

The Problem of Vagueness

The topic of vagueness has been the focus of intense philosophical debate over the last four or five decades. There have been numerous articles and books on the topic, the usual back and forth, and the usual lack of consensus at the end of it all.

But when I tell non-philosophers—or even fellow academics—that I am working on vagueness, they are often surprised. They seem to have no idea why the topic could be of any interest or significance. If you tell them that you are working on the problem of free will, of how we can be free in a deterministic universe; or the problem of skepticism, of how we can know anything about the external world on the basis of experience; or the problem of consciousness, of how there can be consciousness in what appears to be a purely physical universe, then you can expect the interest and the importance of these problems to be immediately apparent to them. But with vagueness this is not so. Vagueness, they seem to think, is something to be ignored or ameliorated but certainly not something to be studied.

There are a number of reasons I believe the topic to be worthy of study. But before going into this, let me explain what philosophers have meant by vagueness. In ordinary parlance, the term “vague” can mean a number of different things. When your annoying guest says he intends to stay “one or more weeks” he is being vague in the sense of nonspecific. Is it one week, two weeks or, God forbid, even longer? Or when an ardent advocate of evolution says that “evolution is a fact,” they are being vague in the sense of indistinct. By “evolution,” do they mean evolutionary theory or evolutionary history? Or when a birther says that the place of Obama’s origin is vague, he may simply mean that it is uncertain or unconfirmed.

None of these senses is the philosopher’s sense. When a philosopher talks of vagueness he has in mind a certain kind of indeterminacy in the relation of something to the world. Thus the predicate “bald” is vague because it is indeterminate who exactly is bald. If you line up some men, starting with one who is completely bald, then proceed by gradual increments to one with a full head of hair, then it is to some extent indeterminate which of them is bald. Or again, the location of the summit of Mount Everest is vague since it is indeterminate where exactly it begins or ends. If I were to climb Mount Everest, then it would be indeterminate when exactly I would have reached or would have left the summit (assuming I were ever to get that far).

Thus when philosophers talk of vagueness, they are interested in the kind of indeterminacy that is characteristic of the indeterminacy in the application of such terms as “bald” or “heap” or in the boundaries of such objects as Mount Everest

or a cloud. Vagueness, on the face of it, can reside in both language and the world. However, in what follows I shall mainly focus on vagueness in language, although I do hope to say something later about vagueness in the world.

Vagueness, as so understood, is largely of interest because of the issues to which it gives rise. These issues are broad and profound in their scope, but they have also turned out to be very hard to resolve. The problem of vagueness, then, is the problem of finding a satisfactory resolution to these various issues.

We may place these issues under two main heads, which I call “The Soritical Problem” and “The Semantical Problem.” I do not want to suggest that these are the only issues raised by the phenomenon of vagueness and later I shall suggest some others. However, they are central to our understanding of vagueness and no account of vagueness can be considered satisfactory until it makes clear how they are to be resolved.

What I would therefore like to do in the rest of the present chapter is to say what these two issues are, to consider some of the principal ways in which philosophers have attempted to deal with them, and to explain why I consider their responses to be unsatisfactory. In the remaining two chapters, I shall then attempt to develop and defend an alternative account which I hope is free from some of the defects of the other accounts.

Let us begin with the Soritical Problem. Vagueness gives rise to paradox, to an apparent breakdown in reason, in which we appear to be able to derive a contradiction from impeccable premises by means of impeccable forms of reasoning. The most

famous of these paradoxes is the paradox of the heap—the so-called sorites from the Greek for “heap.” But given the male preoccupation with cranial hairiness, it has been more common to pose the argument in terms of baldness rather than heaps. So imagine a group of men lined up before us. The first, Fred, has a completely hairless head; the last, Les, has a full head of hair; and in between is a series of men, each of whom has an imperceptibly greater amount of cranial hair than his neighbor.

