THE UNIVERSITY OF CALGARY

Vague Predicates in Natural Languages

by

Windsor George Viney

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF ARTS

DEPARTMENT OF PHILOSOPHY

CALGARY, ALBERTA

SEPTEMBER, 1996

THE UNIVERSITY OF CALGARY

FACULTY OF GRADUATE STUDIES

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled "Vague Predicates in Natural Languages" submitted by Windsor George Viney in partial fulfilment of the requirements for the degree of Master of Arts.

Supervisor, R. X. Ware, Department of Philosophy

B. F. Chellas, Department of Philosophy

H. J. Stam, Department of Psychology

25 Sept. 1996

Date

ABSTRACT

In their use of natural languages (such as English), people sometimes encounter difficulties that can be attributed to the vagueness of the predicates in such languages. In this thesis this vagueness is characterized, and the question of how many kinds of vagueness there are is discussed. Several logical approaches to the problem of vagueness are reviewed. The pragmatic solution to the problem is presented and defended.

ACKNOWLEDGEMENTS

I. Bougie and S. Figueroa

I. Hansen-Smythe

H. Hartley and the Fees Office, University of Calgary

J. J. MacIntosh

F. Marshall

G. X Stacey

B. F. Chellas

A. Demeter

Document Delivery Services, MacKimmie Library

K. Griffin and the staff of the University of Calgary *Gauntlet*

D. F. Hamilton

L. Hurst-Archer and S. Strobl

M. Kohl

D. E. McKerlie

J. J. Miles

L. M. Schnell

H. J. Stam

R. X. Ware

A. R. Bailey

W. Bricken

S. F. Thomas

In memory of my parents

Jean Viney

1920–1975

Henry Viney

1910–1980

CONTENTS

Rohatsu Sesshin at Tassajara

Scalding coffee from a freezing cup.
At the rim no telling
Which is which.

– Lou Hartman

The right attitude in philosophy is to accept aims that we can achieve only fractionally and imperfectly, and cannot be sure of achieving even to that extent. It means in particular not abandoning the pursuit of truth, even though if you want the truth rather than merely something to say, you will have a good deal less to say.

– Thomas Nagel
Introduction to *The View From Nowhere*

ONE

In our day-to-day affairs, we need to ask, direct, cajole, persuade, or order other human beings, and some animals and even machines, to perform, or refrain from performing, actions. At least on some occasions we succeed in using a natural language—English, here—to accomplish these purposes, by means of which we might accomplish our other purposes in the world.

Though we sometimes fail to accomplish our purposes through using our natural language either to bring about action/inaction in others or to direct action/inaction on our own part, it is incoherent to suppose that we could *never* succeed in doing so. For we would then not be able to continue to exist—persistent frustration of our fundamental purposes, such as obtaining food and so on, would quickly result in our demise.

Without reflection, we suppose that the categories that exist in our language, expressed by the *predicates* we use, really do apply to the world, and that it is at least possible to apply such categories correctly. But even a little reflection reveals that there are difficulties with such suppositions.

Here is a sentence many users of English will have uttered at one time or another:

It is blue.

When is that sentence used correctly? In an account of the correctness of applying the predicate "is blue," we would mention at least

- what "it" in the sentence refers to,
- the circumstances of the utterer of the sentence at the moment of uttering it, and especially
- the intention(s) of the utterer of the sentence.

Suppose I am playing marbles with my neighbour's children, and just as I am about to shoot, I am challenged by one of the other players. "Which marble are you aiming at?"

"The blue one," I say.

"*Which* one?"

"That one," I say, pointing.

"That's not blue," exclaims the child, "that's green!"

"Uh, I really think it's blue...."

"No way. Green."

"Blue."

"You're not playing by the rules—you have to say what colour the marble you're shooting for is."

"Fine—okay. It's green."

I might wonder what the cause of the dispute here is. Suppose the marble does, indeed, *look* blue to me, and there seems to be nothing about the circumstances—colour or amount of light, condition of my eyes and nervous system, awareness of possible confusion on my part about the application of the word "blue," and so on—that suggests I might be wrong about what colour the marble is. And, as far as I can tell, none of these possibly confusing factors affects the child, either. Suppose, too, that I am not trying to *deceive* the child. I must conclude that our dispute at least might be an honest one: that is, whatever colour the marble is, I think the word "blue" is the correct name for it and the child thinks the word "green" is.

In general, having two or more expressions for the same thing does not cause difficulty as long as the expressions are recognized by users of a language as being synonyms. "The Morning Star" and "the Evening Star" are both phrases that we now know refer to the same celestial object. "Dove" is the name of a particular light shade of grey, and would be recognized as such by people who are conversant with the vocabulary of the garment trades, so "dove" and "light grey" can, under certain circumstances, be words for what we take to be the same colour.

"Blue" and "green" are *not* generally taken to be synonymous, however. They are names of *different* colours, not the same colour, and we suppose objects in the world cannot simultaneously be two colours all over; hence, we are justified on the face if it in being confused when competent users of English use these two words apparently referring to the same colour.

The things to which a predicate applies can be said to be its *extension*. The predicates "blue" and "green" seem to have mutually exclusive extensions. One and only one of these predicates can, we suppose, apply correctly to any particular object. The sentence "It's blue" is correctly uttered concerning some objects, and incorrectly concerning others, as is the sentence "It's green," but both sentences cannot be correct simultaneously with regard to the same object.

Yet there still can be, and are, disputes—the story told above does not seem at all implausible, and disputes very similar to it are part of most people's experience.

