Philosophical Studies (2005) 122: 1–25 

JOHN HAWTHORNE

# VAGUENESS AND THE MIND OF GOD*

ABSTRACT. This paper examines the mind and language of an omniscient being from a supervaluationist perspective. Two questions hall receive special attention. How ought the supervaluationist explicate the concept of omniscience? And what ought the supervaluationist expect an omniscient speaker to say about a Sorites series?

How ought we to conceive of an omniscient mind in light of the phenomenon of vagueness? Working within the framework of one of the more promising approaches to vagueness – supervaluationism – this paper explores the mind (and language) of God. In section 1 I sketch the barebones of supervaluationism, familiar to most readers. Section 2 is concerned with how the supervaluationist ought to define omniscience. Section 3 describes, from a supervaluationist perspective, how the divine mind would engage with a Sorites series. Finally, in section 4, I address a thought experiment concerning omniscient speakers offered by Timothy Williamson, designed to support the view that there are epistemically inaccessible hidden boundaries associated with ordinary vague predicates. I conclude by noting the broader philosophical significance of the themes explored here.

## 1. SUPERVALUATIONISM

The supervaluationist traces the existence of vagueness to the phenomenon of semantic indecision. Consider a simple example from David Lewis:[1] It is vague whether an external carport is part of a house. Let us call the fusion of the main body of some house with its external carport 'Big', the body of the house 'Little'. It is plausible to suppose that the English language delivers no verdict as to whether it is Big or Little that merits the predicate 'is a house'.

That indecision of this sort exists will often not matter. If I say 'The house is majestic', that may well be unobjectionable on

either way of making the predicate 'is a house' precise, for both Big and Little may be eminently majestic. The supervaluationist suggests that we semantically evaluate sentences beset by semantic indecision by examining admissable ways of making such sentences precise.[2] A sentence is supertrue just in case it is true on all ways admissable ways of making it precise. A sentence is superfalse just in case it is false on all admissable ways of making it precise. Clearly, some sentences will be neither supertrue nor superfalse. According to the supervaluationist, we can make sense of the phenomenon of borderline cases using this semantic machinery: a borderline case of say, baldness, is one where an ascription of the predicate 'is bald' is neither supertrue nor superfalse.

As Kit Fine urged in his seminal paper on the topic,[3] the main interest of this approach is that it can validate logical connections between sentences that, taken individually, are neither supertrue nor superfalse. Thus if Joe is a borderline case of baldness, it may nevertheless be supertrue that if Joe is bald then he is not not bald and supertrue that either Joe is bald or Joe is not bald. If a carpet is borderline between yellow and orange, it is neverthless supertrue that if it is yellow it is not orange (since the semantic decision that 'yellow' and 'orange' cannot be simultanously true of an object *has* been made) ... and so on. This already makes for a big advantage over any approach to vagueness grounded in any of the standard truth tables for many valued logics, since the latter do not validate logical connections between vague sentences. Of particular note is the fact that the supervaluationist is able to embrace classical logic, on the grounds that classically valid sentence forms will have all and only supertrue instances.[4]

Supervaluationism is motivated by the idea that when there is semantic indecision, there is no inscrutable hidden boundary, accessable only to God (or worse still to no possible being). Of course, the world often contains boundaries that are hidden from us. We were at one point altogether unable to detect the joint in nature that separates fool's gold from gold, and that which separates Parkinson's disease from other superficially similar degenerative disorders. But borderline cases do not, according to the supervaluationist, point to a whole new range of hidden boundaries. This motivation puts further constraints on the shape of supervaluationist

semantics: If we are, as semanticists, to avoid positing boundaries of which we are in principle ignorant, then we must inevitably deploy a vague metalanguage. Just as there is no sharp line between the people that are bald and the people that are not, so there is no sharp line between the properties that are acceptable precisifications of 'is bald' and the ones that are not; and so on. Formal models have been developed suitable to the intuitive picture of "vagueness all the way up," built by analogy with modal logics that deny the S4 and S5 axioms.[5]

It is often noted that supertruth and superfalsity do not obey the disquotational schemas standardly held to be canonical for truth and falsity: One cannot accept

If 'u' means P, then ('u' is supertrue iff P)[6]

or

If 'u' means P, then ('u' is superfalse iff ∼P)

This renders pressing the issue whether or not to identity truth and falsity with supertruth and superfalsity and correlatively whether or not to abandon the standard axiom schemas for truth and falsity. On this issue – a real trouble spot for supervaluationism – the following discussion will remain neutral.

## 2. OMNISCIENCE

How should the supervaluationist characterize omniscience? As we shall see, the issue is both delicate and instructive. There seem to be three alternative conceptions available. The first conception opts for the following definition:

(1) ∀x (x is omnscient iff ∀P ((x knows P iff x believes P) and (x believes P iff P)))[7,8]

(The more straightforward '∀x (x is omniscient iff ∀P (x knows P iff P))' has the drawback of allowing an omniscient being to *believe* certain falsehoods.)

Supervaluationists standardly introduce an operator 'Definitely' into their language, governed, inter alia, by the schema: If 'u' means P, 'u' is supertrue iff definitely P.[9] This permits the acknowledgement of borderline cases without resorting to semantic ascent. Thus,

the claim that there is no sharp line between bald and non-bald people can be expressed by

$$\sim\forall x \text{ (definitely x is bald or definitely x is not bald)}$$

A second conception of omniscience can be articulated using such an operator, viz:

(2) $\forall x$ (x is omniscient iff $\forall P$ ((x believes P iff x knows P) and (x believes P iff Definitely P))).

Since the schema

P iff Definitely P

is unacceptable, conception (1) and (2) cannot be regarded by the supervaluationist as equivalent.

