Vagueness, Logic, and Ontology

Achille C. Varzi
Department of Philosophy, Columbia University, New York

(Published in *The Dialogue. Yearbooks for Philosophical Hermeneutics* 1 (2001), 135–154)

Introduction

Remember the story of the most-most? It's the story of that club in New York where people are the most of every type. There is the hairiest bald man and the baldest hairy man; the shortest giant and the tallest dwarf; the smartest idiot and the stupidest wise man. They are all there, including honest thieves and crippled acrobats. On Saturday night they have a party, eat, drink, dance. Then they have a contest. "And if you can tell the hairiest bald man from the baldest hairy man—we are told—you get a prize."

The story is from Saul Bellow's *Herzog*, a novel published almost forty years ago but still very modern in its philosophical provocations.[1] It is a funny and provoking story because so is the idea of a contest like that of the most-most. There is no sharp boundary demarcating the category of bald men, no precise number of hairs separating the bald from the hairy. Hence it makes no sense to suppose that one can identify the hairiest bald man. It does not even make sense to suppose that such a person exists, as if the difficulty were merely epistemic. Some people are clearly bald and some are clearly hairy, but between these two sorts of people there exist a variety of borderline cases: baldish guys, men wearing toupees, hirsute beatniks with a shiny spot on the top of their heads. Our concept of baldness and our linguistic practices do not specify any precise, general criterion for saying in each case whether we are dealing with a bald man or with a hairy one. 'Bald' and 'hairy' are vague concepts and their ranges of application have vague boundaries.

Shall we distinguish three categories, then? Some people are clearly bald (Picasso), some are clearly hairy (the count of Montecristo), and some are borderline cases. That's a third category of its own, and we may suppose that the New York club includes the most of each of these three categories.

Unfortunately things are not so easy. Two boundaries aren't any better than one. And if it is impossible to identify the hairiest bald man (i.e., to draw a boundary

[1] See Bellow [1964], pp. 295–296.

between bald and non-bald), it is also impossible to identify the first *clear instance* of a bald man (i.e., to draw a boundary between clearly-bald and borderline-bald). The vagueness of these concepts does not reduce to the existence of borderline cases: it is the absence of a sharp boundary in their range of application that makes them vague. Imagine we find ourselves in a room with the count of Montecristo and suppose we start plucking his hairs, one at a time. At the beginning of the process the count is pretty hairy. At the end he will be bald. But when exactly will he cease to be hairy? When will he begin to be bald? These are questions that we could not answer even if we were omniscient.

Increasing the number of intermediate categories won't do, either. We could distinguish between borderline cases and borderline borderline cases, or borderline borderline borderline cases, but things would only get worse. Multiplying the number of relevant boundaries amounts to making an even stronger commitment to precision than that of the members of the most-most club. And this is a serious problem because it gives rise to a genuine logical puzzle. Gradually, but ineluctably, we find ourselves in a situation that appears to be contradictory. On the one hand, we find it natural to agree with the following two statements:

(1) Upon removing 1 hair, the count of Montecristo is still hairy.
(2) For every n: if the count of Montecristo is still hairy upon removing n hairs, then he is still hairy upon removing $n + 1$ hairs.

(What difference can a single hair make?) On the other hand, we certainly want to deny the statement

(3) Upon removing all hairs, the count of Montecristo is still hairy.

Yet (3) follows logically from the conjunction of (1) and (2). (This can be shown by repeated applications of the rules of universal instantiation and *modus ponens*.) Hence, either we reject one of the premises, (1) or (2), contrary to our intuitions about the meaning of 'hairy'; or we find ourselves forced to give up some elementary logical principles so as to block the inference to (3). Either way, the picture is worrisome.

Vagueness and the Normativity of Logic

This dilemma was known, in some form or other, already to the ancient world. (The first version of the paradox is attributed to Eubulides of Miletus.[2]) More important,

[2] For a history of the paradox and its numerous variants, see Sainsbury & Williamson [1997] and Keefe & Smith [1997].

it was well known also to Gottlob Frege, one of the fathers of modern symbolic logic. For Frege the rigorous delimitation of every concept was one of the fundamental provisos for the possibility of applying the rules of logic:

> A definition of a concept (of a possible predicate) must .... unambiguously determine, as regards any object, whether or not it falls under the concept (whether or not the predicate is truly assertible of it). Thus there must not be any object as regards which the definition leaves in doubt whether it falls under the concept .... We may express this metaphorically as follows: the concept must have a sharp boundary.[3]

Indeed, beginning with Frege and for a very long time, logicians have been assuming that vagueness may bring logical disaster in its wake. Vague concepts have been regarded as illegitimate, at least to the extent that logic is understood as a normative discipline (rather than a descriptive discipline, like psychology). And this, in turn, has been taken to imply that vagueness should be eliminated altogether from the realm of logic.