Now surely Fred (the first man) is bald and Les (the last man) is not bald. It is also very plausible that if a given man in the lineup is bald then so is the man to his immediate right. For one thing, there is no perceptible difference between the two of them. Look as hard as you like, you would not be able to see any difference in the degree to which they were bald. But being bald is a matter of how someone looks; and so if the one man is bald and looks the same as the next man, then so is the next man.

For another thing, a minuscule difference in cranial coverage would not appear to make any difference to whether someone is bald. Imagine a vain man, call him “Donald,” who is worried about whether he is bald. He looks in the mirror one morning and says to himself “damn, I’m bald.” He is then told by his friendly doctor, Harold, that a certain product, Hair-On, will improve his cranial coverage to an imperceptible degree. Does he now have an incentive to take the product? Surely not, since a minuscule change in cranial coverage can make no difference to whether he is bald.

We now have the wherewithal to state the paradox. Fred is bald. But given that there is no relevant difference between him

and the second man, the second man is bald. There is likewise no relevant difference between the second man and the third, and so the third man is bald. Continuing in this way, we arrive at the conclusion that Les is bald. But he is not bald. Just look!

What makes this a paradox is that we appear to have derived a contradiction from seemingly impeccable assumptions by means of a seemingly impeccable pattern of reasoning. The assumptions are the two “extremal” premises, that Fred is bald and that Les is not bald; and the “transitional” premises, each to the effect that if a given man in the lineup (before Les) is bald then so is the next man. The pattern of reasoning consists of successive applications of the rule of modus ponens, according to which from premises of the form “p” and “if p then q” one can validly infer a conclusion of the form “q.” Thus given that Fred is bald and that if Fred is bald then the second man is bald, we can infer that the second man is bald by modus ponens; given that the second man is bald and that if the second man is bald then the third man is bald, we may infer that the third man is bald; and proceeding in this way, we may eventually draw the conclusion that Les is bald, in contradiction to evident fact that he is not bald.

We must somehow have made a mistake. Either we should not have accepted one of the premises of the sorites argument or one of the instances of modus ponens or the successive instances of modus ponens by which the contradiction was derived.

Here then is the first of our two issues. What goes wrong with the sorites argument? And given that something goes wrong, why are we so readily taken in by the argument?

We turn to the second of the two issues, the Semantical Problem. Philosophers and linguistics alike are interested in providing a semantics for natural language and for various formal languages. A semantics for a given language is a systematic account of how the expressions of the language derive their meaning. Such a semantics may take various forms but the most common form is one in which semantic values, the formal counterpart of meanings, are assigned to the meaningful expressions of the language. The simple expressions of the language are simply stipulated to have an appropriate semantic value; and rules are then given for determining the semantic value of complex expressions on the basis of the simpler expressions from which they are composed.

Many philosophers suppose that a precise language, one composed entirely of precise expressions, should be given a classical semantics. This, at its simplest, is a semantics in which the semantic value of a name is its referent, the semantic value of a predicate is its extension, i.e., the set of objects of which it is true, and the semantic value of a sentence is one of the truth-values, True or False. There are then some obvious rules for assigning semantic values to complex expressions. So, for example, a subject-predicate sentence Fa will be assigned True when the referent of the name a belongs to the extension of the predicate F and will otherwise be assigned False; and a negative sentence $\neg A$ will be assigned True when A is assigned False and assigned False when A is assigned True.

Given a semantics for a given language, one will then be in a position to say when an inference stated in the language is valid. Thus in the case of the classical semantics, we can say that an inference from the sentences A and B to the sentence C will be valid if C takes the value True whenever A and B take the value True under any appropriate assignment of semantic values to the simple expressions of the language. To give a simple illustration, the inference from $\neg\neg A$ to A will be valid under the classical semantics since, given the rule for negation, $\neg\neg A$ can only take the value True when A takes the value True.

