  and  $\langle \beta \rangle$  then (because  $*$  is conservative—or at least because  $*$  is minimally conservative and  $\alpha$  and  $\beta$  have minimal subject matters) you fix how things are with  $\langle \alpha * \beta \rangle$ .
3. If you fix how things are with  $\langle \alpha * \beta \rangle$  then (by the basic idea about truth) you fix  $[\alpha * \beta]$ .
4. Thus (from 1–3) if you fix  $[\alpha]$  and  $[\beta]$  then you fix  $[\alpha * \beta]$ . Hence  $*$  is truth-functional (at least in cases where the component propositions have minimal subject matters).<sup>10</sup>

Finally, it is overwhelmingly plausible that conjunction and disjunction are (at least minimally) conservative connectives. 'Bill is bald and Ben is tall' isn't about anything that neither 'Bill is bald' nor 'Ben is tall' is about; similarly for 'Bill is bald or Ben is tall'. Of course the 'and' statement and the 'or' statement make different claims concerning *how things are* with their subject matter: they have different truth conditions. The point is simply that the subject matter of the conjunction does not extend beyond the combined subject matters of its conjuncts; and likewise for the disjunction and its disjuncts.<sup>11</sup>

In light of the foregoing, this means that conjunction and disjunction must be truth-functional (in certain circumstances). The supervaluationist treatment of conjunction and disjunction, however, is not truth-functional (in those circumstances). This point applies to all forms of supervaluationism mentioned in §2. (As we shall discuss further in §4, it does not apply to plurivaluationism, in which the treatment of conjunction and disjunction is truth-functional.) Here we illustrate the point using the version of supervaluationism in which the supervaluation is defined by reference to the admissible extensions of an initial partial valuation. Suppose that on the partial valuation, the referent of  $a$  is sent nowhere by the extensions of  $P$  and  $R$  and the referent of  $b$  is sent nowhere by the extension of  $Q$ . Hence  $Pa$ ,  $Ra$  and  $Qb$  lack a truth value. Suppose also that on every admissible extension of the partial valuation, the extensions of  $P$  and  $R$  are mutually exclusive and jointly exhaustive, while there are no particular constraints on the relationship between the extensions

of  $P$  and  $Q$ . Hence  $Pa \vee Ra$  will be true on every admissible extension and hence true, while  $Pa \vee Qb$  will be true on some admissible extensions and false on others and hence will lack a truth value. But  $Pa \vee Ra$  and  $Pa \vee Qb$  are both disjunctions where the disjuncts both lack truth values—so  $\vee$  is not truth-functional (if it were, both disjunctions would get the same value or lack thereof). Similar remarks apply to conjunction.  $Pa \wedge Ra$  will be false on every admissible extension and hence false, while  $Pa \wedge Qb$  will be true on some admissible extensions and false on others and hence will lack a truth value. But  $Pa \wedge Ra$  and  $Pa \wedge Qb$  are both conjunctions where the conjuncts both lack truth values—so  $\wedge$  is not truth-functional (if it were, both conjunctions would get the same value or lack thereof).

Let's take stock. I have presented an argument that shows that where  $*$  is a connective that is conservative or at least minimally conservative,  $*$  must behave truth-functionally—at least when the components of the compound propositions formed using  $*$  have minimal subject matters. Now it is overwhelmingly plausible that one-place atomic predications have minimal subject matters and that conjunction and disjunction are conservative (if not in general then at the very least when the disjuncts and conjuncts are one-place atomic predications). Hence, truth values—values that represent presence, absence or quantity of truth—must distribute value-functionally across one-place predications and conjunctions and disjunctions thereof. But the values assigned by supervaluations do not distribute value-functionally across propositions of these kinds. Hence the values assigned by the supervaluation are not truth values: supervaluationism does not provide an adequate model of truth.

To understand the abstract form of my position here, consider an analogy. Values that are supposed to represent *lengths* must distribute themselves across objects in certain ways. For example, if this pencil can be placed alongside this pen in such a way that their ends coincide, then they must be assigned the same length value; if, placed end-to-end, they reach exactly from one end of the ruler to the other, then the length value of the ruler must be twice that of the pen (which we are supposing is the same as the length value of the pencil); and so on. If we assign values to objects in some way that does not respect these sorts of constraints, then these values cannot be modelling lengths. They might be useful for some other purpose—say, modelling the values placed on objects by some agent, or modelling weights—but they cannot be length values. Now I am making a claim of a similar sort about truth values. How true a sentence is varies with how some bit of the world is (this is the basic guiding idea about truth, and the relevant bit of the world is the sentence's subject matter). Which bit of the world is relevant to the truth of a conjunction (disjunction) is related in certain ways to which bits of the world are relevant to the truth of its conjuncts (disjuncts). Together, these facts induce a constraint on truth values: that is, on values representing the presence, absence or quantity of truth. The truth values of conjunctions (disjunctions) cannot vary just any old how with respect to the truth values of their conjuncts (disjuncts): a certain kind of relationship is required. The values assigned by supervaluations do not, in general, meet this constraint. Hence the values assigned by supervaluations are not truth values: they are not values that track presence, absence or quantity of truth.

Before closing this section—and in order to understand the argument presented here more fully—let’s see what the argument does *not* show. First, it does not show that all connectives must be truth-functional. Second, it does not show that conjunction and disjunction (and other conservative connectives) must be value-functional whatever the values assigned to sentences are supposed to represent: the argument applies only to truth values. We discuss these points in turn.