The trouble is that this way of proceeding is itself very problematic. For vagueness is an extremely pervasive phenomenon. As Michael Dummett put it, it gets into everything, like dust.[4] Perhaps the concepts employed in mathematics and some concepts used in the so-called exact sciences are exempted from its grasp. But the vast majority of the concepts that we use in ordinary discourse are vague. And this is not only true of adjectives like 'bald' and 'hairy' and 'short' and 'wise'. The same goes for concepts that find expression in many other grammatical categories such as nouns (how many grains of sand does it take to make a *heap*?), verbs (how slowly can you *run*?), adverbs (how *fast* can you drive?), and so on. So the problem is: on the one hand, vagueness leads to a logical paradox; on the other hand, we applied Frege's criterion and protected logic by eliminating all those words and those concepts which exhibit some degree of vagueness, then we would end up with a language that is so poor as to be utterly useless. Moreover, it is not even clear whether the criterion is actually applicable. It would be applicable (at least in principle) if there existed a sharp boundary between the vague and the non-vague. But even this may turn out to be an unwarranted presumption: some concepts are clearly vague (*bald*) and some are clearly not vague (*circular*), but there are concepts that appear to lie at the borderline. Is there a precise moment at which a person ceases to be *alive*? Is there an exact moment at which a woman becomes a *mother*? Are these concepts vague or are they precise?

---

[3] Frege [1903], §56.

[4] Dummett [1995], p. 207.

Even without taking such problematic cases into account, the vagueness of our concept of vagueness can be appreciated by construing a paradoxical argument that makes 'vague' similar to such predicates as 'bald' and 'hairy', as Roy Sorensen has pointed out.[5] Consider, for example, the following series of concepts: 'either bald or with less than 1 hair on his head', 'either bald or with less than 2 hairs on his head', ..., 'either bald or with less than 500.000 hairs on his head'. (We may suppose that 500.000 is approximately the number of hairs on the head of the count of Montecristo at the time when he enters the room in our earlier experiment.) Evidently, the first term of this series is vague: it is as vague as 'bald', since both are clearly true of hairless people, and for all other people they are exactly alike. Moreover, there is no point in the series where we encounter a vague predicate followed by a sharp one: what difference can a single hair make? So we find it natural to agree with the following two statements:

(4) The adjective 'either bald or with less than 1 hair on his head' is vague.

(5) If the adjective 'either bald or with less than n hairs on his head' is vague, then so is 'either bald or with less than n+1 hairs on his head'.

On the other hand, we certainly want to deny the statement

(6) The adjective 'either bald or with less than 500.000 hair on his head' is vague.

For the last term of our sequence is clearly not vague. It is, in fact, just as precise as the predicate 'with less than 500.000 hairs on his head', and this is true of all men with less than 500.000 hairs on their heads and false of all others. As before, however, the statement which we want to deny follows logically from those which we want to assert. And if we take this situation as a sign of vagueness, we must conclude that 'vague' is just as vague a concept as 'bald' and 'hairy'.

The moral is that it is impossible to draw a line between vague and precise concepts just as it is impossible to draw a line between bald and hairy people. Hence the strategy advocated by Frege is ultimately inapplicable. Nor can we just assume that the precise languages postulated by classical logic are idealizations to be realized in the future, for our argument shows that vagueness dies hard. If we actually try to stipulate it away, our stipulations will themselves be made in less than perfectly precise terms and the regimented language will inherit some of that vagueness. So if vagueness is incompatible with classical logic, classical logic is ultimately inapplicable to ordinary thought and language. So much for the normativity of classical logic.

---

[5] In Sorensen [1985].

Vague Reference

The pervasiveness of vagueness goes even farther than this. We have focused on conceptual vagueness—the vagueness exhibited by general terms such as adjectives, nouns, or verbs. But there can be vagueness also in the case of singular terms, i.e., names or descriptive phrases purporting to refer to individual objects or events rather than concepts or classes. Perhaps the referents of 'Bill Clinton' or 'the mayor of New York' are well defined. But what about such terms as 'Everest', 'Toronto', or 'that cloud in the sky'? Surely the referents of such terms are to some extent indeterminate. And this sort of indeterminacy does not reduce to ambiguity: it bears the mark of vagueness—for example, it gives rise to paradoxical arguments parallel to those defined by (1)–(3) and (4)–(6). You are on the top of Mount Everest and you begin descending. After each step you ask yourself whether you are still on the mountain. At the beginning you have no doubts, and you have no doubts at the end, when you find yourself in the center of Katmandu. But of course there is no point where you can confidently assert: Here is where Everest ends. So again you find yourself in a predicament. You are inclined to agree with the following two assertions:

(7) After 1 step I am still on Everest.

(8) For every n: if I am still on Everest after n steps, then I am on Everest after $n + 1$ steps.

while denying their logical consequence:

(9) At the end of the descent I am still on Everest.

