

# **Supervaluationism Without Gaps**

John T. Roberts

DRAFT

## **1. Introduction**

The account of vagueness that goes by the name of ‘supervaluationism’ says that there are truth-value gaps. In particular, whenever an object is a penumbral case of a vague predicate, it is neither true nor false that the predicate applies to the object. For example, if Mark is a borderline case of a bald man, then ‘Mark is bald’ is neither true nor false.

But I shall argue that the central idea behind supervaluationism does not automatically imply truth-value gaps. Moreover, gaps are not essential to how supervaluationism resolves the sorites paradox (which resolution is arguably the central plank in the case for supervaluationism). So we can reformulate supervaluationism in a way that eschews truth-value gaps, without abandoning the account’s central thought or losing its key virtue. In fact, this reformulation of supervaluationism is simply a more thorough and consistent implementation of the central idea that motivates supervaluationism in the first place; the standard version of supervaluation’s commitment to truth-value gaps represents a failure of nerve on the part of the supervaluationist. (Thus, the reformulation of supervaluationism is ‘supervaluationism without gaps’ in two distinct senses.) Furthermore, this reformulation deftly avoids some of the most prominent

objections to supervaluationism found in the literature. In this paper, I will call the standard version of supervaluationism ‘gappy supervaluationism,’ and the reformulation that I favor ‘resolute supervaluationism.’

Timothy Williamson (1994) anticipated this reformulation of supervaluationism, and argued that it resulted in an epistemic view of vagueness: ‘Of supervaluationism, nothing remains articulate’ (p. 164), he concluded. I shall argue here that Williamson was mistaken about this: We can be resolute supervaluationists, reject truth-value gaps, maintain the disquotationality of truth, and keep classical logic, without collapsing into epistemicism.

## **2. Supervaluationism: The Core Ideas<sup>1</sup>**

In classical truth-conditional semantics, a model  $M$  of a language  $L$  associates each singular term of  $L$  with an object in  $M$ ’s domain, and it associates each predicate of  $L$  with a certain set of objects in that domain, called the *extension* of the

---

<sup>1</sup> The formal machinery of supervaluationism has also been used to give accounts of many other topics, but here I always use “supervaluationism” as the name of a certain theory of the semantics and logic of vague language. Fine 1975 is an early, thorough exposition of supervaluationism about vagueness; see Williamson 1994 (Chapter 5) and Sorensen 2012 for more recent expositions, and Keefe 2000 (Chapters 7-8) for a defense of it as well.

predicate.<sup>2</sup> An atomic sentence 'b is F' is true just in case 'b' refers to an object that belongs to the extension of 'F'. Then, more complex sentences are assigned truth values relative to the model in the familiar compositional way. The central idea behind supervenience is to modify standard truth-conditional semantics in the following ways. First, a model M of a language L consists not just of one assignment of extensions to predicates, but to a set of them, each of which can be called an *allowable interpretation* of L relative to model M. Second, each precise predicate is assigned the same extension by every allowable interpretation, but a vague predicate is assigned different extensions in different interpretations; the many different extensions assigned a given vague predicate by a model can be called that

---

<sup>2</sup> For ease of exposition, I will write as if all vagueness were a matter of vague monadic predicates. The extension of the treatment here to cover vague relational predicates and vague singular terms is straightforward: In classical semantics, n-place predicates are assigned sets of ordered n-tuples of members of the domain of quantification as their extensions, and singular terms are assigned members of the domain as theirs; in supervenience semantics, each allowable interpretation in a model assigns the same kinds of extensions to these sorts of expressions as a classical model would. Also, for the sake of brevity I will neglect higher-order vagueness. Williamson 1994 (pp. 156-161) explains how supervenience can accommodate higher-order vagueness; the modification to supervenience that I will propose below does nothing to preclude treating higher-order vagueness in the same way.

predicate's *allowable sharpenings*<sup>3</sup> relative to that model. Each allowable interpretation determines a truth value for each sentence of L. Finally, if a sentence containing one or more vague predicates is true relative to every allowable interpretation in model M – in other words, if it is *supertrue* in M – then the sentence is definitely true in M; if the sentence is false relative to every allowable interpretation in M – i.e. it is *superfalse* – then the sentence is definitely false in M; if a sentence is neither supertrue nor superfalse in M, then according to M it represents a penumbral case. For example, 'Mark is bald' would be true if Mark belonged to all of the allowable sharpenings of 'is bald,' and it would be false if he belonged to none of them. As it is, though, (in the intended model of our language) Mark is a borderline case of a bald man, so he is a member of some of the allowable sharpenings but not others. 'Allowable sharpening' is a technical term, but it has an intuitive meaning: An allowable sharpening of a predicate is a set that includes every clear case of an object that satisfies that predicate, that excludes every clear case of an object that fails to satisfy that predicate, and that does not include one object and exclude another when the first clearly fits the predicate better than the second (e.g., no allowable sharpening of 'is bald' includes Mark but excludes someone with 5% more hair than Mark).