First, consider a paradigmatic non-truth-functional connective: ‘necessarily’ (or ‘It is necessarily the case that’ or  $\Box$ ). The argument says nothing about  $\Box$  because  $\Box$  is not conservative. On any reasonable understanding of modality, ‘Necessarily, Helen Clark is Prime Minister of New Zealand in 2000’ has a broader subject matter than ‘Helen Clark is Prime Minister of New Zealand in 2000’. (For example: the latter says something about the state of the actual world at a particular time whereas the former says something about the states of all possible worlds at a particular time. Or: the latter says something about *whether* Helen Clark possesses a certain property while the former says something about *how* she possesses it, e.g. essentially.) Thus there is no reason why fixing the truth value of the latter should fix the truth value of the former. If we tried to show this using the above argument, step 2 would fail: fixing how things are with  $\langle \alpha \rangle$  does not mean fixing how things are with  $\langle \Box \alpha \rangle$ .

Second, consider a paradigmatic case where we assign values to sentences in such a way that the value assigned to  $\alpha \wedge \beta$  (and likewise for  $\alpha \vee \beta$ ) is not a function of the values assigned to  $\alpha$  and  $\beta$ : probability theory. Conjunction and disjunction are conservative connectives: the subject matter of  $\alpha \wedge \beta$  is exhausted by the subject matters of  $\alpha$  and  $\beta$ , and likewise for  $\alpha \vee \beta$ . However, the assignment of a probability value to a sentence—unlike the assignment of a truth value—is *not* determined by how things are with the subject matter of the sentence. ‘The die comes up even’ is about the state of the die at some particular time. How true this sentence is is determined by how things are with its subject matter (this is part of the basic guiding idea about truth). But how probable it is is not so determined: to fix a probability, we need to look not just at the actual state of the die at the time in question but at all its possible states (and in particular at the distribution among them of states in which the die shows an even number). Conversely, fixing the probability value of a sentence does not fix how things are with its subject matter. Consider the claim ‘The coin comes up heads’. It is about the state of the coin at a certain time. I can know that the probability of the claim is 0.5 without knowing how things are with the subject matter of the claim (i.e. whether the coin comes up heads or not). Thus, considering the argument above, if we took  $\langle \alpha \rangle$  to be the probability value of  $\alpha$  (as opposed to its truth value), steps 1 and 3 would fail. The basic idea of truth is that how things are with the subject matter of a sentence fixes its truth value and (when subject matters are minimal) vice versa. There is thus a two-way interchange between the state of the subject matter of a sentence and the truth state of the sentence. The argument does not go through for values—for example, probability values—that lack this feature: values of sentences that are not determined by, or do not determine, how things are with the subject matter of the sentence.



Summing up: In the case of  $\Box$ , the truth value of  $\Box\alpha$  is indeed determined by how things are with the subject matter of this sentence. However this subject matter extends beyond the subject matter of  $\alpha$  (i.e.  $\Box$  is not conservative) and so we cannot argue that  $\Box$  should be truth-functional. In the case of probability values, the subject matter of  $\alpha \wedge \beta$  does not extend beyond the combined subject matters of  $\alpha$  and  $\beta$  (and likewise for  $\alpha \vee \beta$ ), but probability values are *not* determined by subject matter in the way that truth values are and so again we cannot argue that conjunction and disjunction must be probability-functional. The argument goes through only when we are talking about conservative connectives and only when we are talking about truth.<sup>12</sup> Thus the argument strikes against the use of the supervaluationist framework as a model of truth, without striking against all frameworks in which there are non-truth-functional connectives and without striking against frameworks in which values that do not represent truth are assigned to conjunctions and disjunctions in a non-value-functional way.<sup>13</sup>

#### 4. Ways Forward

I have argued that the supervaluationist framework cannot provide a good model of truth. So can it provide a good model of something else—or should the framework be abandoned entirely? There is one thing of which the supervaluationist framework certainly provides a good model (although how useful it is to have a model of this thing is a further question that still remains open)—and one thing of which it arguably provides a good model.<sup>14</sup>

##### 4.1 Supervaluationism without Truth I: Probability of Truth

The thing that supervaluationist frameworks definitely model well is *probability of truth under precisification*.<sup>15</sup> Suppose we were to precisify our vague language entirely: make it so that each predicate corresponds to a unique crisp set (its extension) and in general there is a unique classical model that correctly represents the semantics of the (fully precisified) language. There are many legitimate ways of doing this—yet it is not the case that absolutely anything goes. For example:

- Suppose that  $x$  is a borderline case of both ‘red’ and ‘orange’. It is legitimate to put  $x$  in the extension of ‘red’ or to leave it out; similarly for ‘orange’. But if we precisify ‘red’ in such a way that  $x$  comes out as red then we must not precisify ‘orange’ in such a way that  $x$  comes out as orange.
- Bill, who is borderline tall, could legitimately be put in the extension of ‘tall’ or left out; similarly for Ben, who is just a tad taller than Bill; but it is not legitimate to put Bill in and leave Ben out.
- Suppose that Bill is 16, Ben is 18, and Bob is 20. If we precisify ‘juvenile’ and ‘adult’, it must turn out that each of Bill, Ben and Bob falls in exactly one of these categories, and it must not be the case that Bill and Bob fall in one category, while Ben falls in the other.

The first example concerns how one object should be classified relative to several predicates; the second example concerns how several objects should be classified relative to one predicate; and the third example concerns how several objects should

be classified relative to several predicates. This is just a tiny sample of the constraints on legitimate precisification—but it should be enough to give the idea. Now suppose we have a partial valuation that corresponds to our actual use of vague language. Persons generally agreed to be tall are sent to 1 by the extension of ‘tall’; persons generally agreed not to be tall are sent to 0; and the borderline cases are sent nowhere (and similarly for all other vague predicates). An *admissible* extension of this partial valuation will be a classical model that corresponds to one of the *legitimate* ways of precisifying the entire language. And now suppose we ask of a given sentence  $\alpha$ : how likely is it that  $\alpha$  would be true, were the entire language precisified? The value assigned to  $\alpha$  by the supervaluation in the degree-theoretic supervaluationist framework directly answers this question. In the non-degree-theoretic form of supervaluationism we get a more coarse-grained answer: 1 if  $\alpha$  would definitely be true; 0 if  $\alpha$  would definitely be false; and no answer (or a third value) in all other cases (i.e.  $\alpha$  has some positive probability of truth under complete precisification and some positive probability of falsehood).