One might argue that in this case the vagueness lies in the predicate 'I am still on', or perhaps in the relational predicate 'after n steps I am still on', but it is easy to reformulate the argument using precise predicates instead. For instance, we can imagine a purely geometric version of the paradox in which the region of space occupied by Mount Everest is compared with a sequence of precisely demarcated spatial regions, each slightly larger than its predecessor, starting from a small region which comprises the peak of the mountain and ending with a very large region that extends all the way to downtown Katmandu. At which point of the sequence shall we say that region R_n has a smaller volume than the volume of Mount Everest while the next region, R_{n+1}, has an equal or greater volume? Evidently the problem is that the term 'Everest' does not refer to a precise chunk of reality. It does not refer to a volume of matter sharply demarcated from its surrounding and it therefore makes no sense to compare the region it occupies with a series of precise regions. Simply, 'Everest' is a vague term. And so are 'Toronto', 'that cloud in the sky', and many other names and descriptions that we use in ordinary discourse.

Indeed, even an apparently innocuous proper name such as 'Bill Clinton' is arguably vague: What exactly are the spatial boundaries of its referent? Surely they comprise Clinton's heart, and surely they do not comprise my left foot. But what about the candy that Clinton is presently chewing: Is it part of Clinton now? Will it be part of Clinton only after he has swallowed it? After he has started digesting it? After he has digested it completely? And exactly when did Clinton come into existence? Exactly when will it be correct to say that Bill Clinton does not exist any longer?

Ontological Vagueness vs. Semantic Vagueness

Vagueness is such a pervasive phenomenon that one can hardly overestimate the threat that it represents to traditional logical and philosophical theories. There is, in addition, dispute concerning the nature of this phenomenon, and different conceptions lend themselves naturally to different ways of coping with it. Consider our statement to the effect that the referent of a vague term *t* is not sharply demarcated. This very statement is ambiguous and admits a *de re* reading, as in (10a), or a *de dicto* reading, as in (10b):

(10) a The referent of *t* is such that it is indeterminate whether certain chunks of reality lie within its boundaries.

b. It is indeterminate whether certain chunks of reality lie within the boundaries of the referent of *t*.

On the first reading the indeterminacy is ontological: vague terms refer to vague objects, objects which lack precise spatial or temporal boundaries. For example, 'Everest' is vague insofar as the boundaries of Everest do not sharply divide the matter composing it from the matter outside it. In Michael Tye's words:

> Everest's boundaries are fuzzy. Some molecules are inside Everest and some molecules outside. But some have an indefinite status: there is no objective, determinate fact of the matter about whether they are inside or outside.[6]

Likewise, on this view there is no objective fact of the matter about whether certain water droplets are inside a cloud, whether a certain drop of rain fell on Downtown Toronto, whether the candy is part of Clinton. Mountains, clouds, neighborhoods, houses, forests, deserts, islands, and perhaps even people, on this view, are all genuinely vague denizens of reality. Like the figures of an impressionist painting, they do

[6] Tye [1990], p. 535.

not fit in the standard topological picture according to which every object has an interior surrounded by an exterior. Like old soldiers, "they just fade away".[7]

On the second reading—the *de dicto* reading—the indeterminacy exhibited by a vague term is exclusively semantic, or cognitive at large. It lies in the representation system (our language, our conceptual apparatus), not in the represented entity, and to say that the referent of a term is not sharply demarcated is to say that the term vaguely designates an object, not that it designates a vague object. When we say 'Everest' (or when the founder of the Indian Geodetic Office baptizes a certain piece of land, at the border between Tibet and Nepal, 'Everest') we simply do not specify exactly which piece of land we are referring to. The referent of our term is vaguely fixed. If we wish, we can add that it is ultimately the vagueness of the relevant sortal concept (the concept *mountain*, in this case) that is responsible for the way in which the referent of 'Everest' is vaguely fixed. But it is not the stuff out there that is vague. Each one of a large variety of slightly distinct chunks of reality has an equal claim to being the referent of our newly introduced name. And each such thing is precisely determinate. To use a different example, from Henryk Mehlberg,

> The term 'Toronto' is vague because there are several methods of tracing the geographical limits of the city designated by this name, all of them compatible with the way the name is used. It may be interpreted, for instance, either as including some particular tree on the outskirts of the city or as not including it. The two areas differing from each other with respect to the spot where this tree is growing are two distinct individual objects; the word 'Toronto' may be interpreted as denoting either of these objects and is for that reason vague.[8]

This opposition between ontological and semantic vagueness applies also in the case of names and descriptions for other sorts of entity. An expression such as 'Sebastian's walk' is a vague event designator: there is indeterminacy concerning the exact spatiotemporal location of the designated event. On the ontological conception this means that the event itself is vague, that its spatial and temporal boundaries are genuinely fuzzy. On the semantic conception, by contrast, 'Sebastian's walk' is vague only insofar as it vaguely designates an event. It's not that there is this event, Sebastian's walk, with imprecise boundaries. There are plenty of things going on inside and outside Sebastian as he moves along the sidewalk, each with a precise location in space and time, and many of them qualify as legitimate referents of the phrase 'Sebastian's walk'. The phrase itself is too vague to discriminate among them.

---

[7] The phrase is from Sylvan & Hyde [1993], p. 19.

[8] Mahlberg 1958, p. 257.