A third conception is motivated by the idea that an omniscient being only has thoughts that involve precise ingredients: a belief ascription that involves a vague singular term or a vague predicate cannot be true of an omniscient being. Now there are plenty statements that are definitely true that involve either vague singular terms or vague predicates. For example: If I am bald, then I am bald. The third conception contends that if 'that P' involves any vague terms or predicates, then any belief ascription using it will be false of the omniscient being.[10] In short: only where 'P' is *both* precise and supertrue can we use it to describe the contents of God's mind. Let us introduce an operator 'Precisely', governed by the schema,

> If 'u' means P, then (Precisely P iff ('u' is supertrue and all the constituents of 'u' are precise)).[11]

We can now articulate the third conception as follows:

(3) $\forall x$ (x is omniscient iff $\forall P$ ((x believes P iff knows P) and (x believes P iff Precisely P)))

Given that nearly all of our predicates are vague, we can, on this view, make very few positive knowledge attributions to God that are correct. Whatever its merits, the shortcomings of the third approach are obvious: it entails, for example, that while *we* know that bald people are bald, an omniscient being does not know any such thing.[12] Before retreating to such a compromised conception of omniscience, we should at least investigate whether either of the

first two conceptions can be made to work. (I shall not be exploring the third conception any further in this paper.)

Which of the first two conceptions is preferable? Return to the toy example with which we began. The supervaluationist is committed to each of the following:

(4) It is indefinite whether Big is a house.[13]
(5) It is indefinite whether Little is a house.
(6) Either Little is a house or Big is a house.

If it is indefinite whether P, for some P, then it would not seem that I am in a position to know that P. Knowledge is unavailable in a borderline case. Thus, given that Little and Big are both borderline cases of being a house, the following pair of claims seem right.

(7) I don't know that Little is a house.
(8) I don't know that Big is a house.

With (4)–(8) in view, it is easy enough to recognize the costs of our first two conceptions. Supposing we embrace the first conception and claim that God is omniscient. We are forced to the conclusion:

(9) Either (Little is a house and God knows it and I do not know it) or (Big is a house and God knows it and I do not).

That seems bad. As far as Big and Little's claim to be houses goes, it doesn't seem that God knows any more than I do. But (9) suggests he does. Against (9), it seems intuitively clear that God doesn't know that Little is a house (and that he doesn't know that Big is a house). And if you asked Him why, say, he does not know that Little is a house he would tell you: It is because the concept house does not definitely apply to Little. So it seems that

(10) God doesn't know that Little is a house

and

(11) God doesn't know that Big is a house.

Combine (10) and (11) with (9) and contradiction ensues.[14]

Further if we combine the first conception of omnisicience with (4), (5) and (6), we get the bizarre

(12) Either (it is indefinite whether Little is a house and God knows that Little is a house) or (it is indefinite whether Big is a house and God knows that Big is a house).[15]

Initially, the second conception may seem to raise equally serious problems for the supervaluationist. If (2) rather than (1) is taken as the proper characterization of omniscience, and we assume that God exists, then we are compelled to accept

(13) (Either Little is a house and God does not know that Little is a house) or (Big is a house and God does not know that Big is a house).

Prima facie, this is disturbing, for it suggests that there is something that God does not know. Recall, however, that supervaluationists must already be willing to tolerate:

(14) Either (Little is a house and it is indefinite whether Little is house) or (Big is a house and it is indefinite whether Big is a house).

If one has learned to live with (14), (13) is likely to strike one as less disturbing. So while (13) is surprising at first, it does not seem to offer grounds for abandoning the second conception that are nearly as compelling as those that were offered for abandoning the first. So I suggest, albeit somewhat tentatively, that of the three candidate characterizations of omniscience, it is (2) that best accords with the supervaluationist's other commitments. Since definitions are, plausibly, supertrue, we can now endorse

(15) Def[16] ∀x (x is omniscient iff VP ((x knows P iff x believes P) and (x believes P iff Def P)))

To lend further support to this choice of definition, let me address another objection to the second conception. Supervaluationists have said little about how supervaluations work for terms as they occur within belief contexts. But in choosing among the candidate definitions of omniscience, it matters crucially how belief contexts are treated. We have imagined that there are two precisifications of 'house', one which holds of Little (and not Big), the other of which holds of Big (and not Little). Call these precisifications house[1] and house[2] respectively. Let us agree that God believes that Little is a house[1] and that he does not believe that Little is a house[2] It might

now be suggested that any that-clause involving 'is a house' is beset with semantic indecision. Since propositions themselves are never indefinite, the line of thought runs, it must be that the ascription is vague as to which proposition is being ascribed as the object of belief. One precisification of 'God believes that Little is a house' is 'God believes that Little is a house$^1$' another is 'God believes that Little is a house$^2$'. Accordingly it will be suggested (contrary to the second conception), that it is indefinite whether God believes that Little is a house, since that belief ascription is correct on one precisification, incorrect on another.

This model is hardly faithful to how ordinary belief ascriptions work. We can all agree that Clinton believes that he has a house. According to the model just adumbrated, the latter has two precisifications – Clinton believes that he have a house$^1$; Clinton believes that he has a house$^2$ – each being required to hold in order for the original belief ascription to definitely hold. But we are not nearly so willing to accept either precisification as we are the original belief attribution! Something is wrong with the model.[17]

A more promising supervaluationist approach to belief attributions maintains that the vague concept *house* is definitely distinct from the precise concept *house*$^1$ and that one may have a belief involving the latter and not the former (and vice versa). That one concept is an acceptable sharpening of another does not require that the concepts are indefinitely distinct. In passing,[18] Kit Fine remarks that the claim 'Casanova believes that he has had many mistresses' can be taken in two ways: As a "precise report of a vague belief or as a vague report of a precise belief"[19] (say, that he has had more than 15 mistresses). I take it that the first use is more typical: it would in most contexts be misleading to use the ascription if Casanova believed he had had more than 15 mistresses but had no idea as to whether more than 15 mistresses was a lot of mistresses. To avoid confusion, we should not lapse into the second use without warning. Adhering now to the first, more standard use, we should insist that while the predicate 'is a house' expresses a vague concept in that it admits borderline cases, the claim that I believe I have a house serves as a perfectly precise report of my state of mind.[20] Meanwhile, ordinary people definitely don't believe that they own a house$^1$, since they don't have the concept *house*$^1$. And God definitely does

not believe Little is a house, since while he ascribes the concept *house*[1] to Little, he definitely does not ascribe the concept *house* to Little. (Of course he definitely does not ascribe the complex concept *Not a house* to Little either.) It is a happy consequence of the second conception that God's omniscience is not compromised by such failures to believe.