The question still remains whether it is in any way useful to have a model of probability of (classical) truth under complete precisification of the language—we shall say a little more on this below. However at this point we can at least say: from a conceptual point of view, there is nothing wrong with viewing the supervaluationist framework as a model of probability of truth were the language completely precisified. The degree-theoretic version models this directly; the tripartite version provides a coarser model in which non-extreme probabilities are collapsed to a single middle value. Problems arise only if we try to view *probability* of truth (were the language precisified) as *quantity* of truth (now). Compare the famous Problem of Points, where two players need to decide how to divide the prize money—all of which would have been won by one of them, had the game proceeded to completion—given that the game was called off prematurely. In this context it is perfectly reasonable to move from a probability of winning all the money to a fair share of the money now. My point at present is that the corresponding move is not reasonable in the case of truth. Probability of truth cannot be regarded as quantity of truth, because values representing probabilities of truth are assigned to conjunctions and disjunctions in a non-value-functional way, whereas genuine truth values must be assigned to conjunctions and disjunctions in a value-functional way.<sup>16,17</sup>

#### 4.2 Supervaluationism without Truth II: Assertability

We turn now to the second thing of which the supervaluationist framework might provide a model. From Fine and Kamp in the 70’s, through Osherson and Smith in the 80’s, Kamp and Partee in the 90’s and Keefe in the 00’s, up to Sauerland in the present decade, the case for supervaluationist treatments of vagueness has always rested heavily on intuitions about the use of compound statements in the presence of borderline cases.<sup>18</sup> For example, suppose that a certain blob is on the border of pink and red and let  $P$  be the sentence ‘the blob is pink’ and  $R$  be the sentence ‘the blob is red’—so  $P$  and  $R$  are neither clearly true nor clearly false. Fine (1975, 269–70) thinks that  $P \vee R$  is clearly true and that  $P \wedge R$  is clearly false. On a related note, Osherson and Smith (1981, 45–6) think that where  $Ax$  means that  $x$  is an apple,  $Aa \wedge \neg Aa$

should be true to degree 0 and  $Aa \vee \neg Aa$  should be true to degree 1, whatever  $a$  is.<sup>19</sup> The arguments are often stated in terms of *truth* (e.g. such and such statements are clearly true, such and such other statements are clearly false, and so on) but the data on which they rest concern what speakers find it natural to say. Now I think it is far from clear that the data really line up in the way that supervaluationists claim.<sup>20</sup> But suppose (for the sake of argument) that they do. Then one thing we could take the supervaluationist framework to model is the *assertability* of vague sentences. In the degree-theoretic version, assertability would be a graded matter; in the non-degree-theoretic versions, assertability would be a tripartite matter (assert, deny or hedge). There could even be a meaningful link with the idea that supervaluationism models probability of truth under precisification. Rather than it simply being a brute fact that assertability behaves in certain ways, it could be argued that assertability goes by probability of truth under precisification. The situation would then be rather similar to a familiar picture of conditionals: their assertability conditions *but not their truth conditions* are specified in probabilistic terms.<sup>21</sup>

Of course, on such a view we would still need a separate account of the *truth conditions* of vague sentences—and, given the argument of §3, that is something the supervaluationist framework cannot provide. For example, the truth conditions of sentences on the partial valuation could be derived by adding a truth-functional tripartite logic—what we called in §2 the recursive way of extending a partial valuation to a model. Thus we would add two separate pieces of machinery to the partial valuation: a recursive apparatus for assigning truth values; and an apparatus of classical extensions and a supervaluation (either tripartite or degree-theoretic) for assigning assertability/probability values. A variation on this approach would be to replace the tripartite valuation with a fuzzy valuation, and again add two separate pieces of machinery to it: a truth-functional fuzzy logic to assign truth values to sentences;<sup>22</sup> and an apparatus of classical extensions and a supervaluation (either tripartite or degree-theoretic) for assigning assertability/probability values.<sup>23</sup>

### 4.3 Plurivaluationism as a Model of Truth

Nothing that we have said counts against taking plurivaluationism as a model of truth—for in the plurivaluationist framework, conjunction and disjunction are truth-functional. In plurivaluationism, there is nothing like a supervaluation: an assignment of further values to sentences, based on but (in general) distinct from the values they are assigned within classical models. There are only the classical models. The only places that truth values are assigned are within these models—and in classical models, conjunction and disjunction are truth-functional.

Of course, plurivaluationism lends itself very naturally to a non-value-functional story about *assertability*—but as we have seen, there is no reason why conjunction and disjunction should behave in a value-functional way when the values assigned to sentences represent assertability (not truth). A natural story about assertability within the plurivaluationist framework goes as follows. Plurivaluationism models semantic indeterminacy or plurality. When I utter a sentence—say, ‘Bob is bald’—I speak relative to multiple classical models simultaneously. Relative to each model taken individually, I make a particular claim: that *this guy* (the referent of ‘Bob’ on

that model) is a member of *this set* (the extension of ‘is bald’ on that model). But overall, I do not make a single claim: we can say that it is indeterminate which claim I make (semantic indeterminacy); or we can say that I make all of them at once (semantic plurality)—both are reasonable glosses on the situation. Now it is natural to say that if *all* the claims I make are true (each on its respective model) then I can simply assert ‘Bob is bald’ without hesitation: the indeterminacy or plurality doesn’t matter. Likewise if all the claims I make are false: in that case I can simply deny ‘Bob is bald’ without hesitation. But if some of the claims are true and some are false, then neither outright assertion nor outright denial seems appropriate: instead I should hedge. On a more subtle version of the story, instead of a catch-all hedging response, we have appropriate degrees of assertion: the greater the proportion of true claims (relative to the total number of claims) the more confident my assertion should be. Now to model this story about assertability in a formal way, we can add a supervaluation to the classical models (but note that there is still no partial model in the picture, so formally the account is still not the same as supervaluationism). On the first version of the story, the supervaluation assigns one of three values; on the more subtle version of the story, the supervaluation assigns a continuum of values, as in the degree-theoretic form of supervaluationism. In both cases, the values assigned represent quantities of assertability (not truth), and so there is no problem in the fact that the assignments will not be value-functional.<sup>24</sup>