Finally, we can introduce an opposition between ontological and semantic interpretations also in the case of conceptual vagueness, i.e., the vagueness exhibited by predicates such as 'bald' and the like. These are predicates whose extensions do not have sharp boundaries, and the assertion that the extension of a predicate *P* lacks sharp boundaries can be given a *de re* reading, as in (11a), or a *de dicto* reading, as in (11b):

(11) a The extension of *P* is such that it is indeterminate whether certain objects fall within its boundaries.

b. It is indeterminate whether certain objects fall within the boundaries of the extension of *P*.

The *de re* reading yields ontological vagueness. A predicate such as 'bald' would then designate a fuzzy set, a set whose membership function allows for borderline values: some people, such as Picasso, are definitely *in* this set; others, such as the count of Montecristo, are definitely *out*; and others have an indefinite status: there is no objective, determinate fact of the matter about whether they are in or out. By contrast, the *de dicto* reading corresponds to a purely semantic conception of vagueness. The set of bald people is not a vague set at all. There are exactly 2^n sets of people (where *n* is the number of all people at the present time), each with its perfectly precise membership function, but it is indeterminate *which* of those sets can do duty for the extension of the predicate 'bald'. There are several good candidates but the predicate itself is too vague a description to successfully pick out a unique one of them.

Ways Out

At this point one can look at the available options. How can vagueness be reconciled with logic without forgoing the normative value of the latter?

A popular stance is to take the ontological turn and to abandon classical logic in favor of some kind of "fuzzy logic".[9] To the extent that the referents of our names and the extensions of our predicates can have hazy boundaries, the truth conditions of our statements can also be blurred; some statements will be definitely true, some will be definitely false, and some will have a truth value which is somewhat intermediate between true and false. For example, after removing a few thousands hairs from the count's scalp, the statement

(12) The count of Montecristo is still hairy

[9] Fuzzy logic can be traced back to Zadeh's fuzzy set theory [1965]. Two representative applications to vagueness are Machina [1976] (on vague predicates) and Tye [1990] (on vague names).

will no longer be fully true because the count will no longer be a clear member of the fuzzy set of hairy people (and not yet a clear member of the fuzzy set of bald people). And after a few hundred steps your statement

(13) I am still on Everest

will no longer be fully true because the land under your feet will no longer be a clear part of the vague mountain called 'Everest'. Typically one allows for an infinity of intermediate truth values, to do justice to the intuition that there is no limit to the degree to which a statement can be truer than others. For example, (13) will be slightly truer if uttered after n steps than after n+1 steps. It is actually customary to allow for a continuum of intermediate truth values, to do justice to the intuition that vagueness goes hand in hand with lack of discontinuity. (One can utter (13) at any point between the nth and the (n+1)th step.) And once an infinity of truth values is available, one can easily block the paradox involved in reasoning as in (1)–(3) or (7)–(9). In each case the second premise of the argument—the one expressing the intuition that vague predicates and vague names are, in Crispin Wright's phrase, tolerant to marginal change[10]—will not be fully true. And a piece of reasoning which relies on repeated applications of a rule of inference (modus ponens) to premises which are not fully true is a piece of reasoning that is not fully sound. Its soundness decreases as the number of applications increases.

This is a popular account but it is not without its own problems. In fact I think its problems are much worse than the troubles the account is supposed to handle. For one thing, there is something bizarre in the idea that the imprecision of certain boundaries is to be explained away by reference to an infinity of perfectly precise degrees of truth (membership, parthood). To what degree, exactly, is it true that the count of Montecristo is still hairy upon removing 10,000 hairs? To degree 0.8? Perhaps to degree 0.81? Or maybe 0.8123456? To what degree is a certain borderline molecule part of Mount Everest? Practically one can ignore such details, but the theory itself requires precise answers in each case. Second, there is the embarrassing presupposition that a point exists where one goes from full truth to partial truth, and from partial truth to full falsehood. At what point, during the process that we have envisaged, does the count of Montecristo cease to be a clear case of a hairy person? What is the last step after which it is no longer true (i.e., fully true) that you are on Mount Everest? Evidently the assumption of a boundary separating the clear cases from the borderline cases is just as problematic as the assumption of a boundary separating the true from the false. Finally, the very idea that the vagueness of our words lies in the vagueness of the world—that vague words refer to vague objects

[10] See Wright [1975].

and sets—is deeply puzzling. For what would the relevant objects and sets be? *What* exactly is the vague mountain corresponding to the name 'Everest'? How could we be so precise as to designate *it*?

I think these are serious difficulties that every ontological conception of vagueness is bound to face, and I am not at all optimistic about the possibility of overcoming them. Besides, the costs are very high: one is to give up classical logic as well as classical set theory and mereology. Accordingly, I will now focus on an alternative way out, which is based not on the ontological (*de re*) conception of vagueness but on the semantic (*de dicto*) conception. I will argue that such a conception is superior to the ontological conception and I will show how it can, if properly understood, provide a way out of the Fregean dilemma. To repeat, then: on this view there is no such thing as a vague mountain. Rather, there are many things where we conceive the mountain to be, each with its precise boundary, and when we say 'Everest' we are just being vague as to *which* thing we are referring to. Likewise, there is no such thing as a vague set: to say that a predicate is vague is to say (at the very least) that its extension is incompletely specified, and where there is incomplete specification of extension there is indeterminacy between various ways of picking out a precise extension.