Let me end by noting how the second conception makes trouble for a much discussed proposal concerning propositional attitude ascriptions. It is sometimes maintained that proper names occur transparently within the scope of attitude ascriptions, and thus that coreferring proper names are (surprisingly) substitutable salva veritate in attitude contexts.[21] Proponents of such a thesis will endorse the following schema (henceforth 'the transparency' schema):

S knows a is F iff ∃x (x = a and S knows x is F)

But this schema cannot be accepted by a proponent of the second conception. Assume with such a proponent that ∼God knows that Little is a house. Suppose further that the house figuring in our focal example has a proper name, say 'Balmoral'. It is clear enough that we want to say that (i) Definitely, Balmoral is a house (ii) it is indefinite whether Balmoral is identical to Little and (iii) it is indefinite whether Balmoral is identical to Big. Now we should not reckon valid any schema that allows us to derive a claim that is not definitely true from a set of claims that are definitely true. But the transparency schema, in combination with the second conception, allows us to do just that. As follows:

Given that definitely, Balmoral is a house, we should agree that:

(16) God knows Balmoral is a house

Applying the transparency schema we now get

(17) Ex (x = Balmoral and God knows x is a house)

The following is a logical truth:

(18) Ex ((x = Balmoral and God knows x is a house and Balmoral = Little) ⊃ Ey (y = Little and God knows y is a house))

which, assuming the transparency schema, is equivalent to

(19) Ex ((x = Balmoral and God knows x is a house and Balmoral = Little) ⊃ God knows Little is a house))

But we have agreed that, given the second conception

(20) ∼God knows Little is a house.

From (18), (19), (20), we can now derive

(21) ∼Little = Balmoral

But (21) is not supertrue. Something has gone wrong. Assuming the second conception, we should not endorse the transparency schema. There is a more general point to be made, of course: if knowledge is absent in borderline cases, we should not endorse the transparency schema. (After all, the above derivation could have been run for anyone that clearly knows that Balmoral is a house but clearly does not know that Little is a house.) It thus seems that considerations of vagueness, as applied to proper names, provide compelling grounds to give up the transparency schema. Some no doubt will be inclined to run the argument in the other direction, arguing that if I know that Balmoral is a house, then at worst it is indefinite whether I know Little is a house. That diagnosis seems less plausible, refusing as it does to acknowledge that vagueness brings a lack of knowledge in its wake.

These matters are admittedly very delicate; I cannot pursue them further here. Having motivated the second conception, I shall adhere to it in what follows. Some readers will no doubt be unconvinced. Fortunately they should have little trouble seeing how the remarks that follow might be adapted to fit the first conception.

## 3. GOD ENCOUNTERS A SORITES SERIES

Recall that a key motivation for supervaluationism is the idea that borderline cases do not generate hidden boundaries. On such a picture, suitable reflection would enable us to know all there is to know about how, say, the concept *bald* pertains to some Sorites series. If there is nothing that we are not in a position to know about such a series (at least in connection with questions to do with baldness), we should not, it seems, find any great conceptual strain in allowing an all-knowing being to come into the picture.[22] Further, if there is nothing to know about baldness that can't be known by us, then it doesn't seem that we could learn anything from such a being

that we couldn't, in principle, learn by inquiring into the matter for ourselves.

It turns out that matters are rather more complicated than this initial picture would suggest. There are two key points that are easy to overlook but which are vital to a proper understanding of the matter. I shall present them in turn.

### 3.1. *Pristinity and Tracking*

Suppose that we have encountered a being – God – who definitely fits the second conception of omniscience. And so:

(22) Def God is omniscient.

We have adopted the following conception of omniscience:

> Def ∀x (x is omniscient iff VP ((x knows P iff x believes P) and (x believes P iff Def P)))

Let an omniaccurate being be one that meets the following weaker characterization:

(23) Def ∀x (x is omniaccurate iff VP (x believes P iff Def P))

God, let us assume, is no mere borderline case of omniaccurateness. And so:

(24) Def ∀P (God believes P iff Def P)

In what follows, we shall focus on the implications of omniaccurateness: the additional requirements of omnisicience will not matter. I shall assume throughout the following standard rule for logics of definiteness:

(25) $\vdash \mathrm{Def}(\alpha \supset \beta) \supset (\mathrm{Def}\,\alpha \supset \mathrm{Def}\beta)$

and also that any logical truth is definitely true.[23]

Suppose we sit God down and run through a familiar Sorites series, removing one hair at a time from a person that is, at the outset, clearly not bald, asking at each step 'Is he bald?' How are we to describe what goes on?

It is important to notice that in order to describe this situation in detail, we need to make some decision about the nature of God's language (or more precisely, the relevant fragment of language that is being put to use in conducting this test). In particular, we need to decide whether or not the language in question is "slippery" or

"pristine". An utterance type u is pristine just in case it is always a sharp matter whether or not u has been uttered by God. An utterance type u is slippery just in case it is not pristine.

What would it be like if God deployed a slippery language? Consider the following story: I say to God: "Ok, I'm going to take away one hair at a time and at each step tell me if the person in front of me is not bald." God replies: "Ok I'll draw a picture of a bald person if and only if the person is bald." As I remove progressively more hairs, God draws increasingly hair-free depictions. After a while, it will be plausible to say that God has stopped drawing pictures of people who are definitely not bald. And a while after that, it will be plausible to say that God has started drawing pictures of people who are definitely bald. As for drawings made during the intervening period, we will say on each occasion that God has drawn a picture of someone who is neither definitely bald nor definitely not bald. The utterance type which is used to say that the person is bald is the act of drawing a bald person. Obviously there will be cases in which it is indefinite whether the utterance is performed. A slippery language is in play.