This is good news for many who think of themselves as ‘supervaluationists’: for really, the formal framework that models their semantic picture is (in the terms of this paper) plurivaluationism, not supervaluationism. I am thinking of those who find the following kinds of remark congenial:

I regard vagueness as semantic indecision: where we speak vaguely, we have not troubled to settle which of some range of precise meanings our words are meant to express. (Lewis, 1986a, 244, n.32)<sup>25</sup>

Broadly speaking, supervaluationism tells us two things. The first is that the semantics of our language is not fully determinate, and that statements in this language are open to a variety of interpretations each of which is compatible with our ordinary linguistic practices. The second thing is that when the multiplicity of interpretations turns out to be irrelevant, we should ignore it. If what we say is true under all the admissible interpretations of our words, then there is no need to bother being more precise. (Varzi, 2003, 14)

The picture here is one in which a vague predicate such as ‘tall’ is not associated with one inherently vague property: it is associated with many inherently precise properties. In the plurivaluationist framework, this is modelled by associating a vague language with many classical models: the vague predicate is then associated with many (different) precise extensions, one in each model. That exhausts the semantic story. There is no partial valuation here, and no supervaluation. The machinery is all perfectly classical. It is just that we consider many classical models of the language in parallel, rather than considering only a single classical model of the language.

Of course, as already noted, we can take a step back and consider all the classical models side by side, so to speak, and say things such as ‘If your claim is true on all of the classical models, just proceed as if you are making a single determinate true claim’. This further story can be formally modelled by bolting a supervaluation onto the classical models. But this further story is not semantic: it is a story about assertability. The values issued by the bolt-on supervaluation are assertability values, not truth values.

In sum, either we model this picture of how vague language works (i.e. the picture expressed in the quotations above from Lewis and Varzi—a picture which many seem to find congenial) using only the plurivaluationist machinery, in which case we have only classical models, and no values that are assigned in a non-value-functional way; or we also bolt on a supervaluation (although still there is—as noted—no partial valuation in the picture, so formally this is still not the same as the supervaluationist framework), but the supervaluation assigns values representing assertability, not truth. Thus, either there are *no* values that are assigned in a non-value-functional way, or there are values representing something other than truth that are assigned in a non-value-functional way—but on no way of spelling out the picture are there *truth* values that are assigned in a non-value-functional way. Either the formal tool of a supervaluation is absent—or the values it assigns are taken to represent quantities of assertability, not of truth.<sup>26</sup>

#### 4.4 *Supervaluationism without “Truth = Supertruth”*

Finally, let’s consider how one might try to retain the supervaluationist machinery as a model of truth (not merely as a model of probability of truth under precisification, or as a model of assertability).

Of course, one way in which a supervaluationist might try to do this is by rejecting the idea that truth is saying it how it is. The onus would then be on the supervaluationist to say what truth is—if not saying it how it is—and to show that her framework models truth in this sense. Given the points made above about supervaluationism as a model of assertability, I suppose one obvious move at this point would be to say that truth is some kind of idealised assertability.<sup>27</sup> It seems to me that the cost of such a move would be too high—but this is not the place to debate theories of truth. As far as the present paper is concerned, just seeing that saving supervaluationism as a model of truth involves abandoning the traditional realist view of truth as saying it how it is in favour of some kind of irrealist view of truth as idealised assertability would be progress, from the point of view of clarifying concepts and positions.

However, even sticking with the idea of truth as saying it how it is, there is still one way in which one might try to retain the supervaluationist machinery as a model of truth. Someone might say that when you assertively utter a sentence  $\alpha$ —say, ‘Bob is tall’—what you are talking about is not Bob and his tallness or lack thereof, but all legitimate ways of precisifying the language: you are really making a quantified metalinguistic claim of the form ‘On all legitimate precisifications of the language, ‘Bob is tall’ comes out true’. (Note that this is very different from the plurivaluationist view according to which you are simultaneously making many

claims, one relative to each acceptable classical model: making many claims, each about Bob and his possession or lack thereof of a particular precise property, is quite different from making a single claim that quantifies over the legitimate precisifications of ‘tall’.) On this view, we may associate a sentence  $\alpha$  with the following three quantified claims:

1. On every legitimate precisification  $\alpha$  is true.
2. On some legitimate precisifications  $\alpha$  is true and on some legitimate precisifications  $\alpha$  is false.
3. On every legitimate precisification  $\alpha$  is false.

Assuming a nonempty set of legitimate precisifications, exactly one of these three claims must be true (and the other two false).<sup>28</sup> The supervaluation keeps track of which one is true: assigning 1 to  $\alpha$  means that the first claim is true; assigning no value means that the second claim is true; assigning 0 to  $\alpha$  means that the third claim is true. Now the fact that the supervaluation assigns its three values in a non-value-functional way does not mean, on this interpretation of the machinery, that some connectives are non-truth-functional. The values assigned by the supervaluation do indeed provide a model of truth, on this view (as opposed to providing a model of probability or assertability)—but they are not themselves truth values. Rather, the value assigned to  $\alpha$  is a code telling us which of the three associated quantified statements is true. Thus, on this interpretation of the formalism, we can retain the supervaluationist machinery as a model of truth: but in this picture, there are no *truth values* that are assigned in a non-classical way.  $\alpha$  gets a truth value only within each classical model. The quantified statements get truth values outside such models—but there are only *two* such values (True and False) and they are assigned in an entirely classical way.