So far, the supervaluationist is committed to saying that supertruth is a sufficient condition for truth, and that superfalsity is a sufficient condition for falsity.

---

<sup>3</sup> Here I substitute "sharpening" for the more common but aesthetically objectionable term "precisification."

Nothing yet has been said about the truth values of sentences that are neither supertrue nor superfalse; we will come to them in good time.

One of the most important tasks for a theory of vagueness (and perhaps the most important) is giving a satisfactory treatment of the sorites paradox. Here is one form of that paradox:

If  $n$  grains of sand bunched together comprise a heap of sand, then so do  $n-1$  grains of sand bunched together.

One grain of sand by itself is not a heap.

Therefore, no finite number of grains of sand bunched together comprise a heap of sand.

The conclusion is obviously false, and the second premise is obviously true. But in classical logic, the conclusion follows validly from the two premises. So we seem to face a forced choice between three options: We may deny the first premise; we may revise classical logic in a way that renders the inference invalid; or we may accept that there are true contradictions (for example, that the conclusion of the inference is both true and false).

The supervaluationist response to the paradox is a version of the first strategy: reject the first premise. Anyone who uses this strategy must face *the problem of arbitrariness*: It seems that to deny the first premise is to affirm that there is some number  $n$  such that  $n$  grains of sand makes a heap but  $n-1$  grains do not, and this seems appallingly arbitrary.

The supervenientist appears to be able to make a case that the first premise is false that does not involve any objectionable arbitrariness. First of all, in the intended model of our language, every allowable sharpening of 'is a heap of sand' is a set of sand-aggregations that does not include any single, isolated grains of sand, that includes every pile of a million or more grains of sand, and that does not exclude any aggregation that is smaller than one it includes. It follows that for every allowable sharpening, there is an  $n$  such that there is a pile of sand containing  $n$  grains that belongs to it, whereas no pile of sand containing  $n-1$  grains or fewer belongs to it. Therefore, on every allowable interpretation, this sentence is true:

'There is a number  $n$  such that  $n$  grains of sand make a heap but  $n-1$  do not'

and this one is false:

'For every  $n$ , if  $n$  grains of sand bunched together comprise a heap of sand, then so do  $n-1$  grains of sand bunched together.'

Thus, the first premise of the sorites paradox is superfalse; therefore it is false.

Notice that this argument uses only the assumption that superfalsity is sufficient for falsity (and supertruth sufficient for truth); it depends on no assumption at all about the truth values of sentences that are neither supertrue nor superfalse.

But what about the problem of arbitrariness? There is no number  $n$  such that it is supertrue that  $n$  grains of sand make a heap but  $n-1$  do not, though it is

supertrue that there is some such  $n$ . So, the supervenient semantics declines to draw the arbitrary boundary between heaps and non-heaps at any particular place. Thus, the supervenient appears to evade the problem of arbitrariness, by declining to single out any arbitrary border. A worry remains, though: Is it enough simply to decline to say where the sharp boundary between heaps and non-heaps is drawn? Or is it also required that one deny that there is any such sharp boundary? We will, in effect, consider one answer to this question in section 3, and a second one in sections 4 and 5.

### **3. Gappy Supervenientism**

All prominent expositions of supervenientism<sup>4</sup> attribute to it a stronger claim than I have been doing so far: They attribute to it the identification of truth with supertruth and the identification of falsity with superfalsity. This has the consequence that any sentence that is neither supertrue nor superfalsity – that is, any sentence that comes out true on some allowable interpretations but false on others – is neither true nor false. This provides a neat solution to the arbitrariness worry: For any number  $n$  such that aggregations of  $n$  grains of sand are in the penumbral region between heaps and non-heaps, ‘A pile consisting of  $n$  grains of sand is a heap’ is neither true nor false. Likewise, for any such number  $n$ , ‘If  $n$  grains of sand

---

<sup>4</sup> For example, Fine 1975, p. 266 and pp. 272-273; Williamson 1994 p. 144; Keefe 2000 p. 202; Sorensen 2012, section 5.

bunched together comprise a heap of sand, then so do  $n-1$  grains of sand bunched together' is neither true nor false. Thus the supervaluationist not only declines to draw the boundary between heaps and non-heaps in any arbitrary place, but also denies that the world draws such a boundary: Any statement purporting to say where the boundary lies is untrue (albeit also unfalse) according to her.