We now face the problem of providing a semantics for a vague language, one partly composed of vague expressions. Suppose, for example, that our language is one in which we can talk about the baldness of men of varying cranial disposition. Then what kind of semantics should we give for a sentence such as “Max is bald,” where Max is someone in the middle of a sorites series for the predicate “bald,” and what rule should we adopt for negation?

One might think that the answer to this question was straightforward. For why not just take over the classical semantics for a precise language and apply it to a language that is vague? Thus we can say, in particular, that a predicate, such as “bald,” has a certain extension and that the sentence “Max is bald” takes the value True when Max is in the extension of the predicate “bald” and otherwise takes the value False.

However, it is a presupposition of the classical account that every sentence be either true or false, the so-called Principle of

Bivalence. It was for this reason that the negative sentence $\neg A$ was taken to be false (or have the truth-value False) when A was not true and that a subject-predicate sentence Fa was taken to be false when the referent of a was not in the extension of F . But the Principle of Bivalence does not sit easily with the idea that the predicate “bald” is indeterminate in its application to the members of a sorites series. For how can the application of the predicate be indeterminate if it is either true or false to say of anyone of them that he is bald? And how, in this case, would a vague predicate, like “bald,” differ from a precise predicate, such as “electron”?

This then is the second of our two issues. How, if at all, should the usual classical semantics for a precise language be modified so as to allow for the presence of vague terms? The two issues are connected. For, as we have seen, a semantics for a vague language should deliver an account of valid inference and, on the basis of that account, we should then be in a position to say where the soritical reasoning goes wrong. Thus a satisfactory account of vagueness should, in this way, provide a unified approach to the semantics and logic of vague language.

I would like in the remainder of the chapter to discuss three lines of solution to the two issues we have raised. I shall call them Degree-ism, Supervaluationism, and Epistemicism, though I should apologize right away for thrusting such horrible sounding isms on you. These three lines of solution are by no means the only ones that have been proposed, and my discussion of them will be somewhat superficial and far from

complete. But I hope to say enough to give the reader a feel for the various positions and how they might be found wanting.¹

The first line of solution, Degree-ism, is perhaps the simplest and most natural. A precise predicate will be true or false of any object in its range of application. We may take it to be characteristic of a vague predicate, by contrast, that its application will not in general be bivalent; there will be objects of which it is neither true nor false. The predicate “bald,” for example, will presumably be neither true nor false of certain men in the middle of our lineup.

We therefore posit further truth-values “intermediate” between the extreme values of Truth and Falsehood. Under the simplest version of Degree-ism, there will be one such truth-value, the Indefinite, and so the sentences of our language will now be capable of taking three truth-values, True, False, and Indefinite, in place of the two classical values, True and False. The classical rules for determining the truth-value of complex sentences must now be extended to take account of the third truth-value. Thus under the most natural rule for negation, a negative sentence $\neg A$ will take the value Indefinite when, and only when, A takes that value. More elaborate versions of Degree-ism may posit varying degrees of truth ranging from False, or most false, to True, or most true, through small—perhaps even

1. An early version of Degree-ism is Goguen [1969], and a more recent version, attempting to deal with various objections, is Smith [2008]; supervaluationism was developed in Fine [1975] and has recently been defended, in relation to the other views, by Keefe [2000]; presentations of epistemicism are to be found in Sorensen [2001] and Williamson [1994]; and a useful general reader is Keefe & Smith [1997]

continuously small—increments; and, in such a case, we should provide corresponding rules for determining the degree of truth of complex sentences in terms of the degree of truth of its component sentences. Thus if degrees of truth are taken to be real numbers between 0 and 1 (with 0 corresponding to False and 1 to True), then we might take the degree of truth of the negative sentence $\neg A$ to be 1 minus the degree of truth of A .