Furthermore, it should be noted that the idea just mooted—as opposed to the plurivaluationist picture from which it was distinguished—is a decidedly odd one. In saying that someone who says ‘Bob is tall’ is really talking about all the legitimate ways of precisifying the language, it seems to demand of speakers a far more sophisticated set of conceptual resources than we would normally think necessary for successful communication using vague sentences.

## 5. Conclusion

I have argued that if truth is saying it how it is, then truth values—values that represent presence, absence or quantity of truth—must distribute value-functionally across (at least) one-place predications and conjunctions and disjunctions thereof. This argument does not impact upon plurivaluationism as a model of truth, because in plurivaluationism truth values are assigned in an entirely classical way. (This is so within each acceptable model, which is classical—and truth values are not assigned anywhere else except inside acceptable models, on the plurivaluationist view.) It does impact upon supervaluationism as a model of truth: the values assigned by a supervaluation (whether tripartite or continuum-valued) are not distributed

value-functionally across one-place predications and conjunctions and disjunctions thereof and hence cannot be taken to be truth values. This is not to say that the formal supervaluationist frameworks have no use at all: they can be used to model probability of truth under precisification of the language—and they can be used to model assertability (although whether this will result in an empirically adequate model is still an open question). It does however mean that the central slogan of many supervaluationist views must be abandoned: supertruth is not truth.<sup>29</sup>

## Notes

<sup>1</sup> The formal technique of supervaluations is originally due to van Fraassen (1966). For seminal supervaluationist treatments of vagueness and future contingents (respectively) see Fine (1975) and Thomason (1970).

<sup>2</sup> It also applies straightforwardly to subvaluationism. Because this will be obvious to anyone who understands the duality between supervaluationism and subvaluationism (on which see e.g. Varzi (1999)) and because subvaluationist views are much less prominent in the literature than supervaluationist views, I shall not discuss subvaluationist views explicitly in this paper. For subvaluationist treatments of vagueness and future contingents (respectively) see Hyde (1997) and Ciuni and Proietti (2013).

<sup>3</sup> The main purpose of this section is to fix terminology—not to give an introduction of these views suitable for someone completely unfamiliar with them. For a full introduction (and references to the relevant literature) see Smith (2008, ch.2).

<sup>4</sup> See Smith (2012b) for an introduction to some of the main ones.

<sup>5</sup> See Decock and Douven (2014), Douven and Decock (forthcoming) and Decock and Douven (forthcoming). This work builds on Douven et al. (2013).

<sup>6</sup> The parenthetical remark here refers to the possibility of retaining supervaluationism as a model of truth by giving up the idea that the values assigned by the supervaluation are truth values. This option will be discussed in §4.4. In the meantime I shall not always repeat the caveat made parenthetically here.

<sup>7</sup> Both literatures are large and still growing. In order to give an idea of which literatures I mean, let me mention just a few (among many other) important works. On truthmakers: Armstrong (2004), Mulligan et al. (1984), Fox (1987), Bigelow (1988), Beebe and Dodd (2005) and Lowe and Rami (2009). On semantic groundedness: Herzberger (1970), Kripke (1975), Yablo (1982) and Leitgeb (2005). There are also other works relevant to the basic idea mentioned in the text (that each claim is about something—it has a subject matter) that do not fall neatly into either of the two literatures just mentioned. For example (and again among many others): Putnam (1958), Goodman (1961) and Yablo (2014).

<sup>8</sup> For convenience I write as though  $*$  is a two-place connective but the idea is completely general.

<sup>9</sup> Note that this does *not* mean that  $\alpha * \beta$  has a minimal subject matter. No such assumption is required in the following argument: we assume only that  $\alpha$  and  $\beta$  have minimal subject matters.

<sup>10</sup> To get from the claim that fixing  $[\alpha]$  and  $[\beta]$  fixes  $[\alpha * \beta]$  to the claim that  $*$  is truth-functional we require furthermore that the fixing of  $[\alpha * \beta]$  is uniform: that is, if  $[\alpha]$  and  $[\alpha']$  are fixed in the same way and  $[\beta]$  and  $[\beta']$  are fixed in the same way then  $[\alpha * \beta]$  and  $[\alpha' * \beta']$  are fixed in the same way. But given that  $\alpha$ ,  $\alpha'$ ,  $\beta$  and  $\beta'$  have minimal subject matters, this further claim is extremely plausible. To fix ideas—and because it is the case that will concern us below in the argument that supervaluationism does not provide an adequate model of truth—let's consider the case where  $\alpha$ ,  $\alpha'$ ,  $\beta$  and  $\beta'$  are one-place atomic predications. Where  $\gamma$  is a one-place atomic predication,  $\langle \gamma \rangle$  will be something like: how a certain object stands with respect to a certain property. So all the world has to go on in settling  $\langle \alpha * \beta \rangle$  and  $\langle \alpha' * \beta' \rangle$  is, in each case, a pair of facts concerning how an object stands with respect to a property. We are supposing that  $[\alpha]$  and  $[\alpha']$  are fixed in the same way and that  $[\beta]$  and  $[\beta']$  are fixed in the same way—so these facts are, respectively, the same (i.e. if the object involved in  $\langle \alpha \rangle$  possesses the property involved in  $\langle \alpha \rangle$  then the object involved in  $\langle \alpha' \rangle$  possesses the property involved in  $\langle \alpha' \rangle$  and so on). But then there is no way for  $\langle \alpha * \beta \rangle$  and  $\langle \alpha' * \beta' \rangle$  to diverge. For even if it is the case, say, that the properties involved in  $\langle \alpha \rangle$  and  $\langle \beta \rangle$  are compatible (i.e. it's possible for an object to possess both) while the properties involved in  $\langle \alpha' \rangle$  and  $\langle \beta' \rangle$  are incompatible, these further facts about compatibility and incompatibility cannot get into the act of fixing  $\langle \alpha * \beta \rangle$  and  $\langle \alpha' * \beta' \rangle$ : because  $*$  is conservative,  $\langle \alpha * \beta \rangle$  is already fully

fixed by  $\langle \alpha \rangle$  and  $\langle \beta \rangle$  and  $\langle \alpha' * \beta' \rangle$  is already fully fixed by  $\langle \alpha' \rangle$  and  $\langle \beta' \rangle$ . But now if  $\langle \alpha * \beta \rangle$  and  $\langle \alpha' * \beta' \rangle$  cannot diverge, then neither can  $[\alpha * \beta]$  and  $[\alpha' * \beta']$ .