This solution comes with various costs. First of all, the supervaluationist must deny that truth is disquotational in general; she must reject as untrue some of the T-biconditionals, i.e. statements of the form:

'S' is true iff S

A quick way to see why is to note first that since every classical tautology is true on every classical model, it follows that each one is true on every allowable interpretation in any supervaluationist model, and so supertrue in every supervaluationist model, hence a logical truth according to supervaluationism. Thus, every instance of the law of excluded middle, 'A or  $\sim A$ ,' is a logical truth according to supervaluationism. But the combination of excluded middle with truth-value gaps means that there will be untrue T-biconditionals. For example, 'Mark is bald or Mark is not bald' is true, but "'Mark is bald' is true or 'Mark is not bald' is true' is false (since each of 'Mark is bald' and 'Mark is not bald' is neither true nor false). But if both of these T-biconditionals were true:

'Mark is bald' is true iff Mark is bald



'Mark is not bald' is true iff Mark is not bald

then it would follow from excluded middle that one of 'Mark is bald' and 'Mark is not bald' was true; contradiction. So at least one of these T-biconditionals is false; presumably, the reasonable supervaluationist must say that both are false, as are all T-biconditionals for sentences that are neither true nor false.

What is more, the supervaluationist who identifies truth with supertruth must place restrictions on classical logic. As just noted, every classical tautology is a logical truth according to supervaluationism. But there are some classically valid inferences that are not valid according to supervaluationism. These inferences involve the 'definitely' operator, which is used to express the claim that someone is a clear case of satisfying (or failing to satisfy) some vague predicate, not a penumbral case. For example, 'Patrick Stewart is definitely bald, but Mark is not.' The sentence 'b is definitely F' is true just in case b belongs to every allowable sharpening of F, and it is false otherwise. So 'b is F' and 'b is definitely F' are true in exactly the same supervaluationist models (namely, all those in which 'b is F' is supertrue); thus, the inference from 'b is F' to 'b is definitely F' is necessarily truth-preserving, as is the converse inference; 'b is F' and 'b is definitely F' are inter-derivable. However, whereas 'b is F or b is not F' is logically true according to supervaluationism, 'b is definitely F or b is definitely not F' is false on some supervaluationist models. This implies that the classical rule of proof by cases is invalid in the supervaluationist framework. According to that rule, if you can derive C from A, and you can derive C from B, then you can derive C from (A or B). We get a

counterexample if we let A = 'Mark is bald,' B = 'Mark is not bald,' and C = 'Mark is either definitely bald or definitely not bald.' For, as we have seen, 'b is F' and 'b is definitely F' are inter-derivable, so from 'Mark is bald' we can derive 'Mark is definitely bald' and thence 'Mark is either definitely bald or definitely not bald.' Similarly, from 'Mark is not bald' we can derive 'Mark is definitely not bald,' and thence 'Mark is either definitely bald or not bald.' But we cannot derive 'Mark is definitely bald or definitely not bald' from 'Mark is bald or Mark is not bald,' since the latter is a tautology whereas the former is false. For similar reasons, the supervenientist must reject certain instances of the classical rules of *reductio ad absurdum*, conditional proof, and contraposition.<sup>5</sup>

These consequences are unattractive. At the very least, they limit the would-be supervenientist's options with respect to theorizing about truth and logic; at worst, they commit the supervenientist to things we have good reason to believe are false.

What is worse, accepting the identity of truth with supertruth commits the supervenientist to accepting significant restrictions on the central idea of supervenientism itself. Suppose we extend the supervenientist scheme to include a semantics for a language containing its own semantic metalanguage. Suppose we do this by going on in the same way: A supervenientist model consists of a set of allowable interpretations, each of which is itself equivalent to a

---

<sup>5</sup> See Williamson 1994, pp. 151-152 for details. But see also Keefe 2000, pp. 179-181 for an argument that this consequence is not nearly as bad as it seems.

classical model that classifies all the non-penumbral cases correctly<sup>6</sup>, and a sentence is supertrue (and therefore true) if it comes out true on every allowable interpretation. Now consider the sentence, 'The sentence 'Mark is bald' has a truth value.' On every allowable interpretation (in the intended model of our language), this is true. So shouldn't a supervaluationist count the sentence true? Similarly, on every allowable interpretation, the sentence 'The predicate 'is bald' has a certain definite extension' is true – so shouldn't the supervaluationist be prepared to count it as true too? Yet, the standard version of supervaluationism denies these things: The sentence in question has no truth value, and the predicate in question lacks a determinate extension.