Even if this view provides an acceptable account of the truth-value of simple subject-predicate sentences, it is far from clear that it is able to provide a satisfactory account of the truth-value of more complex sentences. Take a patch, called Pat, on the border between red and orange and let us suppose that the sentences “Pat is red” and “Pat is orange” are both indefinite. What truth-value should we assign to the conjunctive sentence “Pat is red and Pat is orange”? Intuitively, we would like to say that the sentence is false. After all, red and orange are exclusive colors. This means that the rule for conjunction should say that the conjunction of two indefinite sentences is false. But this then has the consequence that the sentence “Pat is red and Pat is red” should also be false. Yet surely it should have the same truth-value, Indefinite, as the sentence “Pat is red.”

This is the so-called problem of “penumbral connection.” Pat lies on the penumbra of “red” and “orange”; it is a borderline case of both predicates. But even though Pat lies on the penumbra, there can still be logical connections between its being red and its being orange; the one can exclude the other. And it is unclear how the degree-theorist is able, within the degree-theoretic framework, to explain such connections.

I turn to the second view, Supervaluationism. Take a vague predicate, such as “bald.” We suppose, as with the degree-theorist, that there may be borderline cases of the predicate, cases of which the predicate is neither true nor false. However, under the present approach, we take account of something else: the different ways in which we might acceptably make the predicates of our language completely precise. There are three important technical terms here—“precise,” “completely,” and “acceptably.” In making the predicates of our language *more precise*, we leave the clear, or non-borderline, cases as they are and change the status of only the borderline cases, changing some or all of them from borderline cases to clear cases. In making the predicates of our language *completely* precise, we make them more precise in such a way as to leave no borderline cases—each borderline case becomes a clear case. We also insist that the way we make the predicates more precise should be *acceptable* in the sense of being in conformity with our intuitive understanding of the predicates.

Let us illustrate with our lineup. Suppose, for simplicity, that there are just two borderline cases: Max, who is the less hairy of the two, and his neighbor Ned. There are then three acceptable ways of making the application of the predicate “bald” to the men in the lineup completely precise. We can take both Max and Ned to be bald, take both not to be bald, or take Max to be bald and Ned not to be bald. There is one unacceptable way to make the predicate completely precise, with Ned, the more hairy, bald and Max, the less hairy, not bald.

Given any way of making the predicates of the language completely precise, we can provide a classical semantics for the

language in the usual way. We can say, for example, that the sentence “Max is bald” is true under the first and third of the three ways above for making the predicate completely precise but false under the second of the three ways.

But this is only to assign truth to the sentences of the language relative to some particular way of making it completely precise. We now say, under the supervenient approach, that a sentence is true simpliciter if it is true for all acceptable ways of making it completely precise and that it is false simpliciter if it is false for all acceptable ways of making it completely precise. So, for example, “Fred is bald” will be true since it is true under all acceptable ways of making the predicate completely precise, “Les is bald” will be false since it is false under all acceptable ways of making the predicate completely precise, and “Max is bald” will be neither true nor false since it is true for some of the acceptable ways of making the predicate precise and false for others.

Truth for the degree-theorist is truth from “below,” since we compute the truth-value of a complex sentence from the truth-values of its component sentences. Truth for the supervenient, by contrast, is truth from “above,” since it is only by first looking up at the different ways of making the language completely precise that we are in a position to say whether the sentence is true or false in the original unprecisified language.

One advantage of Supervenientism over Degree-ism is that it is able to account for penumbral connection. Consider again our patch that was on the border of red and orange. We

wanted the sentence “Pat is red and orange” to be false, which we could not have under the degree-theoretic approach. But note now that no way of making the predicates “red” and “orange” precise will be acceptable if it renders an object both red and orange. Thus under any acceptable way of making these predicates completely precise, the sentence “Pat is red and orange” will be false; and so the sentence will be false simpliciter, just as we wanted.