Such a conception of vagueness is, I think, intuitive as it stands. It is a natural correlate of the idea that vagueness could in principle be removed by careful stipulations, which is exactly what Frege thought we should always do to prevent logical chaos. It could be removed, that is, if our stipulations could be made in precise terms. More importantly, however, the semantic conception of vagueness is intuitive because it combines very naturally with a supervaluationary account of the sort advocated by Kit Fine, David Lewis, and others.[11] For to the extent that vagueness can in principle be removed, it can be removed in many different ways. (There are many objects and sets that we could choose to assign to 'Everest' and 'bald'.) Thus, when evaluating a statement involving vague expressions it is natural to consider the many possible ways in which those expressions can be made precise. If the statement is true under all such "precisifications", then we may take it to be true *simpliciter*; the unmade semantic stipulations don't matter. In other words, it makes no difference what those expressions could mean had their semantic values been defined more precisely: what the statement says is true regardless (or super-true, as Fine has it). Likewise, if the statement comes out false under every precisification then we may regard it as false (or super-false) in spite of its vagueness. It is only when the statement comes out true under some precisifications and false under others that there is trouble. In such cases, the statement suffers a truth-value gap. In Lewis's words:

---

[11] See e.g. Fine [1975] and Lewis [1986].

> Whatever it is that we do to determine the "intended" interpretation of our language determines not one interpretation but a range of interpretations. . . . What we try for, in imparting information, is truth of what we say under all the intended interpretations.[12]

To illustrate, although 'Everest' is vague, it is nevertheless true that after 1 step you are still on Everest because this statement is true regardless of how we suppose the referent of 'Everest' to be made precise. Likewise, it is false that you are still on Everest at the end of your descent. On the other hand, there is no way we can settle the issue when it comes to the intermediate regions, since the land under your feet may turn out to be inside Everest or outside it depending on how we carve out a precise referent for 'Everest'. In those cases nothing will settle the issue for us, and the statement that you are still on Everest will fail to receive a definite truth value. This also allows us to explain why, for example, we can confidently assert that Mount Everest is in Asia and deny that it is in Europe, though we must suspend judgment when it comes to saying whether Everest is mostly in Tibet: the truth value of such a statement depends crucially on how much land one includes in the referent of 'Everest'. And, of course, we can by the same pattern explain why we feel doubtful when it comes to evaluating the statement that the count of Montecristo is still hairy after the removal of, say, 10,000 hairs: the truth value of such a statement depends crucially on how we imagine the extension of 'hairy' to be precisified. Still, the count is clearly hairy at the beginning of the process and bald at the end, since every precisification must agree with that. (The predicates 'hairy' and 'bald' are vague, but certain facts about them are perfectly clear and every precisification must comply with these facts.)

Why Supervaluationism is a Better Account

The supervaluationary account is attractive because it reflects a deep, preanalytical intuition concerning vagueness as it arises in ordinary language. We speak vaguely because in ordinary circumstances the vagueness of our words does not matter. We know that what we say would be true *if* we were speaking precisely, no matter how we imagine this precision to be reached, and therefore we don't care. The unmade semantic stipulations don't affect the truth of what we say.

To be sure, one could now object that just as fuzzy logic fails to overcome the presumption that there exist sharp boundaries—boundaries demarcating the borderline cases, if anything—so does supervaluationism. After all, supervaluationism says that the truth value of a vague statement is a function of the truth values of its precisi-

[12] Lewis [1993], p. 22.

fications, and this presupposes that the set of precisifications is itself precise. It presupposes, for example, that there exists a unique set of precisely demarcated areas which qualify as all and only the admissible referents of 'Everest', and a unique set of precisely demarcated sets which qualify as all and only the admissible extensions of 'hairy'. This sounds counterintuitive. However, this counterintuitiveness does not constitute a genuine threat for the account. Surely, intuitively it is impossible to draw a sharp line around the set of the admissible precisifications of a term, or of a set of expressions. If a certain piece of land α counts as an admissible referent of 'Everest', then so does any slightly larger piece of land obtained from α by adding a tiny hunk of matter along the border. (What difference can a tiny hunk of matter make?) If a certain set of people X counts as an admissible extension of 'hairy', then so does any set obtained from X by removing a single hair from the head of any member of X. (What difference can a single hair make?) But this only means that the notion of a precisification is itself vague. It means that the semantic machinery of supervaluationism suffers itself from the phenomenon of vagueness. And bad news as this may be, it comes as no surprise. We already know that the language in which the theory is formulated—the semantic metalanguage—is itself vague because we already know that 'vague' is vague. And if the semantic notion of vagueness is vague, so is the semantic relation of precisification. (This is, after all, the reason why we cannot hope to eliminate vagueness by means of *actual* stipulations.) So, by treating vagueness as a semantic phenomenon supervaluationism is bound to suffer from higher-order vagueness. But this higher-order vagueness is itself a semantic phenomenon and supervaluationism does not, therefore, succumb to the objection raised above against the ontological conception of vagueness.