As we have seen, the move of identifying truth with supertruth (and falsity with superfalsity) is not an automatic consequence of the basic motivating thought behind supervaluationism and plays no role in the way that supervaluationism blocks the sorites paradox. Nevertheless, it is undeniably tempting, primarily because it give the supervaluationist a neat reply to the problem of arbitrariness. But as we have seen, making this move commits the supervaluationist to unattractive positions in semantics and logic. And we have also just seen that it involves allowing a gap within supervaluationism itself: The central idea is not

---

<sup>6</sup> Of course, 'non-penumbral case' is itself a vague predicate. Hence, so are 'allowable interpretation,' 'allowable sharpening,' and 'supertrue.' Nonetheless, there can be clear cases of all of these predicates; for some sharpenings, it is supertrue that they are allowable sharpenings, etc.

applied to semantic vocabulary. Is there really no other way for the supervenientist to flesh out her view and reply to the arbitrariness worry?

#### **4. An Alternative: Resolute Supervenientism**

I suggest that the supervenientist should not have given in to the temptation to identify truth with supertruth, but should have instead remained resolute in her commitment to the core supervenientist idea, and extended it to semantic and logical metalinguistic vocabulary. I will call the resulting view *resolute supervenientism*.

Here is a sketch of the formal program: We start with a first-order language  $L$  containing classical logical vocabulary as well as metalinguistic names for all of its own terms and sentences, and metalinguistic vocabulary including ‘ $\_$  is true,’ ‘ $\_$  satisfies  $\_$ ,’ ‘the extension of  $\_$ ,’ and ‘ $\_$  follows validly from  $\_$ ,’ as well as the modifiers ‘a clear case of’ and ‘a penumbral case of.’ A *resolute supervenientist model* (RSM)  $M$  of  $L$  is a non-empty set of classical models of this language that may be called the *allowable interpretations* of  $L$  on  $M$ . Each sentence of  $L$  is true or false relative to each allowable interpretation, and truth on an interpretation is defined in exactly the same way as truth on a model is in classical semantics. In some cases, the conditions under which a sentence is true on a given classical model involve quantification over all classical models – e.g. ‘ $C$  follows validly from  $P$ ’ is true on a given classical model is true just in case  $C$  is true on every classical model on which

P is true; the conditions under which such a sentence is true on a given allowable interpretation thus involves quantification over all allowable interpretations on all resolute supervaluationist models (RSMs). (So, 'C follows validly from P' is true on a given allowable interpretation of a given RSM just in case C is true on every allowable interpretation from any RSM on which P is true.) In other cases, no such quantification over all allowable interpretations is involved: For example, in a classical model, 'S is true' is true just in case S itself is true in that model; similarly, 'S is true' is true on a given allowable interpretation just in case S itself is. Similarly, 'b belongs to the extension of F' is true on a given allowable interpretation just in case b belongs to the extension that interpretation assigns to F. The sentence 'b is a clear case of F' is true on a given allowable interpretation of a given RSM M just in case 'b is F' is true on every allowable interpretation in M; the sentence 'b is a penumbral case of F' is true on a given allowable interpretation of M just in case 'b is F' is true on some of the allowable interpretations in M and not on others. Finally, if a sentence is supertrue on an RSM M (i.e., true on all of the allowable interpretations in M) then it is true on M; if it is superfalse on M, then it is false on M. A sentence is true or false simpliciter just in case it is true on the intended model of our language.

There we stop: We do not add any assumption about the truth or falsity, in an RSM M, of any sentence that is neither supertrue nor superfalse relative to M. We agree that such sentences have truth values, but we make no comment on which truth value any of them has.<sup>7</sup> Let's see how far this gets us.

---

<sup>7</sup> Worry: Doesn't this amount to acknowledging that such sentences have definite truth values but *we simply do not know* which truth values they have? And doesn't

Note first that there is almost nothing novel in any of this: The key idea of supervaluationism was to let a model of a language consist of a whole set of classical models, and define truth relative to the individual classical models in the familiar classical way. The current proposal is just to keep on doing that, even when the language contains its own metalanguage. The only element of the present framework that is not imported directly from classical semantics via the core idea of supervaluationism is the semantic rules for the metalinguistic vocabulary concerning vagueness itself – namely ‘clear case’ and ‘penumbral case’ – but this element itself seems an inevitable consequence of the key intuition behind supervaluationism. (In fact, it is arguably nothing more or less than a statement of the central thesis of supervaluationism about vagueness.)

An immediate consequence of the proposal is that truth is disquotational:  
Every instance of the schema

‘S’ is true iff S

is true on every allowable interpretation of every RSM, so it is supertrue on every RSM; hence it is true on every RSM, so the resolute supervaluationist is committed to saying that it is true simpliciter.<sup>8</sup>

---

that mean the view being expounded reduces in the end to a form of epistemicism about vagueness? -- No, this worry is unfounded; see section 6.