<sup>11</sup> A terminological note. One property of connectives that one might want to investigate is as follows (to put it very roughly): adding the connective to a language does not enable the expression of any propositions that were not already expressible before the addition. A natural name for this property might be 'conservative'. But I am not interested in this property in this paper and I am certainly not claiming that connectives that are conservative in my sense have this other property. For example, disjunction is conservative in my sense, but however the new notion mentioned in this footnote is made precise, disjunction would seem to lack it.

<sup>12</sup> If there were another property that behaved like truth in the relevant ways then of course the argument would go through for that property also.

<sup>13</sup> A referee comments that my argument seems to rule out intuitionistic logic as a model of truth in the same way that it rules out supervaluationism—and raises the worry that if this is indeed the case then many readers will be tempted to reject my argument for that very reason (rather than rejecting intuitionism as a model of truth). In response, the first thing that I need to do is clarify that my argument does not apply to logics: it applies to semantics. Given a standard first order language, supervaluationist logic (thought of as a set of valid formulas or as a consequence relation on the language) is the same as classical logic (Smith, 2008, 82). My argument is not against this logic: it is against the idea that the supervaluationist semantics—involving a supervaluation that makes tripartite assignments of values to formulas—can be seen as a model of truth (given the traditional realist conception of truth outlined at the beginning of §3). So in the case of intuitionism, we need to consider not intuitionist logic in itself but semantics for intuitionist logic. Furthermore, we need to consider the intuitive interpretation placed on the semantics: that is, we need to ask whether the semantics is intended to model truth in the traditional realist sense. My argument counts against views that involve the following two features: (i) they employ a semantics that assigns values to conjunctions and disjunctions (of one-place atomic predications) in a non-value-functional way; and (ii) the values assigned to formulas are supposed to model presence, absence or quantity of truth—where 'truth' is understood according to the traditional realist conception. When we turn to semantics for intuitionist logic—together with intuitive glosses thereon—we in fact do not seem to find any that combine both these features. Hence there is no conflict between the argument of this paper and intuitionism. Of course I cannot here examine every extant approach to intuitionist logic—let alone every possible approach. What I shall now do is consider the two features mentioned above and for each one, give examples of approaches to intuitionism that lack this feature. In the end it will turn out that all of the mainstream approaches to intuitionist logic—and some others besides—have been mentioned (some more than once). Hence none of these approaches is in the firing line of the argument of this paper—or vice versa.

Feature (i). The Brouwer-Heyting-Kolmogorov (BHK) interpretation of the meanings of the logical operators in intuitionist logic (Troelstra and van Dalen, 1988, 9) lacks feature (i) because it does not involve assigning values to formulas at all: rather it says, for each kind of formula, what a proof of such a formula would consist in. Semantics for intuitionist logic that involve assigning values in a Heyting algebra (Stone, 1938) (Tarski, 1938) lack feature (i) because they assign values to formulas in a value-functional way. Kripke's semantics for intuitionist logic (Kripke, 1965) lacks feature (i) because, while it does not *in general* assign values to formulas in a value-functional way, it *does* assign values to conjunctions and disjunctions in a value-functional way; exactly the same can be said of Fine's semantics (Fine, 2014).

Feature (ii). A striking feature of many approaches to intuitionism is that either they do not mention truth at all—they talk, for example, about proof or justification or warrant or assertibility—or else they do talk about truth but make clear that they intend 'truth' in an epistemically infused sense (e.g. as verifiability or provability of some sort) which is in contrast to the traditional realist conception of truth. For example: The BHK interpretation is framed entirely in terms of proof (not truth). In Kripke's semantics, the values **T** and **F** do not represent truth and falsity in the traditional sense—they represent an *epistemic* notion of verification (or lack thereof) by a body of evidence (Kripke, 1965, 98). A similar point applies to the standard interpretation of Beth models (Beth, 1965, §145) (Kripke, 1965, 107) (Rabinowicz, 1985, 192, 210, 212). Dummett's account of intuitionism involves an understanding of truth as provability—and this intuitionist conception of truth is explicitly contrasted with the realist or platonist conception (Dummett, 2000). A similar point applies to Prawitz (1980)—and



according to Rabinowicz (1985), intuitionistic truth consists in *verifiability*. McDowell (1976) gives an account of intuitionism in which *truth* is distinguished from *known truth*: whether a complex statement is true is determined by whether its components (taken individually) are true—*truth* is distributed truth-functionally; but whether a complex statement is *known* to be true is (in general) not determined by whether its components (considered individually) are known to be true. It should be noted that there *are* interpretations of intuitionism that take it to be concerned with something like truth in the traditional realist sense. For example, Hazen (1982, 129) proposes a “less blatantly epistemological” interpretation of Kripke’s semantics—and Fine (2014, 549) argues for a view on which “intuitionistic logic might . . . be tied to a realist conception of the relationship between language and the world”. However Fine’s and Kripke’s semantics have already been mentioned under (i) above.

<sup>14</sup>Of course there may also be other uses for the supervaluationist framework, besides the two to be considered below.