One could argue that there is another problem with the semantic notion of a precisification. For isn't there a hidden presumption in the very idea that a vague term is one which *can* be precisified in many ways? Consider again Mount Everest and let Alpha be any precisely demarcated, mountain-shaped piece of land culminating in Everest's peak. After a few hundred steps during your descent it is indeterminate whether you are still on Everest, but it is not indeterminate whether you are still on Alpha, for 'Alpha' has a very precise meaning. Hence—one could argue—Mount Everest must be distinct from Alpha by Leibniz's law (specifically by the principle of the indiscernibility of identicals). That is, when we say 'Everest' we cannot be talking about Alpha, not even in principle. And by generalization it follows that we cannot be talking about *any* precisely demarcated piece of land. By a similar pattern, one could argue that no precise set would serve as a precisification of a vague predicate such as 'bald'. So if this line of reasoning were correct, then the supervaluationary account would be illegitimate and we would be forced to reconsider the basic issues.

The reasoning, however, is incorrect.[13] To put it briefly, the use of Leibniz's law is fallacious in this context. It involves a fallacy analogous to a familiar one arising in the presence of intensional operators. We know that the two statements in (14) have different truth values:

(14) a. It is contingent that 9 is greater than the number of planets.
b. It is contingent that 10 is greater than 9.

Yet this is not enough to conclude that 9 and the number of the planets have different properties (hence that they are distinct, by Leibniz's law) unless we also assume the equivalence between statements of the forms (15a) and (15b):

(15) a. It is contingent that 10 is greater than *t*
b. *t* is an *x* such that it is contingent that 10 is greater than *x*.

And, of course, this equivalence holds when '*t*' is replaced by '9' (a rigid designator) but not when it is replaced by 'the number of planets'.[14] Likewise, if 'Alpha' picks out a precisely demarcated piece of land, then the two statements in (16) may have different truth values:

(16) a. It is indeterminate whether after *n* steps you are on Everest.
b. It is indeterminate whether after *n* steps you are on Alpha.

This is not enough to conclude that Everest and Alpha have different properties (different spatial properties, in this case) unless we also assume the equivalence between statements of the forms (17a) and (17b):

(17) a. It is indeterminate whether after *n* steps you are on *t*
b. *t* is an *x* such that it is indeterminate whether after *n* steps you are on *x*,

where '*t*' can be replaced by 'Everest' or by 'Alpha'. And clearly enough we have no reason to make such an assumption. The equivalence holds when '*t*' is a precise designator such as 'Alpha', just as the equivalence between the two statements in (15) holds when '*t*' is a rigid designator. But when '*t*' is a vague designator such as 'Everest' the equivalence holds only on a *de re* conception of vagueness. On the *de dicto* conception (17b) is bound to be false even when (17a) is true.

One could also formulate this defense of supervaluationism with the help of an analogy. As Lewis has pointed out, supervaluational precisifications are a bit like

---

[13] In Collins & Varzi [2000] a different line of argument is considered to the effect that *some* vague terms cannot be precisified. However, the terms in question are rather special, so I will ignore that complication in the present context.

[14] The *locus classicus* is Smullyan [1948].

possible worlds, with super-truth playing a role analogous to necessary truth—truth in every possible world. The analogy between the operators 'it is contingent that' and 'it is indeterminate whether' is then immediate and the familiar diagnosis of the opposition in (15) extends directly to (17).[15] We cannot infer that Everest and Alpha are distinct, but only that it is indeterminate whether they are the same. And this is perfectly coherent with supervaluationism.

Indeed, from this point of view the objection under examination can be turned into a fatal objection against ontological vagueness. This is the gist of a nice little argument going back to Gareth Evans.[16] Let me reproduce it in a slightly modified version, as follows. Suppose, toward a *reduction*, that the vagueness of 'Everest' is *de re* and suppose that Beta is any object for which the identity statement

(18) Everest = Beta

is indeterminate. If indeed there are molecules such that it is neither definitely true nor definitely false that they are inside Everest, it should not be hard to find such an object. For example, our earlier Alpha could do. Now, the indeterminacy of (18) implies the truth of

(19) It is indeterminate whether Everest = Beta.

On the other hand, it is evident that in spite of all vagueness Everest is determinately identical to itself. Everything is identical to itself, whether its boundaries are sharp or not. Hence the statement

(20) It is indeterminate whether Everest = Everest.

must be false. But then Leibniz's law allows us to conclude that Everest and Beta are in fact distinct. This would be a fallacious move if the vagueness of 'Everest' were understood *de dicto*, as we have just seen. But it is not fallacious if this vagueness is understood *de re*. For in that case (19) and (20) imply, respectively:

(19') Beta is an x such that it is indeterminate whether x = Everest.
(20') Everest is an x such that it is indeterminate whether x = Everest.