<sup>8</sup> This might seem worrisome, since paradoxes of self-reference seem to pose serious problems for some instances of this schema. You can’t solve every problem

Another consequence is that the law of bivalence – ‘S is true or S is false’ – is true in every instance, even if S itself is a penumbral case. This is because on every allowable interpretation, either ‘S is true’ is true or ‘S is false’ is true, and in either case, the disjunction of these sentences is true on that interpretation. Every instance of bivalence is supertrue on every resolute supervaluationist model; hence every instance is true, simpliciter, according to the resolute supervaluationist. Similarly, metalinguistic statements such as ‘The predicate ‘is bald’ has a unique extension’ are true according to the resolute supervaluationist.

Another consequence is that ‘C follows validly from P1, P2, ... Pn’ is true in a given RSM iff the inference from {P1, P2, ... Pn} to C is classically valid. In other words, an inference is valid according to the resolute supervaluationist just in case its conclusion is true on every allowable interpretation on which its premises are true. This is equivalent to the criterion of validity that Williamson 1994 calls ‘local validity.’

Williamson distinguishes ‘local validity’ – a criterion of validity that an inference meets just in case its conclusion is true on every allowable interpretation of any model on which the premises are true – and ‘global validity’ – a criterion that an inference meets just in case its conclusion is supertrue on every

---

at once. But there is good reason to hope that the present view of vagueness could be combined with the correct resolution of the paradoxes: Take your favorite proposal for revising classical semantics to handle the paradoxes and let the allowable interpretations in a resolute supervaluationist models function exactly as models do under that proposal, and leave everything else here the same.

supervaluationist model on which all of its premises are supertrue. Global validity is the standard supervaluationist criterion of validity. Williamson notes that going with local validity instead would result in a logic equivalent to classical logic, and thus save the supervaluationist some headaches. But Williamson argues that this option should be unattractive to the supervaluationist, since the supervaluationist identifies truth with supertruth, and validity is standardly defined as necessary truth-preservingness; hence, a consistent supervaluationist should define validity as global validity. Were she to identify validity with local validity, while identifying truth with supertruth, she would cut off the notion of validity from the notion of truth, leaving it mysterious just what validity is supposed to be.<sup>9</sup> But this is no objection to the resolute supervaluationist, who endorses local validity as the criterion of validity while rejecting the identification of truth with supertruth.

## **5. Doubts About Resolute Supervaluationism**

Williamson 1994 anticipated the view I have just called ‘resolute supervaluationism’; he considered the possibility that a supervaluationist might reject global validity in favor of local validity as the criterion of logical validity, and reject the identification of truth with supertruth in favor of identifying truth with a predicate that is universally disquotational.<sup>10</sup> He claimed, as I do, that this move would free the supervaluationist from numerous difficulties. But he argued that a

---

<sup>9</sup> Williamson 1994, p. 148.

<sup>10</sup> Williamson 1994, pp. 163-164.



supervaluationist who made this move would no longer be a supervaluationist: ‘Of supervaluationism, nothing remains articulate.’<sup>11</sup> On such a view, the notion of clear case of F, like the complementary one of a penumbral case of F, have lost all connection with the notion of truth,<sup>12</sup> and if we cannot get a grip on them via their connection to truth, how can we get a grip on them? Williamson suggests that this can only be done by understanding a ‘penumbral case’ of, say, baldness as a case of a man of whom competent speakers cannot know whether he is bald or not; in short, the view collapses into the epistemic view of vagueness. Epistemicism and supervaluationism are supposed to be competing theories of vagueness, based on very different conceptions of how the phenomenon of vagueness works; surely, to show that a version of supervaluationism collapses into epistemicism is to show that it is no version of supervaluationism at all.

We might add to Williamson’s case: The resolute supervaluationist says, explicitly, that it is either true or false that Mark is bald; that ‘Mark is bald’ definitely has a truth value, and that truth value is either True or False; that the predicate ‘is bald’ has a unique, determinate extension and either Mark is in that extension or else he isn’t. But the resolute supervaluationist won’t say which way it is. Likewise, the resolute supervaluationist is committed to saying that there is some number  $n$  such that it is true that  $n$  grains of sand make a heap but  $n - 1$  grains do not. (For,

---

<sup>11</sup> p. 164.

<sup>12</sup> Williamson actually speaks of the “definitely” operator instead of the “clear case of” and “penumbral case of” modifiers, as I do here, but these are clearly notational variants.