However, the approach suffers from problems of its own. Under this approach, it will be true to say that there is a last bald man, one who is preceded by men who are bald and succeeded by men who are not bald. For under any acceptable way of making the predicate “bald” completely precise, we must draw a line between the men who are bald and those who are not bald and, from among the men who are bald, there will always be one who is last. But given that there is a last bald man in the lineup, can we not legitimately ask “who is it?” And if we can legitimately ask such a question, then should it not be correct to say of some particular man that *he* is the last to be bald. But there is no such man. For whichever man we choose, there will always be some acceptable way of making some other man be the last man to be bald. Thus it looks as if, under the supervaluational approach, questions which should have answers will have no answers.

I turn to the last of the three views, Epistemicism. The two previous accounts took vagueness to require the failure of Bivalence, the principle that every sentence should be true or

false, and hence called for a modification to the classical bivalent semantics. The epistemicist disputes that there is a failure of Bivalence in the case of vagueness: every sentence that expresses some content, even a vague sentence, will be true or false. He will therefore be willing to say that the sentences of a vague language have exactly the same semantics as the sentences of a precise language. Vagueness makes no difference to the semantics. By the same token, the logic appropriate to a vague language will be the same as the logic for a precise language. Vagueness makes no difference to logic.

But what then for the epistemicist is the difference between vague and precise terms? And what, in particular, is it for a predicate to be indeterminate in its application to a range of objects?

Before, under the degree-theoretic and supervaluational approaches, we took indeterminacy to be a semantic matter, a gap in truth-value or meaning. The epistemicist now takes indeterminacy to be an epistemic matter, a gap in our knowledge. When a predicate indeterminately applies to some objects, it truly or falsely applies to those objects; it is just that we do not know to which objects it applies.

But the gap is not any old gap in our knowledge. I do not know how many people are in the room right now, but not because of any vagueness in the expression “number of people in the room.” In this case I know what it would take for there to be a 100 people in the room, say, rather than 101. The thought is that, in the case of vagueness, we are ignorant of what it would even take for the predicate to have application in any

given case; we are ignorant of the *criteria* for its application. There will indeed be a complete and precisely defined criteria for when the predicate “bald,” say, does or does not apply to any given man; exactly this number of hairs of such and such a length and thickness and distribution, etc., will guarantee his baldness; exactly this number, etc., enough to guarantee non-baldness; and so on, through all the different possibilities. And yet we do not know what the criteria are.

The epistemic view has certain advantages over the gap-theoretic and supervaluational views. Like the supervaluational approach, it can account for penumbral connection. Indeed, on the epistemic view there is, strictly speaking, no penumbral connection, since there are no penumbral cases to connect. But what are commonly regarded as penumbral truths will still be true. It will be true, for example, that Pat is not both red and orange since one of the two predicates, “red” and “orange,” will in fact be true of Pat and the other not.

The epistemic view, in contrast to the supervaluationist view, will not suffer from the existence of unanswerable questions. The epistemicist, like the supervaluationist, will want to say that there is a last bald man in our lineup. If now we ask, “but who is he?” then there will indeed be a correct answer: we can indeed truly say of one of the men in the lineup that he is the last to be bald. It is just that we are in no position to know that what we say is true.

Despite its merits, many philosophers (myself included) have thought this view to be unbelievable. There are meant to be some facts—presumably facts of usage—that determine an

absolutely precise criterion for being bald. But how could they do this? We simply have no idea.

We can see Epistemicism as the result of a common—and what I believe is usually a misguided—philosophical tendency to identify a cause with its symptoms. We are presented with some underlying phenomenon, and because we are not sure how it should be characterized, we are tempted to identify it with the symptoms by which it is made manifest. Thus we are not sure how to characterize the mental and so we identify it with the corresponding physical behavior or we are not sure how to characterize a law and so we identify it with the regularities to which it gives rise. In the present case, we are not sure how to characterize vagueness and so we identify it with the resulting ignorance rather than attempting to get at the underlying phenomenon by which the ignorance might be explained.