Such a story seems intelligible enough. But it does seem rather unsatisfying. Our natural conception of a Sorites-related conversation with God is one where God uses a pristine language – so, for example, supposing that 'Yes' is part of the language fragment in use, it is sharp matter whether or not God utters 'Yes' in response to a question.[24] (I shall return to slippery languages later.)

Suppose I take someone – call him 'Dave' – and remove one hair from his head at eleven A.M. each day. Lets imagine further that at least it is a sharp matter whether the noises: 'Dave is not bald yet' has been uttered by God at any given time. (Suppose that the time of the utterance is given by when the utterance is completed). So for any time, either it is definite that 'Dave is not bald yet' has been uttered by God at that time or else it is definitely not the case that 'Dave is not bald yet' has been uttered by God at that time. So, for example,

(26) Def God utters 'Dave is not bald yet' at noon on Tuesday or Def $\sim$God utters 'Dave is not bald yet' at noon on Tuesday.[25]

We haven't yet given those noises semantic significance. But now we tell God what to do with the noises – we explain to him a fragment of the language that we want him to use in talking to us. The rule is simple: utter 'Dave is not bald yet' at noon on a given day iff David is not bald at at noon on that day. Each day, at about eleven A.M., we remove a hair and await God's pronouncement.

Can God use the language in accordance with our directive? Suppose he definitely accepts the instructions. Then, definitely, God believes that he will utter 'Dave is not bald yet' at noon on Tuesday iff Dave is not bald yet at noon on Tuesday. So we have

(27) Def God believes (God utters 'Dave is not bald yet' at noon on Tuesday iff Dave is not bald yet at noon on Tuesday).

It follows that

(28) Def (God utters 'Dave is not bald yet' at noon on Tuesday iff Dave is not bald yet at noon on Tuesday).

(From (24) and (27) using (25) and Modus Ponens.[26])

Combine (28) with (26) – the assumption that there are no borderline cases of uttering 'Dave is not bald yet' – and we get the conclusion that

(29) Either Def Dave is not bald yet at noon on Tuesday or Def ∼Dave is not bald yet at noon on Tuesday.[27]

This is a terrible result for our supervaluationist. Tuesday might well be one of the days when Dave was a borderline case of bald.

What is the supervaluationist to say? Let's retrace our steps and see what has given rise to the situation. Reflection reveals that the problem has arisen from two sources working in tandem. First, we assumed that the relevant fragment of the language God was speaking was pristine. Second we assumed that God definitely believed that his use of the fragment would conform to a certain biconditional of the form

(30) God will utter 'u' at t iff a certain condition C obtains at t.

Call the first assumption – that it is sharp whether or not the relevant utterances in the story are made – *Pristinity*. Call the second – that the relevant utterances are believed by God to conform to certain

biconditionals of form (30) – *Tracking*. Simultaneous commitment to Pristinity and Tracking is tantamount to a commitment to sharp boundaries. (This is the first of the two key points that I wish to bring out.[28]) So if we are committed to a condition C's admitting of borderline cases, and we wish to maintain our commitment to a being who definitely satisfies the second conception of omniscience, then we must abandon our commitment to that being's simultaneous conformity to Pristinity and Tracking. We have already seen how abandoning Pristinity can help in this regard. Let us see how the same holds for the abandonment of Tracking. Suppose, for example, that God merely intends (and believes) that he will conform to the rule:

(31) God utters 'Dave is not bald yet' only if Dave is not bald yet

then, even with a commitment to Pristinity and borderline cases, no incoherence ensues. For God could conform to (31) by staying silent throughout the entire event, or by uttering 'Dave is not bald yet' on a couple of occasions of definite baldness, staying silent the rest of the time. But, as before, Tracking offers a natural gloss on what it would take for an omniscient being to conform to the conversational maxim of being maximally informative. If God is to use 'Dave is not bald yet' as a vehicle for conveying the information that Dave is not bald yet and is to be maximally informative, then we would not expect Him to be silent or to be sparing in his utterance of 'Dave is not bald yet' If so, however, then the supervaluationist's God must give up on Pristinity: if there are vague boundaries of the sort the supervaluationist claims there are, then an omniscient being who intends to be maximally informative will find a pristine language unusable.

Note that what goes for God's language also goes for God mind. Omniaccuracy gives us:

(32) Def (God believes Dave is bald on Tuesday iff Def Dave is bald on Tuesday).

Assume that God's state of mind is sharp, that it is definite whether God has some given belief on not. This gives us

(33) Def God believes Dave is bald on Tuesday or Def $\sim$God believes Dave is bald on Tuesday

We can now derive

(34) Def (Def Dave is bald on Tuesday) or Def ∼(Def Dave is bald on Tuesday)

Where does this all leave us? As we noted in the opening sketch, supervaluationists wish to make room for second order vagueness. There will be cases where it is indefinite whether definitely P. If there is definitely an omniaccurate being and his belief states are sharp, then we can rule out second order vagueness for any given P by a derivation of the sort just given. Thus a supervaluationist who wishes to accommodate a being who is definitely omniaccurrate will take care not to assume that God's mind is sharp. (In effect, we deny pristinity for God's language of thought.) Supposing it is indefinite whether Def Dave is bald on Tuesday, the supervaluationist will insist that it is also indefinite whether God believes that Dave is bald on Tuesday.

### 3.2. *Penumbral Connections*

Let me pursue a further complicating thread that has been kept out of view thus far.