'There is some number  $n$  such that it is true that  $n$  grains of sand make a heap but  $n - 1$  grains do not' is true on every allowable interpretation.) But the resolute supervenientist refuses to say which number it is. Not only is this a particularly irritating bit of caginess, it suggests that the resolute supervenientist has really given up on the key idea of supervenientism and joined the forces of epistemicism. For the resolute supervenientist says that there is a boundary between the heaps and the non-heaps, and a boundary between the bald and the non-bald – it's just that her theory leaves it open precisely where that boundary lies. What could the resolute supervenientist mean, then, if not that there really is a precise boundary somewhere but we simply do not know (and perhaps cannot know) where it lies?

Yet another worry about resolute supervenientism is that it seems to have no answer to the arbitrariness worry. The worry is that if the first premise of the sorites paradox is false, then it would seem to follow that there is some precise borderline separating the heaps from the non-heaps, but any such precise border would be arbitrary: Why should the border be here, rather than a little bit to one side or the other? There seems to be nothing in the nature of the world or in our usage of our language that singles out one precise boundary over many other candidates, so the existence of such a boundary would appear to be an ungrounded brute fact. The world must contain ungrounded brute facts somewhere, of course; the outcomes of irreducible chance processes, the fundamental laws of nature, and the boundary conditions of the universe might be among them. But these things are all plausible candidates for basic features of the universe, which can be as they are without being grounded in anything else. By contrast, if there were a precise

location for the heap/non-heap boundary, this would be a case of a truth whereof it is unintelligible how anything about reality or our language could make it true. It is hard to understand how our universe could contain such a truth at all.

Standard supervenience has a neat answer to this worry: The sentence 'there is a number  $n$  which marks the minimum number of grains of sand an aggregation must contain in order to be a heap' is true, since it is supertrue. But every substitution instance of '\_\_\_ is the number that marks the heap/non-heap boundary' is untrue; if we fill in the blank with a number that is too big or too small, then we get a falsehood, whereas if we fill it in with a number in the penumbral range, we get a sentence that is neither true nor false. So there is no ungrounded fact in the world about the precise location of the boundary. This is an instance of the principle that an existential generalization can be true even if none of its substitution instances are true – a consequence of the supervenience semantics and the identification of truth with supertruth. But the resolute supervenience is committed to saying that every true existential generalization has a true substitution instance. (For the sentence 'Every true existential generalization has a true substitution instance' is itself true on every allowable interpretation; hence it is supertrue in every resolute supervenience model, and therefore true simpliciter.) And indeed, she is committed to saying that there is a precise boundary between the heaps and the non-heaps. Nothing in her theory tells us where the boundary lies, but the theory itself clearly implies that the boundary exists, and that it is sharp. So it seems that the resolute supervenience has no reply to the arbitrariness worry; the conclusion seems inescapable that she simply must bite the bullet here, and

allow that the universe contains ungrounded brute facts about the boundaries of our predicates. Thus, it sacrifices one of the key virtues of standard supervenience.

These objections appear to add up to a very strong case against resolute supervenience. But they can all be overcome. The resolute supervenientist has a coherent view of the semantics of vague language that is quite different from that of the epistemicist; she can offer an account of how we understand the notions of clear cases and penumbral cases without connecting these notions to that of truth; she has a compelling reply to the arbitrariness worry. The key to all of this is the notion of there being (or failing to be) a *fact of the matter* about something, a notion that plays a key role in the intuitive picture behind supervenience, and one which the resolute supervenientist can adequately accommodate within her formal framework.

## **6. Resolution of These Doubts**

The supervenientist (whether of the gappy or the resolute variety) wants to say the following: A predicate is vague when all the facts in the world that are relevant to fixing the semantics of predicates – facts that perhaps include regularities in usage, selection histories, baptism ceremonies, inferential roles, causal connections between utterances and things, the joints in nature which make some sets more ‘eligible’ than others, and so forth – are not up to the job of nailing

its extension down precisely. (This is quite different from what the epistemicist wants to say – namely, that these facts do indeed nail down the extension precisely, but in such a way that we are incapable of knowing precisely how they do so.)