Each of the views that we have considered has certain advantages and disadvantages over the others. But they also have some common failings. One of these, it seems to me, is that they are incapable of providing a satisfactory solution to the sorites paradox. Consider a degree-theoretic approach in which the degrees of truth are the real numbers between 0 and 1, with 0 being Falsity and 1 being Truth. The degree of truth of a major premise such as the conditional premise $p_k \supset p_{k+1}$ will sometimes be a little less than 1 (since the truth-value of p_{k+1} will be a little less than the truth-value of p_k). Our mistake, according to the degree-theorist consists in treating as true a statement that is not true but almost completely true.

But suppose we had used $\neg(p_k \wedge \neg p_{k+1})$ (it is not the case that the k -th man is bald and yet the $(k+1)$ -th man is not bald) in place of $p_k \supset p_{k+1}$, as the major premise. Then when the truth-values of p_k and p_{k+1} are close to $1/2$, the truth-value of $\neg p_{k+1}$, and hence the truth-values of $(p_k \wedge \neg p_{k+1})$ and of $\neg(p_k \wedge \neg p_{k+1})$ will also be close to $1/2$ —at least on standard ways of computing these truth-values. But we are as much, if not more, inclined to regard $\neg(p_k \wedge \neg p_{k+1})$ as true as $(p_k \supset p_{k+1})$. The degree-theorist seems unable to offer any plausible explanation of why this is so; and this suggests that their explanation of our inclination to accept the conditional premises is also mistaken.

The supervaluationists and epistemicists have what is perhaps an even more serious problem. Both take it to be true that there is a sharp cutoff, with a given man in the series bald and the man next to him not bald. But then why are we so inclined to think otherwise?

It is often supposed that we somehow confuse the statement that some man is bald while his neighbor is not with the statement that some man is determinately bald while his neighbor is determinately not bald. But why should we be inclined to reinterpret the statement in this way? One can imagine reinterpreting a statement that was false so as to be true. But why reinterpret a statement (that there is a sharp cutoff) that is true so as to be false? Nor is it that, once we take care to read the statement so as not to presuppose the presence of an implicit determinately-operator, we are somehow freed from any inclination to regard it as true.

There is another common failing, far more sweeping in its scope. But to explain what it is, I should first make a distinction that has so far only been implicit in what I have said. I have loosely talked of indeterminacy in the application of a predicate to some objects. But there are two kinds of indeterminacy that may be in question—*local* and *global*. Local indeterminacy is indeterminacy in the application of the predicate to a single object. We have a man, say, in the middle of our lineup and it is indeterminate whether or not he is bald. This notion of indeterminacy is just the same as our previous notion of a borderline case; for something to be a borderline case of a predicate is simply for the predicate to be locally indeterminate in its application to that object.

Global indeterminacy, by contrast, is indeterminacy in the application of the predicate to a *range* of cases. The term “range” is important here; we have not a single case, but a number of cases, two at the very least. Thus we may say in this sense that the predicate “bald” is indeterminate in its application to the men in our lineup. Even though the predicate may be determinate in its application to some of the men, it is not completely determinate in its application to all of the men.

The “common failing” to which I alluded concerns the relationship between global and local indeterminacy; for there appears to be no satisfactory way, under existing views of vagueness, of explaining how they relate. In arguing for this conclusion, I will need to make use of two assumptions, which I call the Compatibility Requirement and the Incompatibility Requirement.

Suppose I am presented with a sorites series for the predicate “bald” and I consider, for each of the men in the series, whether or not he is bald. I can either say “yes” or “no” or refrain from giving an answer. This is a so-called forced march.

The Compatibility Requirement states that a global claim of indeterminacy should be compatible with the minimal response in which I provide a positive answer in the first case and a negative answer in the last case while refraining from giving any answer in the other cases. There should be no coherence in my asserting both that the predicate “bald” is not completely determinate in its application to the men in the series while asserting that the first man is bald and that the last man is not bald.