<sup>15</sup>Similar remarks apply to supervaluationist treatments of statements about future contingents: the foregoing argument shows that the values assigned to such statements by the supervaluation cannot be truth values; however the supervaluationist framework can provide a model of *probability of turning out true*.

<sup>16</sup>Kamp (1975, 547) finds it surprising that it should have taken so long for machinery familiar from probability theory to be applied in semantics. From the point of view of this paper, however, declining to use this machinery for semantic purposes was entirely appropriate: probability measures provide a good model of *probability* (including probability of truth under precisification)—not of *truth*.

<sup>17</sup>There is a structural similarity between the definition of supertruth and the definition of truth in a model as satisfaction in the model relative to every value assignment on the model (henceforth ‘satisfaction-truth’). I have argued that supertruth cannot provide a model of truth but can provide a model of probability of truth. On a related note, I argue in Smith (2012a, 503–4, n.25) that satisfaction-truth (there referred to as ‘trewth’) does not provide an adequate analysis of truth (understood as saying it how it is): rather, satisfaction-truth captures the idea of something that would be true, whatever its free variables denoted, if its free variables were regarded as singular terms. Thus supertruth and satisfaction-truth are both at one remove from genuine truth (understood as saying it how it is): they capture things that *would* be true in certain circumstances, as opposed to things that *are* true.

<sup>18</sup>See Fine (1975), Kamp (1975), Osherson and Smith (1981), Osherson and Smith (1982), Kamp and Partee (1995), Keefe (2000) and Sauerland (2011). For further references and discussion of this literature see Smith (2015).

<sup>19</sup>Osherson and Smith present their argument in terms of degrees of membership of objects in sets rather than degrees of truth of statements.

<sup>20</sup>See Smith (2015).

<sup>21</sup>See for example Lewis (1991, 1991) and Jackson (1991, 1987).

<sup>22</sup>There are many possibilities here; see Smith (forthcoming) for further details.

<sup>23</sup>While the present paper concerns all uses of supervaluationism as a model of truth—not just supervaluationist approaches to vagueness—it is worth noting how the paper fits into the vagueness literature. In that literature, supervaluationists have criticised truth-functional approaches on the grounds that truth-functionality conflicts with ordinary usage of compound statements in the presence of borderline cases. Elsewhere I defend fuzzy approaches against these attacks (Smith, 2008, 2015). The present paper can be seen as turning the tables—moving from defence to offence: not only does the supervaluationist attack fail, truth-functionality (in certain areas) is in fact *essential* to an adequate model of truth.

<sup>24</sup>My point here is that the argument of this paper does not count against plurivaluationist semantics—nor against grafting onto such semantics an additional piece of machinery: a (tripartite or continuum-valued) supervaluation, thought of as assigning not truth values but assertability values. However, I think that there are other reasons why (classical) plurivaluationism cannot provide a good model of vague language (but a version of plurivaluationism built on fuzzy rather than classical models can); see Smith (2008) for details.

<sup>25</sup>Cf. Lewis (1986a, 212).

<sup>26</sup>For more on the differences between supervaluationism and plurivaluationism see Smith (2008, §2.5), which is where the term ‘plurivaluationism’ was first introduced in order to distinguish views that are formally distinct but often conflated in the vagueness literature.

<sup>27</sup> Cf. certain of the interpretations of intuitionist logic discussed in n.13.

<sup>28</sup> The legitimate precisifications are represented by classical models. On any classical model, every closed formula is true or false.

<sup>29</sup> Earlier versions of this paper were presented at a Philosophy Department seminar at the University of Melbourne on 25 September 2014, a Philosophy Department seminar at the Australian National University on 13 November 2014, and the Workshop on Vagueness via Nonclassical Logics at the University of Sydney on 19 December 2014. I am grateful to the audiences on those occasions—and the anonymous referees—for very helpful comments.

## References

- Aristotle. *Metaphysics*, c.350 BC. In Barnes (1984), vol.2.
- Armstrong, David Malet. *Truth and Truthmakers*. Cambridge University Press, 2004.
- Barnes, Jonathan, editor. *The Complete Works of Aristotle: The Revised Oxford Translation*. Princeton University Press, 1984.
- Beebee, Helen and Julian Dodd. *Truthmakers: The Contemporary Debate*. Oxford University Press, 2005.
- Beth, Evert W. *The Foundations of Mathematics: A Study in the Philosophy of Science*. North-Holland Publishing Company, second revised edition, 1965.
- Bigelow, John. *The Reality of Numbers: A Physicalist's Philosophy of Mathematics*. Clarendon Press, 1988.
- Ciuni, Roberto and Carlo Proietti. The abundance of the future: A paraconsistent approach to future contingents. *Logic and Logical Philosophy*, 22(1):21–43, 2013.
- Davis, Steven and Brendan S. Gillon, editors. *Semantics: A Reader*. Oxford University Press, 2004.
- Decock, Lieven and Igor Douven. What is graded membership? *Noûs*, 48 (4):653–82, 2014.
- . Two recent degree-theoretic approaches to the sorites paradox. In Otavio Bueno and Ali Abas-nezhad, editors, \*On the Sorites Paradox\*. Springer, forthcoming.
- . What verities may be. *Mind*, forthcoming.
- Douven, Igor, Lieven Decock, Richard Dietz, and Paul Égré. Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic*, 42(1):137–160, 2013.
- Dummett, Michael. *Elements of Intuitionism*. Clarendon Press, second edition, 2000.
- Edgington, Dorothy. Vagueness by degrees, 1997. In Keefe and Smith (1997), pages 294–316.
- Fine, Kit. Vagueness, truth and logic. *Synthese*, 30:265–300, 1975.
- . Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, 43:549–77, 2014.
- Fox, John F. Truthmaker. *Australasian Journal of Philosophy*, 65(2):188–207, 1987.
- Goodman, Nelson. About. *Mind*, 70(277):1–24, 1961.
- Hamilton, Edith and Huntington Cairns, editors. *The Collected Dialogues of Plato*. Pantheon Books, 1963.
- Hazen, Allen. On a possible misinterpretation of Kripke's semantics for intuitionistic logic. *Analysis*, 42(3):128–33, 1982.
- Herzberger, Hans G. Paradoxes of grounding in semantics. *Journal of Philosophy*, 67(6):145–167, 1970.
- Hyde, Dominic. From heaps and gaps to heaps of gluts. *Mind*, 108:641–60, 1997.
- Jackson, Frank. On assertion and indicative conditionals, 1979. In Jackson (1991), pages 111–35.
- . *Conditionals*. Basil Blackwell, 1987.
- , editor. *Conditionals*. Oxford University Press, 1991.
- Kamp, Hans. Two theories about adjectives. In Edward L. Keenan, editor, *Formal Semantics of Natural Language*, pages 123–55. Cambridge University Press, 1975. Reprinted in Davis and Gillon (2004), pages 541–62. Page references to latter.
- Kamp, Hans and Barbara Partee. Prototype theory and compositionality. *Cognition*, 57:129–91, 1995.
- Keefe, Rosanna. *Theories of Vagueness*. Cambridge University Press, 2000.
- Keefe, Rosanna and Peter Smith, editors. *Vagueness: A Reader*. MIT Press, 1997.
- Kripke, Saul. Semantical analysis of intuitionistic logic I. In J.N. Crossley and M.A.E. Dummett, editors, *Formal Systems and Recursive Functions: Proceedings of the Eighth Logic Colloquium, Oxford, July 1963*, pages 92–130. North-Holland, 1965.