However, the resolute supervaluationist adds that the semantics-fixing facts are indeed up to the task of fixing certain other things precisely – namely, that each predicate has some determinate extension or other, that every meaningful sentence has one truth value or the other, and so on. They do the latter by way of fixing the roles that terms like ‘extension,’ ‘predicate,’ ‘truth value’ and so on play within our language; the semantic facts about those terms make it the case that, for example, every predicate has a unique set as its extension, and so forth. So the semantics-fixing facts are up to the task of nailing it down that ‘Mark is bald’ has one of the truth values, but they are not up to fixing which one it has.<sup>13</sup> In other words, it is

---

<sup>13</sup> Imagine a language game in which participants speak of their own and others’ “favorite U. S. President,” played by the following rules: Each player has authority over his or her favorite President, in the sense that whenever they publicly declare “my favorite U. S. President is \_\_,” what they have said – and all of its consequences, given the actual historical facts – must be accepted as true by all the other players, subject only to the restriction that all such statements be consistent with each other and with the facts about the U. S. Presidents to date. So for example, I can say “My favorite U. S. President is Andrew Jackson,” and thenceforth it is true that my favorite U. S. President is Andrew Jackson, and it follows that it is true that my favorite U. S. President was an opponent of central banks. But I can also make up my mind about my favorite President little by little, so to speak; I can say “My

indeed a fact that 'Mark is bald' has one of the truth values, but there is no fact of the matter about which truth value it has.<sup>14</sup>

---

favorite U. S. President is a twentieth-century President,” without saying anything more to narrow it down. If I do say that then it has to be accepted as true (according to the rules of this game) that my favorite President was a man, and that he was either a Democrat or a Republican (since the only Presidents to belong to neither of those parties served prior to the twentieth century). Thus, this sentence must also be recognized as true in this game: “[Author’s Name]’s favorite President belonged to either the Democratic Party or the Republican Party, but not to both. “ The rules of this game, plus the current state of play and the facts of the surrounding world, suffice to make it correct to assert this disjunction without making it correct (or incorrect!) to assert either disjunct. One wants to say that the rules of the game have run out here – they do not *completely* specify what counts as correct and incorrect play. But this doesn’t stop the game from being played within its rules, any more than the fact that the rules of baseball fail to specify what should happen if a player magically transformed the ball into a butterfly shows that baseball cannot be played within its rules. If the suggestion in the text is true, then discourse involving vague predicates (in other words, all discourse) is something like this language game. (There’s an analogy; take it or leave it.)

<sup>14</sup> Standard supervenience (what I prefer to call “gappy supervenience”) disagrees with this claim. But notice how similar it is to a claim that supervenience is all more than happy to make: Namely that the disjunction

But can the resolute supervenientist consistently say this? It depends on how we explicate the expression ‘there is a fact of the matter.’ I suggest that as it is used in informal English, this expression is ambiguous. By ‘There is a fact of the matter about whether S,’ one thing we can mean is: Either S is true, or else S is not true (i.e. S is false). We might call this *the semantic reading* of ‘fact of the matter.’ If that’s what we mean by the expression, then the resolute supervenientist must agree that there is a fact of the matter about whether, say, Mark is bald. But another thing we can mean by it is: the facts about the world and about the semantics of our language as they now stand – without any additional or ad hoc arbitrary stipulation on our part now – are sufficient to make it the case that S is true, or else they are sufficient to make it the case that S is false. We might call this *the meta-semantic reading* of ‘fact of the matter’; it seems to me that in cases where philosophers want to deny that ‘there is any fact of the matter’ about something, it is more often this second meaning that they have in mind.

On the meta-semantic reading, to say that there is or is not a fact of the matter about whether S is to make a meta-semantic comment – a comment about those features of the world that determine the semantic properties of linguistic expressions, as such. (It isn’t exclusively meta-semantic – it is about the non-linguistic facts of the world as well. But part of the information the comment conveys is about meta-semantics.) Since the resolute supervenientist wants to apply her supervenientism to all discourse, she will apply it to meta-semantic

---

“Mark is bald or he isn’t” is definitely true, though there is no fact of the matter about which one of its disjuncts is true.

discourse as well; so, she will say that a meta-semantic comment is true if it is supertrue. Resolute supervenientism does not include a theory of meta-semantic facts. But presumably, whatever the true meta-semantic facts of our world fix is that our language has the range of allowable interpretations that it does, and that baldness and heapiness, for example, have the ranges of allowable sharpenings that they have. These ranges – and not any specific interpretation or sharpening within them – are the semantic facts about vague language, so they are what the meta-semantic features have to determine. Meta-semantic facts might themselves be vague and it might be a vague matter just what the semantic features of some expressions are. But some things are definitely true; for example it is definitely true that the meta-semantic facts about the world do not suffice to narrow down the extension of ‘is bald’ to a unique set (assuming, that is, that epistemicism is mistaken). So, by a resolute supervenientist’s lights, it must be true on every allowable interpretation of our complete vocabulary (including object-level, semantic, and meta-semantic vocabularies) that the meta-semantic facts suffice to fix only a range of allowable sharpenings for ‘is bald,’ rather than a unique extension. Thus, it may be supertrue that the meta-semantic facts of the world, together with the non-linguistic facts, do not by themselves suffice to determine whether ‘Mark is bald’ is true or false. In other words, it is supertrue (and thus, true) that there is no fact of the matter about whether ‘Mark is bald’ is true or false (on the meta-semantic reading of ‘fact of the matter’). Nevertheless, according to resolute supervenientism, the sentence ‘Mark is bald’ is either true or false’ is true on every allowable interpretation of our entire vocabulary, so it is