To begin, we should remind ourselves that if 'P' is true on some precisifications and 'Q' is false on some precisifications, it does not follow that there is some precisification on which 'P and not Q' is true. This is because there may be logical/conceptual connections between 'P' and 'Q' – what Fine calls 'penumbral connections' – that must be respected by any precisification of that sentence. To take a simple example, 'That is a pale shade of red' might be true on some precisifications, 'That is red' may be false on some precisifications, but 'That is a pale shade of red and that is not red' will not be true on any way of making that sentence precise, since the inferential connection between 'is a pale shade of red' and 'is red' must be respected by any way of making the complex sentence precise. Our supervaluationist theist must similarly hold that there is a penumbral connection between 'God believes that Dave is not bald on Tuesday' and 'Definitely Dave is not bald on Tuesday'. Suppose that it is indefinite whether Def Dave is not bald on Tuesday and also indefinite whether God believes that Dave is not bald on Tuesday. By supervaluatonist lights, there will in such a case be precisifications on which 'God believes that Dave is not bald on Tuesday' is false

and there will be precisifications on which 'Def Dave is not bald on Tuesday' is true. If there is no penumbral connection between the relevant pair of sentences, there will be no reason to deny that there is a precisification on which 'God believes that Dave is no bald on Tuesday iff Def Dave is not bald on Tuesday' is false. But if there is such a precisification, then God cannot, after all, be definitely omniscient or definitely omniaccurate since

(35) Def (God believes Dave is not bald at noon on Tuesday iff Def Dave is not bald at noon on Tuesday)

will not hold. So a penumbral connection between the belief ascription in (35) and the definiteness claim to the right of the biconditional in (35) will have to be instituted.[29] If no such penumbral connection can be made coherent, then the supposition of a being that is definitely omniscient (as opposed to one that is indefinitely omniscient) will begin to look suspect. (I leave it to others to explore this issue further.)

Analogous points can be raised in connection with various stories that invoke definitely omniscient beings that deploy slippery languages. Consider the following story: "A definitely omniscient being is subjected to a Sorites series. He agrees (and believes) that he will issue a loud clap of thunder at noon on any given day iff Dave is not bald on that time. On a certain day – say Tuesday – we think to ourselves that Dave is a borderline case of bald. We listen for guidance and notice that the clap of thunder we hear is only borderline loud." The story seems at first glance coherent by supervaluationist lights. But it is not so clear that it is. Suppose that on Tuesday Dave is a borderline case of baldness. The story had it that

(36) Def (God issues a loud clap of thunder at noon on Tuesday iff Dave is bald at noon on Tuesday)

(since God is definitely omniaccurate and also believes the claim embedded by the Def operator in (36)). Meanwhile, we are supposed to think that on Tuesday it is indefinite whether God issues a loud clap of thunder and also indefinite whether Dave is bald. But there is no special logical or conceptual connection between the loudness of thunder and the baldness of a person. So there would seem to be some ways of making precise the biconditional 'God issues a

loud clap of thunder at noon on Tuesday iff Dave is bald at noon on Tuesday' such that on those precisifications, the biconditional is false. But in that case, (36) cannot hold, contra the assumption of the story. What we learn – and here is the second key point – is that if God is to speak a slippery language and Tracking is to hold, then some penumbral connection has to be in force that connects that question whether the relevant utterance has been performed to the question whether the relevant condition holds. Notice that in the earlier, picture-drawing story, such a penumbral connection did hold – and one can no doubt tell other such stories. But without a penumbral connection between utterance and condition, our story risks lapsing into incoherence. Notice, once again, that if we merely assume that God is a borderline case of omniscience (or omniaccurracy), no threat of incoherence will arise. But if he is to be definitely omniscient, and to satisfy Tracking, then the language he uses will have to be subject to very special strictures: it will have to be slippery, and there will have to be a penumbral connection between utterance and condition.

## 4. WILLIAMSON ON OMNISCIENT SPEAKERS

Timothy Williamson has observed that omniscient speakers make for something of a puzzle for those, like the supervaluationists, who claim that nothing epistemically elusive is marked by ordinary vague predicates. The puzzle that he raises is the following. Consider an omniscient speaker, asked at each step in a Sorites series whether or not a certain condition holds – for example, suppose he is presented with a progressively diminishing pile of sand grains, and asked at each step whether the grains in front of him form a heap. We are sure that the being will say 'Yes' to 'Is it a heap?' when confronted with, say, ten million grains (suitably arranged) – and we are similarly sure that He will not be saying 'yes' to that question by the time we are down to a single grain. So it seems that God would have to stop at some point – and that we are in no position to know where that point will be. Moreover, it is hard to believe that such a stopping point would lack semantic significance. After all, it seems that one could be confident that if we ran the series again, He would stop at the same place. As a result, Williamson suggests, we should

conclude if that there must, after all, be a hidden boundary associated with any given vague predicate. Here is Williamson's statement of the puzzle:

On the view that nothing is hidden, it should be harmless to imagine omniscient speakers, ignorant of nothing relevant to the borderline case. Such a hypothesis of itself carries no commitment to bivalence or excluded middle. It is supposed, for example, that if TW is thin, then the omniscient speaker know that TW is thin, and that if TW is not thin then the omniscient speaker knows that TW is not thin. It is no part of the hypothesis that TW is thin or TW is not thin. . . .

Accompanied by an omniscient speaker of English, you remove grain after grain from a heap. After each removal you ask 'Is there still a heap?'. The omniscient speaker is not required to answer 'Yes' or 'No'; she can say 'That is indeterminate' or 'To degree 0.917' or 'You are asking the wrong question' if she likes. If there is nothing to say, she will remain silent. She will not say 'I don't know', nor will she hesitate, unsure of the best answer. You can ask her not to mumble. She is determined to be cooperative and as relevantly informative as she can.

After the first few removals, what is left remains quite clearly a heap. The first few times you ask 'Is there a heap?', the omniscient speaker answers 'Yes'; she may say other things too, but whatever else she says, on those first few occasions she says 'Yes'. She is not an obscurantist. After sufficiently many of the grains have been removed, she does not say 'Yes'. She is not a liar. For some number n, she says 'Yes' after each of the first n removals, but not after n plus 1. You do not know the value of 'n' in advance. Yet on the view in question, her stopping point cannot represent a hidden line between truth and something less than truth. How can that be?