And in the case of statements such as these Leibniz's law is perfectly applicable. We must therefore conclude that (18) is not indefinite but false, contrary to our initial supposition. We must also conclude, by reasoning in a similar fashion, that on the ontological conception of vagueness *every* statement of the form

---

[15] See Lewis [1988].

[16] See Evans [1978]. Similar arguments can be found in Salmon [1981], pp. 243–246, Wiggins [1986], and Pellettier [1989].

(21) Everest = x

is false, except when 'x' is replaced by 'Everest'. In other words, we must conclude that there are no circumstances under which a statement of this form will be indeterminate. The world is full of vague objects but each such object has identity conditions that are absolutely precise. And this is very strange indeed. Leibniz's law leaves supervaluationism unaffected but is has, apparently, fatal consequences on the alternative, ontological conception.

The Paradox Dissolved

These considerations offer some support in favor of the superaluationary account and against the ontological account. But let us now go back to our starting point and let us see how supervaluationism deals with the paradoxical arguments illustrated by (1)–(3) and (7)–(9). For that is a necessary test: we don't have a way out of Frege's dilemma—the incompatibility of vagueness with the normativity of logic—unless we have an account of this sort of paradoxical argument. So what are we to make of such paradoxical patters of reasoning within the supervaluationary framework?

The answer is that supervaluationally these patterns are (somewhat surprisingly) valid but unsound. They are valid because supervaluationism turns out to be perfectly compatible with cassical logic after all. This follows from the fact that supervaluational truth is defined entirely in terms of truth under a precisification, and precisifications yield classical models.[17] On the other hand, the arguments under examination are unsound because in each case the second premise is not true. In fact it is false (i.e., super-false), for it comes out false on every precisification of the relevant vague terms. No matter how 'hairy' is precisified, there is bound to be a number n such that the count of Montecristo is still hairy upon removing n hairs but not upon removing $n + 1$ hairs. So the statement

(2) For every n: if the count of Montecristo is still hairy upon removing n hairs, then he is still hairy upon removing $n + 1$ hairs.

is bound to be super-false. Likewise, on every precisification of 'Everest' there is bound to be a number n such that after n steps you are still on Everest but after $n + 1$ steps you are not. So your statement

(8) For every n: if I am still on Everest after n steps, then I am on Everest after $n + 1$ steps.

[17] See Fine [1975] for details.

is bound to be super-false. Hence any argument using such statements as premises is bound to be unsound. The conclusion may well follow but it need not be true.

Now, this blocks the paradox but—one could object—the price is unacceptable. For to deny such statements as (2) or (8) is to violate the intuition that the relevant vague terms are tolerant to marginal change. And that intuition is non-negotiable, isn't it?

The answer is that the intuition is non-negotiable but also misleading. Vague expressions such as 'hairy' and 'Everest' are indeed tolerant to marginal change, but not because they verify such statements as (2) or (8). Rather, they are tolerant to marginal change insofar as it is impossible for anybody to exhibit a specific *counterexample* to such statements as (2) or (8).[18] In other words, they can satisfy a semantic condition of the form (22a) without satisfying the corresponding condition in (22b):

(22) a. 'Every *n* is such that . . . *n* . . .' is false
b. Some *n* is such that '. . . *n* . . .' is false.

Of course, this is a distinction without a difference in the ordinary semantics for classical logic. That is why we have a natural impulse to demand a counterexample whenever a generalized statement is denied. It is also a distinction without a difference if vagueness is understood ontologically, at least insofar as such a conception is to provide a basis for the semantics of fuzzy logic. (In fuzzy logic the truth value of a universal generalization is typically defined as the greatest lower bound of the truth values of its instances, so if the generalization gets the lowest possible value—false—there must be at least one instance that gets that value as well.) However, in the presence of semantic vagueness the distinction becomes significant and the impulse to demand a counterexample for every false generalization leads to confusion and paradox. The paradox arises precisely because the impossibility to come up with an *n* that falsifies schematic conditionals of the form (2') or (8')

(2') If the count of Montecristo is still hairy upon removing *n* hairs, then he is still hairy upon removing *n* + 1 hairs.
(8') If I am still on Everest after *n* steps, then I am on Everest after *n* + 1 steps.

induces us to think that the corresponding universal generalizations, (2) and (8), are true. And that is illegitimate.