supertrue and thus true that 'Mark is bald' is either true or false. Thus, the resolute supervenientist can consistently maintain that although 'Mark is bald' has one truth value or the other, there is no fact of the matter (in the meta-semantic sense) about which truth value it has.

Now let's return to the objections against resolute supervenientism that we noticed above. First of all, how are we to get our grip on the notions of a clear case and a penumbral case, if not via their relation to the notion of truth? The answer is that we get them via the notion of there being a fact of the matter, in the meta-semantic sense. Mark is a penumbral case of a bald man; this means that there is no fact of the matter (in the meta-semantic sense of this phrase) about whether Mark is bald. The sentence 'Mark is bald' has a truth value, and that truth value has to be either True or False, but there is no fact of the matter about which one it is. That is to say, the extra-linguistic facts, together with the facts about language-use and so on that fix the semantics of our predicates, fail to determine an answer to the question of whether Mark is bald. We can say this without mentioning truth at all (I just did so), but it gives a clear sense to the notion of a penumbral case (and, indirectly, that of a clear case, or something's definitely being the case) that fits our pre-theoretic sense of what makes a piece of language vague. It is also, as we have seen, easily accommodated within the resolute supervenientist framework.

The resolute supervenientist agrees that the rules for using our semantic meta-vocabulary make it definitely true that "The sentence "Mark is bald" has a truth value' is true – in other words, they make it so that 'Mark is bald' has a truth value; she also agrees that it is definitely true that there is some number  $n$  which is the

precise cut-off point between heaps and non-heap aggregations of sand. But she denies that there is (in the meta-semantic sense) any fact of the matter about what that truth value is, or about which number  $n$  is the cut-off. This provides both a solution to the problem of arbitrariness and an answer to Williamson's question 'What then remains of supervaluationism?'<sup>15</sup> What is appalling in the idea of there being an arbitrary line between the heaps and the non-heaps, or between the bald men and the non-bald men, is that it is hard to imagine what aspects of the world could possibly make it the case that the boundary lies here rather than some other place very nearby, and equally hard to imagine that this could be an ungrounded brute fact. The resolute supervaluationist, however – by contrast with the epistemicist – denies that there is any ungrounded brute fact of the matter here, and also that there is any grounded fact of the matter. It isn't just that we cannot know the fact of the matter in question; God couldn't know it either, because there is nothing there to know. It is correct, on her view, to say that the relevant sentences are either true or false, but incorrect to imagine that there is any feature of the world – whether grounded or fundamental – that makes it one or the other. And the arbitrariness worry is a worry about the latter, not the former.

It cannot be denied that there remains a puzzling thing about resolute supervaluationism: It maintains that a sentence like 'Mark is bald' is either true or false, while not only declining to say which it is, but also denying that there is any fact of the matter about which it is. This makes it seem a bit as if the resolute supervaluationist takes back through one corner of her mouth what she

---

<sup>15</sup> Williamson 1994, p. 164.

simultaneously says through the other corner. This bears a strong family resemblance to a puzzling feature of the standard supervaluationist view, namely that it says that Mark is either bald or he is not, but there is no fact of the matter about which. In each case, there is an appearance of weirdness because we are inclined to think that the truth of a disjunction must 'flow from' the truth of one disjunct or the other; in order to live with the apparent weirdness of standard supervaluationism, we need to learn to stop thinking that, and once we have learned to stop thinking it, we should have no trouble living with the apparent weirdness of resolute supervaluationism either. And from the supervaluationist point of view, we need to learn to live with this appearance of weirdness anyway, because it is nothing other than the apparent weirdness that a mind trained in classical logic is apt to find in vague language as such.<sup>16</sup>

### **References:**

Fine, Kit 1975. 'Vagueness, truth and Logic.' *Synthese* 30: 265-300.

Keefe, Rosanna 2000. *Theories of Vagueness*. Cambridge University Press.

---

<sup>16</sup> [Acknowledgements omitted for review]

Sorensen, Roy 2012. 'Vagueness', The Stanford Encyclopedia of Philosophy (Summer 2012 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2012/entries/vagueness/>.

Williamson, Timothy 1994. *Vagueness*. Routledge.