You repeat the experiment with other omniscient speakers, starting with exactly similar arrangements of grains in exactly similar contexts, then removing the grains in exactly the same way. The trials are independent; there is no collusion between the omniscient speakers. If they all stop at the same point, it evidently does mark some sort of previously hidden boundary, although it may be a delicate matter to say just what it is a boundary between.

. . . You can instruct the omniscient speakers how to use their discretion. For example, you can instruct it to use it conservatively, so that they will answer 'Yes' to as few questions as permissible. . . .

Now if two omniscient speakers stop answering 'Yes' at different points, both having being instructed to be conservative, the one who stops later has disobeyed your instructions, for the actions of the other show that the former could have used her discretion to answer 'Yes' to fewer questions than she actually did . . . Thus if all are instructed to be conservative, all will stop at the same point. You do not know in advance where it will come. It marks some sort of previously hidden boundary . . .

The discussion contains a slip that is easily corrected. The idea of running independent trials on omniscient speakers makes no sense. For how is one to keep the results of various trials a secret from an *omniscient* being! One could, however, run the thought experiment with beings whose omniscience was restricted to matters having to do with heaps,[30] the rest of the details being essentially unchanged. So this hardly takes us to the heart of the issue.

With section 2 in mind, the supervaluationist might also object to the initial characterization of omniscience. Williamson seems implicitly to appeal to the first conception of omniscience – but as we have seen, the supervaluationist would do well to insist on the second conception, defining omniscience in terms of what is definitely the case.

One might also object to Williamson's failure to address the possibility that the instruction "Be maximally conservative" is itself vague owing to borderline cases of maximal conservativeness.

More importantly, though, I want to suggest that, once accepted in their own terms, the preliminary materials of the thought experiment – prior even to the invocation of a multitude of speakers and an instruction to be conservative – already wreak havoc for the supervaluationist. Three features of the thought experiment stand out, corresponding to the untenable trio identified in section 3.1:

(a) It is clearly being taken for granted that our omniscient speaker is no mere borderline case of omniscience – that is, that our omniscient speaker is definitely omniscient.

(b) The "no mumbling" clause is, in effect, an instruction to presume that the utterance of 'Yes' is pristine: the story has it that the omniscient speaker, being cooperative, is not to mumble in such a way that it is not clear whether or not he has said 'Yes'.

(c) The suggestion that the omniscient speaker is to be 'cooperative' and 'as relevantly informative as she can' induces us into supposing that the omniscient speaker is determined to utter 'yes' after a removal if and only if a heap remains after that removal.[31] That is, the story seems to presuppose that the speaker is committed to Tracking.

After all, we are being in effect told that for any removal r:

(37) (A heap remains after removal r and she does not say 'yes' after r) ⊃ She will not be relevantly informative and cooperative (she would be an obscurantist).

and

(38) (A heap doesn't remain after removal r and she does say 'yes' after r) ⊃ She will not be relevantly informative and cooperative (she would be a liar)

For one who accepts the validity of classical logic – as the supervaluationist does – this amounts to the information that

(39) ∀r (She is relevantly informative and cooperative ⊃ (A heap remains after removal r iff she says 'yes' after r)).

Given that it is part of the description of the case that the omniscient speaker is determined to be relevantly informative and cooperative, we are forced to assume that the speaker is determined to make it the case that

(40) ∀r (A heap remains after removal iff she says 'yes' after r).

To assume that the omniscient speaker accepts (34) is to assume that Tracking is in force. We are now, already, caught in incoherence.

Perhaps the later stages of the thought experiment will be important to those who reject classical logic. But the supervaluationist cannot accept even the initial terms of the thought experiment insofar as it involves a determinately omniscient speaker, Pristinity and Tracking.

The earlier result holds: insofar as a determinately omniscient speaker uses a pristine language, he cannot coherently take himself to be governed by the sorts of biconditionals that we would naturally associate with being maximally informative. The result may be surprising to some, admittedly: but I do not see it as a cogent basis upon which to convict the supervaluationist of incoherence.

A variant on Williamson's thought experiment is worth considering,[32] one which makes important use of a multitude of omnisicient speakers, while dropping the tracking assumption. Suppose that each of, say, ten thousand definitely omniscient speakers is given the following pair of instructions: Say "Heap" only if what is in front of you is a heap. Say "No heap" only if what is in front of you is not a

heap. An omniscient speaker who remained silent throughout would obey the instructions. But presumably many would offer sundry "Heap" and "No heap" utterances. One could keep track of which cases received some "Heap" responses, which "No heap" responses, which neither. As we ran more and more trials, evidence would accumulate. True enough, different speakers would have different levels of reticence with regard to their willingness to offer "Heap" and "No Heap" verdicts. But after ten thousand trials, two lines would have emerged – one between those cases that received at least some "Heap" verdicts and those that received none, a second between those cases that received at least some "No heap" verdicts and those that received none.[33] We could not easily anticipate in advance where these two lines would appear. Wouldn't they then, plausibly, mark a pair of semantically significant hidden boundaries? Of course, there is no logical guarantee of such. We could imagine our multitude to conspire to refuse en mass to offer a "Heap" verdict in certain clear cases of heaphood. But wouldn't we have at least accrued considerable evidence of semantically significant hidden boundaries in such a case? And absent some kind of conspiracy, isn't it obvious that such multiple trials will cumulatively reveal hidden boundaries?