The second assumption is a little harder to state. Suppose again that one is presented with a forced march but that in this case one either gives a positive answer “Yes, he is bald” or a negative answer “No, he is not bald” to each of the questions. Where there are 25 men, for example, one might respond “Yes” to the first 12 questions and “No” to the remaining 13 or perhaps “Yes” to the first 13 questions and “No” to the remaining 12. In such a case, there would surely be some kind of incompatibility or incoherence in giving these answers and yet going on to make a global claim of indeterminacy. To draw a line in this way between the men who are bald or not bald is implicitly to concede that the predicate “bald” is *not* indeterminate in its application to the men in the series.

But something more general would also appear to hold. For suppose one were to respond to a forced march by saying

that each of the first nine men was not merely bald but determinately bald; that each of the next three men was borderline bald, i.e., neither determinately bald nor determinately not bald; and that each of the remaining men was not merely not bald but determinately not bald. Then presumably this would still be incompatible with a global indeterminacy claim. For a sharp line is still being drawn, not now between the men who are bald and the men who are not bald, but between the men who are *determinately* bald and the men who are *borderline* bald and, in addition, between the men who are *borderline* bald and the men who are *determinately* not bald. And the existence of sharp lines at this “higher” level would appear to be as much in conflict with a claim of indeterminacy as the existence of a sharp line at the “lower” level.

What goes for sharp lines at this higher level would appear to extend to sharp lines at higher levels still. It would not do, for example, to respond to each question within a forced march with the response that the man is determinately bald to the n th degree or that he is borderline bald to the n th degree or that he is determinately not bald to the n th degree.

The more general point is this. Consider any complete set of responses to a forced march—such as “Yes, . . . , Yes, No, . . . , No” or “Determinately Yes, . . . , Determinately Yes, Borderline, . . . , Borderline, Determinately No, . . . , Determinately No.” Call such a series of responses *sharp* if it draws a contrast between at least two neighboring cases. Then a claim of indeterminacy should exclude any sharp

response to a forced march; it should not be possible to make the indeterminacy claim compatibly with giving a sharp response. As Sainsbury [1989] puts it, “Vague concepts are concepts without boundaries.”

We can now state an Impossibility Result:

Impossibility: No putative claim of indeterminacy can meet the Compatibility and the Incompatibility Requirements.

In other words, no indeterminacy claim, whatever form it might take, can be both compatible with the minimal response to a forced march and yet incompatible with any sharp response. Vagueness would therefore appear to be impossible insofar as there is nothing that can meet the demands upon which its existence would appear to depend.

This is actually a *theorem*, susceptible of mathematical formulation and mathematical proof. I cannot provide a precise statement or proof of the theorem here.² But what the proof does, under the assumption of the Compatibility Requirement, is to construct a sharp response using iterations of the determinately operator. Thus the sharp response, in a particular case, might be that, for the first 20 men, it is determinately, determinately, determinately the case that they are bald, that, for the 21st man, it is neither determinately, determinately, determinately the case that he is bald nor determinately, determinately,

2. A proof of a simple version of the result is given in Appendix A and a more elaborate version is to be found in Fine [2008].

determinately the case that he is not bald, and that, for the last 19 men, it is determinately, determinately, determinately the case that they are not bald.

There are a number of ways, under existing views of vagueness, by which one might attempt to evade this result. One might claim, for example, that only incompatibility with a low level sharp response is required. However, even if there are forms of global indeterminacy that are compatible with a high-level sharp response, surely there are also forms of global indeterminacy that are not. Or again, it is a presupposition of the proof that if one is willing to assert that Fred, say, is bald, then one should also be willing to assert that it is determinately the case that Fred is bald, determinately determinately the case, and so. This too might be questioned though it is hard to see why, in the presence of a completely bald man, one should not be willing to assert that he is determinately bald to the n th degree. My own view is that none of these responses to the impossibility result will ultimately stand up; and, if this is so, then all of the existing approaches to vagueness should be abandoned and some other way of evading the result should be found.