- Kripke, Saul. Outline of a theory of truth. *Journal of Philosophy*, 72:690–716, 1975.
- Leitgeb, Hannes. What truth depends on. *Journal of Philosophical Logic*, 34(2):155–192, 2005.
- Lewis, David. Probabilities of conditionals and conditional probabilities, 1976. In Jackson (1991), pages 76–101.
- . *On the Plurality of Worlds*. Basil Blackwell, 1986a.
- . Probabilities of conditionals and conditional probabilities II, 1986b. In Jackson (1991), pages 102–110.
- Lowe, E. J. and A. Rami. *Truth and Truth-Making*. Acumen, 2009.
- McDowell, John. Truth conditions, bivalence, and verificationism. In Gareth Evans and John McDowell, editors, *Truth and Meaning: Essays in Semantics*, pages 42–66. Clarendon Press, 1976.
- Mulligan, Kevin, Peter Simons, and Barry Smith. Truth-makers. *Philosophy and Phenomenological Research*, 44:287–321, 1984.
- Osherson, Daniel N. and Edward E. Smith. On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9:35–58, 1981.
- . Gradedness and conceptual combination. *Cognition*, 12:299–318, 1982.
- Plato. *Cratylus*, c.360 bc. In Hamilton and Cairns (1963).
- Prawitz, Dag. Intuitionistic logic: A philosophical challenge. In G.H. von Wright, editor, *Logic and Philosophy*, pages 1–10. Martinus Nijhoff, 1980.
- Putnam, Hilary. Formalization of the concept “About”. *Philosophy of Science*, 25:125–130, 1958.
- Rabinowicz, Włodzimierz. Intuitionistic truth. *Journal of Philosophical Logic*, 14:191–228, 1985.
- Sauerland, Uli. Vagueness in language: The case against fuzzy logic revisited. In Petr Cintula, Christian G. Fermüller, Lluís Godo, and Petr Hájek, editors, *Understanding Vagueness: Logical, Philosophical and Linguistic Perspectives*, volume 36 of *Studies in Logic*, pages 185–98. College Publications, 2011.
- Smith, Nicholas J.J. *Vagueness and Degrees of Truth*. Oxford University Press, 2008. Paperback 2013.
- . *Logic: The Laws of Truth*. Princeton University Press, 2012a.
- . Many-valued logics. In Gillian Russell and Delia Graff Fara, editors, *The Routledge Companion to Philosophy of Language*, pages 636–51. Routledge, 2012b.
- . Undead argument: The truth-functionality objection to fuzzy theories of vagueness. *Synthese* (Special Issue on Vagueness and Probability), 2015. doi: 10.1007/s11229-014-0651-7.
- . Fuzzy logics in theories of vagueness. In Petr Cintula, Christian Fermüller, and Carles Noguera, editors, *Handbook of Mathematical Fuzzy Logic*, volume 3. College Publications (Studies in Logic, Mathematical Logic and Foundations series), forthcoming.
- Stone, M. H. Topological representations of distributive lattices and Brouwerian logics. *Časopis pro pěstování matematiky a fyziky*, 67(1):1–25, 1938.
- Tarski, A. Der aussagenkalkül und die topologie. *Fundamenta Mathematicae*, 31:103–34, 1938.
- Thomason, Richmond H. Indeterminist time and truth-value gaps. *Theoria*, 36(3):264–281, 1970.
- Troelstra, A. S. and D. van Dalen. *Constructivism in Mathematics: An Introduction*, volume 1. North-Holland, 1988.
- van Fraassen, Bas C. Singular terms, truth-value gaps and free logic. *Journal of Philosophy*, 63:481–95, 1966.
- Varzi, Achille C. *An Essay in Universal Semantics*. Kluwer, 1999.
- . Indeterminate identities and semantic indeterminacy. Typescript (available at [http://www.columbia.edu/~av72/papers/Parsons\\_2003.pdf](http://www.columbia.edu/~av72/papers/Parsons_2003.pdf)); text of a talk presented at *The Philosophy of Terence Parsons: Logic, Metaphysics, and Natural Language*, University of Notre Dame, 8 February, 2003.
- Yablo, Stephen. Grounding, dependence, and paradox. *Journal of Philosophical Logic*, 11(1):117–137, 1982.
- . *Aboutness*. Princeton University Press, 2014.