I suspect that the force of this last thought experiment relies on our forgetting that the divide between omniscient and nonomniscient speakers (whether or not we are considering omniscience simpliciter or omniscience regarding some restricted subject matter) may itself be vague. Without pretending to offer a final word, I offer the following, counterveiling, intuition pump. Suppose a tribe contains some contingently omniscient speakers. We know that if a member of such tribe drinks ten pints of beer he or she is definitely sufficiently impaired as to cease to be omniscient. Meanwhile we know that if a member of the tribe drinks one pint of beer he or she is definitely still omniscient. We run trials on groups of tribemembers in varying states of inebriation using the same protocol as above, deploying a Sorites series on heaphood and the same pair of instructions. No instance in the Sorites series will be immune from any given verdict. But we will notice that some cases will receive a "Heap" verdict only from tribespeople that surpass a certain level of inebriation. Would we now be able to detect some sharp hidden

boundary of semantic significance? Only if some level of inebriation clearly marked the transition point from omnisicience to less than omniscience. But if that transition point is not sharp, we could not very well extract a hidden boundary from our various trials. Look back to our original ten thousand trials, described in the previous paragraph. Various tribepeople in our new trial will certainly have given a "Heap" verdict in other cases. Call these the "deviants" Assuming that omniscience, as well as definite omniscience, is no sharp matter, we would be in no position to assert that all the deviants are to be classified as not omniscient, or even not definitely omniscient. But insofar as we are in no position to assert *that*, we are in no position to claim that the pair of lines that emerged from the original trial are semantically significant.

In conclusion, then, it remains unclear whether any thought experiment involving omniscient speakers makes real trouble for the supervaluationist who denies the existence of inscrutable hidden boundaries associated with vague predicates. Certainly, Williamson's original thought experiment is inconclusive; nor is any in the vicinity clearly more decisive.

## AFTERWORD

The wider relevance of the preceding remarks should not be overlooked. Even those unsympathetic to the metaphysical possibility of an omniscient being will take seriously the possibility of local omniscience[34] about some restricted, though still Sorites susceptible, subject matters. Not only do we idly recognize such possibilities. As philosophers, we frequently invoke some counterfactual idealized being (or else an idealized "limit of inquiry") as a tool in philosophical analysis. Certain ethicists have suggested that x is good iff x would be reckoned good by an ideally reflective agent. Some philosophers of modality have suggested that P supervenes on base Q just in case ($Q \supset P$) would be a priori for an idealized agent. And so on. Many of the preceding puzzles and concerns have analogues in connection with those philosophical analyses. Consider a simple illustration. In a borderline case of goodness, should we suppose that there would be no fact of the matter as to whether an ideally reflective agent would accept a goodness ascrip-

tion in that case? Or should the proponent of the ethical analysis just mentioned retreat to

> ∀x ((Def x is good) iff (x would be reckoned good by an ideally reflective agent))

(in effect supplanting the first conception of (restricted) omniscience by the second)? Relatedly, in a case where it is definite that either x is good or x is not good but indefinite whether x is good, should we learn to live with the problematic disjunction

> Either (it is indefinite whether x is good and an ideally reflective being would reckon x good) or (it is indefinite whether x is good and an ideally reflective being would reckon x not good)?

Obviously, this is not the place to pursue these connections further.

## NOTES

* Thanks to Cian Dorr, Mark Scala, Ted Sider, audiences at the University of Arizona and King's College, London, and especially Tamar Szabo Gendler and Timothy Williamson for comments and discussion.

[1] See "Many, But Almost One," in *Papers in Metaphysics and Epistemology* (Cambridge University Press, 1999), p. 172. The expression 'semantic indecision' is his.

[2] Lewis' advice to think of a precisification as a way of 'making the unmade semantic decisions' is helpful here. See "Many, But Almost One," p. 172.

[3] "Vagueness, Truth and Logic," *Synthese* 30 (1975), pp. 265–300.

[4] It may be, however, that certain classically valid *inference patterns* cannot be recognized as valid by the supervaluationist. See Williamson's *Vagueness* (Routledge, 1994), Chapter 5. The arguments in the text do not trade on the disputable inference patterns (which accounts for their longwindedness in some cases).

[5] See, for example, Fine's "Vagueness, Truth and Logic" and Williamson's "On The Structure of Higher-Order Vagueness," *Mind* 108 (January 1999), pp. 127–143.

[6] After all, 'P' may be borderline and yet 'u' is supertrue' may be superfalse, in which case there will be some precisifications on which 'P' is true and yet "u' is supertrue' is false.

[7] Those that hold that there can be knowledge without belief will obviously want to do things a little differently.

[8] I shall use propositional quantification throughout. Here is not the place to defend its cogency.

[9] Assuming that we allow that a language can contain its own supertruth predicate. If not, the definiteness operator will have to be explained slightly differently. The relevant technical issues do not bear much on the topic of this paper.

[10] Relevant here is the discussion of supervaluationism and propositional attitude attributions that follows shortly.

[11] Of course, more work would need to be done to convince us that the "Precisely" operator is coherent.

[12] Notice, moreover, that since our semantic and psychological concepts – means, refers, believes, loves and so on – are vague, we could not on this view coherently think of God as believing that we mean anything, refer to anything, believe anything or love anything. I note for the record that in conversation the third conception nevertheless proved attractive to many interlocutors.

[13] Here 'It is indefinite whether P' has its standard meaning, abbreviating '∼Definitely P and ∼Definitely ∼P'.

[14] I assume here and elsewhere that the supervaluationist embraces modus ponens and conjunction introduction as inference rules, and that she accepts as valid all the classically valid sentence forms. (That supervaluationism may not license certain other classically valid inference rules (see Williamson, *Vagueness*, p. 151ff.) is thus not relevant.)

[15] Those supervaluationists who wish to identify supertruth with truth can offer a further consideration against the first conception. Knowledge entails truth. That is: ∀P ∀x (x knows that P ⊃ It is true that P). If truth is supertruth, this is equivalent to: ∀P Vx (x knows that P ⊃ Definitely P). Combine this with the first conception and we get ∀x (x is omniscient ⊃ ∀P (P ⊃ Definitely P)) This is clearly unacceptable, unless we intend our conception of omniscience to be so crafted that it be blatantly infelicitous to believe that something is omniscient.

[16] Henceforth I shall abbreviate 'Definitely' by 'Def'.