This sort of reply is often greeted with a gaze of suspicion, if not incredulity. However this explanation is the right one—and the only possible one—if we agree

[18] This line of response is detailed in McGee and McLaughlin [1995]. My own account is in Varzi [1999].

on the *de dicto* understanding of vagueness. For let us focus again on a concrete example, say 'Everest'. (The case of 'hairy' is similar.) On the one hand, if Everest is not a vague object then it must have sharp boundaries, for every object has sharp boundaries somewhere. Hence the statement that there is no cut-off number *n*, i.e., the statement in (8), must be false. On the other hand, 'Everest' is vague, which means that it is impossible to pick out a specific object and, consequently, to specify an actual cut-off number *n*. For there are only three possibilities, and none of them corresponds to a value of *n* which falsifies the conditional in (8'): (i) If *n* is relatively low, then the antecedent and the consequent of the conditional are both true on all admissible precisifications. (ii) If *n* is relatively high, then the antecedent and the consequent of the conditional are both false on all precisifications. (iii) If *n* is a borderline number of steps (and this may well be a vague issue), then there will be precisifications which verify the antecedent and the consequent as well as precisifications which falsify the antecedent and the consequent, but there will also be precisifications where the antecedent comes out true but the consequent comes out false, thereby yielding a truth-value gap in the supervaluation. Thus (8) is false. In short, supervaluationally it is false that there exists no number *n* of steps that marks the boundary of Everest, but there is no number *n* of steps such that it is true of *it* that it marks the boundary. This is the only reasonable thing to say if all objects are sharp but 'Everest' is vague. And it is precisely the answer delivered by the supervaluationary account. To reject it is to fall for a *de re* account of vagueness, hence for an ontology of vague entities.

I conclude that the supervaluationary account is indeed superior to the ontological account and provides us with a powerful way out of the Fregean dilemma. Vagueness is not incompatible with the normativity of logic, but only with the presumptions underlying the ordinary semantics for classical logic.[19]

References

Bellow, S., 1964, *Herzog*, New York, Viking Press.

Collins, J., and Varzi, A. C., 2000, 'Unsharpenable Vagueness', to appear in *Philosophical Topics*.

Dummett, M., 1995, 'Bivalence and Vagueness', *Theoria* 61: 201–216.

Evans, G., 1978, 'Can There Be Vague Objects?', *Analysis* 38: 208.

Fine, K., 1975, 'Vagueness, Truth and Logic', *Synthese* 30: 265–300.

Frege, G., 1903, *Grundgesetze der Arithmetik, begriffsschriftlich abgeleitet*, Band II, Jena, Pohle; partial Eng. trans. in *Translations from the Philosophical Writings of*

[19] The last two sections of this paper draw in part from material originally presented in Varzi [2000].

Gottlob Frege, ed. by P. T. Geach and M. Black, Oxford: Blackwell, 1952, pp. 21–41.

Keefe, R., and Smith, P., 1997, 'Introduction: Thories of Vagueness', in R. Keefe and P. Smith (eds.), *Vagueness: A Reader*, Cambridge (MA) and London: MIT Press, pp. 1–57.

Lewis, D. K., 1986, *The Plurality of Worlds*, Oxford: Blackwell.

Lewis, D. K., 1988, 'Vague Identity: Evans Misunderstood', *Analysis* 48: 128–130.

Lewis, D. K., 1993, 'Many, but Almost One', in J. Bacon, K. Campbell, and L. Reinhardt (eds.), *Ontology, Causality, and Mind*, Cambridge: Cambridge University Press, pp. 23–38.

Machina, K., 1976, 'Truth, Belief, and Vagueness', *Journal of Philosophical Logic* 5: 47–78.

McGee, V., and McLaughlin, B., 1995, 'Distinctions Without a Difference', *Southern Journal of Philosophy* 33 (Suppl.): 203–252.

Mehlberg, H., 1956, *The Reach of Science*, Toronto: Toronto University Press.

Pellettier, F. J., 1989, 'Another Argument Against Vague Objects', *Journal of Philosophy* 86: 481–492.

Sainsbury, M., and Williamson, T., 1997, 'Sorites', in B. Hale and C. Wright (eds.), *A Companion in the Philosophy of Language*, Oxford: Blackwell, pp. 458–484.

Salmon, N. U., 1982, *Reference and Essence*, Oxford, Blackwell.

Smullyan, A. F., 1948, 'Modality and Description', *Journal of Symbolic Logic* 13: 31–37.

Sorensen, R. A., 1985, 'An Argument for the Vagueness of "Vague"', *Analysis* 27: 134–137.

Sylvan, R, & Hyde, D., 1993, 'Ubiquitous Vagueness without Embarrassment', *Acta Analytica* 10, 7–29.

Tye, M., 1990, 'Vague Objects', *Mind* 99: 535–557.

Varzi, A. C., 1999, *An Essay in Universal Semantics*, Dordrecht: Kluwer.

Varzi, A. C., 2000, 'I confini del Cervino', to appear in V. Fano, M. Stanzione, and G. Tarozzi (eds.), *Logica e filosofia delle scienze. Atti del sesto convegno triennale*, Pisa: Edizioni ETS.

Wiggins, D., 1986, 'On Singling Out an Object', in P. Pettit and J. McDowell (eds.), *Subject, Thought, and Context*, New York, Oxford University Press, pp. 171–182.

Williamson, T., 1994, *Vagueness*, London: Routledge.

Wright, C., 1975, 'On the Coherence of Vague Predicates', *Synthese* 30: 325–364.

Zadeh, L., 1965, 'Fuzzy Sets', *Information and Control* 8: 338–353.