[17] As a way of trying to salvage the model criticized in the text, it may be suggested that there are various ways of precisifying the binary predicate 'believes'. On each way of precisifying 'believes' there is a particular class of precise propositions that that count as the objects of thought. So perhaps 'believes 1' is one precification of 'believes', 'believes 2' another and Clinton believes 1 that he has a house 1 and believes 2 that he has a house 2 (the latter two claims being the acceptable precisifications of 'Clinton believes that he has a house'). We can now maintain consistently that it is indefinite whether Clinton believes that he has a house 1 (since this is only true on one precisification) and definite that Clinton believes that he hs a house. I leave further exploration of such a view to others. I note in passing that a satisfying supervaluationist account of belief ascription will have to be coordinated with an account of the 'means that' construction, given the intuitive link between accepting a sentence and believing a proposition, viz: If S accepts sentence s and s means that P then S believes that P.

[18] See "Vagueness, Truth and Logic," p. 289.

[19] Or at any rate, more precise.

[20] As the beginnings of a toy theory, let us imagine with Frege that when 'house' occurs within belief contexts, what is normally its sense serves as its reference. In ordinary contexts, its reference might in indeterminate, there being a range of extensions that serve as acceptable precisifications of that predicate. But in the context of belief attributions, its reference might be determinate, namely to a particular sense that admits borderline cases.

[21] For an introduction to the topic, see Kripke, "A Puzzle About Belief," reprinted in Peter Ludlow, *Readings I the Philosophy of Language* (MIT 1997), esp. pp. 906–909 (reprinted from Margalit (ed.), *Meaning and Use* (Reidel, 1979)). For a recent extended discussion and defense, see Scott Soames, *Beyond Rigidity* (Oxford University Press, 2002).

[22] Cf. Williamson, *Vagueness*, p. 198ff, whose discussion inspired this paper. I shall discuss Williamson's treatment in detail in section 4.

[23] So here, for example, is how to derive (24) from the assumption (24A) Def God is omniaccurate. We have (24B) Def ((∀x x is omniaccurate iff ∀P (x believes P iff Def P)) ⊃ (God is omniaccurate iff (VP God believes P iff Def P))), since the claim within the scope of the main Def operator is a logical truth). We can now derive (24C) Def (God is omniaccurate iff (∀P God believes P iff Def P)) from (24B) and (23) (using (25) and modus ponens), and in turn (24) Definitely ∀P (God believes P iff Def P) (from (24C) and (24A) using (25) and modus ponens). A derivation of (24) from (22), the claim that Def God is omniscient, is similarly mechanical and straightforward.

[24] Or at any rate, our natural way of envisaging the Sorites engagement is one where we suppose that slipperiness plays no essential role in the story. We shall examine in due course whether this supposition can be maintained.

[25] Let 'noon' stand for 'noon Greenwich mean time'. I assume that it is a sharp matter whether a certain time is noon Greenwich mean time.

[26] There is no special worry about the universal instantiation step – see, by analogy, note 24.

[27] We have 28A Def ((God utters 'Dave is not bald yet' at noon on Tuesday iff Dave is not bald yet at noon on Tuesday) ⊃ (∼God utters 'Dave is not bald yet' at noon on Tuesday iff ∼Dave is not bald yet at noon on Tuesday)) since the claim within the scope of the Def operator is a logical truth. From (28A) and (28) using (25) we get (28B) Def (∼God utters 'Dave is not bald yet at noon on Tuesday' iff ∼Dave is not bald yet at noon on Tuesday). Applying (25) to (28) and (28B) we get (28C) Def (God utters 'Dave is not bald yet at noon on Tuesday') iff Def (Dave is not bald yet at noon on Tuesday), and also, (28D) Def (∼God utters 'Dave is not bald yet at noon on Tuesday') iff Def(∼Dave is not bald yet at noon on Tuesday). It is a logical truth that 28E (((Def God utters 'Dave is not bald yet' at noon on Tuesday or Def ∼God utters 'Dave is not bald yet' at noon on Tuesday) and (Def God utters 'Dave is not bald yet' at noon on Tuesday iff Def Dave is not bald yet on noon on Tuesday) and (Definitely ∼God utters 'Dave is not bald yet' at noon on Tuesday iff Def ∼Dave is not bald yet at noon on Tuesday)) ⊃ Def Dave is not bald yet at noon on Tuesday or Def ∼Dave is not bald yet at noon on Tuesday))). From 26, 28C, 28D and 28E, by modus ponens, we get (29).

[28] Notice that (unlike the point made in 3.2), this point depends simply on the assumption of borderline cases and some simple rules for the definiteness operator, not on any other details of supervaluationist semantics.

[29] As a rough model for such a penumbral connection, we might think of God's mind like a mirror. A borderline case of baldness is reflected in the mirror. The claim that God's mind contains a belief that the person is bald is treated by analogy with the claim that baldness is reflected in the mirror. For consider: the precisifications on which the case is one of baldness will be precisifications on which it is also true that baldness is reflected in the mirror.

[30] And about what 'Yes' means and so on. Admittedly, it is a slightly tricky matter to specify the relevant restriction. I shall not press the concern here.

[31] I note in passing that extra complications are created in connection with this suggestion by the dynamical conception of vague predicates enunciated, inter alia, by Delia Graff, Hans Kamp and Scott Soames. (See Graff, "Shifting Sands: An Interest-Relative Theory of Vagueness," *Philosophical Topics* 28 (2000), pp. 45–81; Kamp, "The Paradox of the Heap," in Monnich (ed.), *Aspects of Philosophical Logic* (Reidel, 1981), pp. 225–277; and Soames, *Understanding Truth* (Oxford University Press, 1999), chapter 7. I lack space to take account of such views here.

[32] I am very grateful here to conversations with Williamson.

[33] Note that our supervaluationist would not tolerate the idea that the "Heap" verdicts would give out at exactly the point at which "No Heap" verdicts kick in. Assuming the second conception, we can be secure that a definitely omniscient being will not give either verdict in a clear borderline case.

[34] Even perhaps, by a human, if the subject matter is sufficiently restricted.

*Department of Philosophy*
*Davison Hall*
*Rutgers University*
*New Brunswick, NJ 08901*
*USA*
*E-mail: jphawtho@rci.rutgers.edu*