Contrast the situation in which someone calls an object I think is blue "pink." I would immediately look for anomalies in that person's eyesight, or some sort of peculiarity in his or her learning of English, or her intention to make a joke, and so on. But with the use of "blue" and "green" disputes occur often enough that we might feel the need to look elsewhere for an explanation—and resolution—for at least some of them.

Let us call predicates such as "blue" and "green" *vague*. Vague predicates are predicates in a natural language of which competent users of that language cannot always determine, even under good conditions of observation and so on, whether or not they apply to some particular object or not, or about which the correct application can be legitimately disputed by such competent users.

Vagueness can be distinguished from other phenomena in language: generality, undecidability, and ambiguity (here I follow Fine 1975, 265). While vagueness can be thought of as deficiency in, an under-determination of, meaning, generality is an unspecificity, a lack

of content; undecidability is possible knowledge; ambiguity, univocal or over-determined meaning.

Kit Fine gives artificial but helpful examples to aid in distinguishing these phenomena. He invites us to suppose that the meanings of the "natural number predicates" nice_1, nice_2, nice_3 are given by these clauses (Fine 1975, 266):

(1) (a) n is nice_1 if $n > 15$
(b) n is not nice_1 if $n < 13$
(2) (a) n is nice_2 if and only if $n > 15$
(b) n is nice_2 if and only if $n > 14$
(3) n is nice_3 if and only if $n > 15$

Note that (2) (a) and (2) (b) are not meant to form a single, contradictory clause. The predicate "nice_1," then, is vague, while "nice_2" is ambiguous and "nice_3" general. And Fine gives a neat example of a sentence that is "possibly undecidable but certainly not vague or ambiguous": "There are infinitely many nice_3 twin primes."

Colour words are vague on Fine's characterization. But the vagueness of predicates seems to have come to the attention of philosophers not immediately as a feature just of predicates themselves, but through the appearance of certain predicates in arguments with a specific structure. These are what we now call pseudo-inductive arguments to distinguish them from the inductions performed legitimately—unless one is of the constructivist or intuitionist persuasion—in mathematics. Simplicius (1108.18; Lee [1936] 1967, 109) reports one of these arguments in the form Zeno is said to have employed:

> By this means he solves the conundrum which Zeno the Eleatic asked Protagoras the sophist. "Tell me, Protagoras," he said, "does a single grain

> of millet or the ten thousandth part of a grain make any sound when it falls?" And when Protagoras said it did not, "Then," asked Zeno, "does a bushel of millet make any sound when it falls or not?" Protagoras answered that it did, whereupon Zeno replied, "But surely there is some ratio between a bushel of millet and a single grain or even the ten thousandth part of a grain"; and when this was admitted, "But then surely," Zeno said, "the ratios of the corresponding sounds to each other will be the same; for as the bodies which make the sounds are to one another, so will the sounds be to one another. And if this is so, and if the bushel of millet makes a sound, then the single grain of millet and the ten thousandth part of a grain will make a sound." This was the way Zeno used to put his questions.

Aristotle (*Phys.* H. 5.250*a* 19; Lee [1936] 1967, 109) is the first to record this argument's attribution to Zeno, while at the same time dismissing it:

> Therefore Zeno's argument is not true, that there is no part of a grain of millet that does not make a sound; for there is no reason why any such part should not in any length of time fail to move the air that the whole bushel moves in falling.

But it is attributed by (the often unreliable) Diogenes Laertius (II, 108; [1925] 1972, 237) to Eubulides:

> To the school of Euclides belongs Eubulides of Miletus, the author of many dialectical arguments in an interrogatory form, namely, *The Liar*, *The Disguised*, *Electra*, *The Veiled Figure*, *The Sorites*, *The Horned One*, and *The Bald Head*. Of him it is said by one of the Comic poets:
>
> > Eubulides the Eristic, who propounded his quibbles about horns and confounded the orators with falsely pretentious arguments, is gone with all the braggadocio of a Demosthenes.

> Demosthenes was probably his pupil and thereby improved his faulty pronunciation of the letter R. Eubulides kept up a controversy with Aristotle and said much to discredit him.

Cicero (*Academica* II 49–51; [1933] 1956, 529–531) remarks at length the complaints against these arguments in still-familiar fashion:

> ... as a first point one must criticize them for employing an exceedingly captious kind of argument, of a sort that is usually by no means approved of in philosophy—the method of proceeding by minute steps of gradual addition or withdrawal. They call this class of arguments *soritae* because by adding a single grain at a time they make a heap. It is certainly an erroneous and captious kind of argument! for you go on mounting up in this way: 'If a presentation put by the deity before a man asleep [i.e., a dream] is of such a character that it is probable, why not also of such a character that it is extremely like a true one? then, why not such that it can with difficulty be distinguished from a true one? then, that it cannot even be distinguished? finally, that there is no difference between the one and the other?' If you reach this conclusion owning to my [i.e., Lucullus's] yielding to you each successive step, the fault will have been mine; but if you get there of your own accord, it will be yours. For who will have granted you either that the deity is omnipotent, or that even if he can do as described he will? and how do you make such assumptions that, if it is possible for x to resemble y, it will follow that only with difficulty can x and y be known apart? and then, that they cannot even be known apart? and finally, that they are identical? for example, if wolves are like dogs, you will end by saying that they are identical. And it is a fact that some honourable things are like dishonourable ones and some good things like not good ones and some artistic things like inartistic ones; why do we hesitate therefore to aver that there is no difference between these? Have we no eye even for incongruities? for there is nothing that cannot be carried over from its own class into another class. But if it were proved that there is no differ-

> ence between presentations of different classes, we should find presentations that belonged both to their own class and to one foreign to them; how can this possibly occur? Consequently there is only one way of routing the difficulty about unreal presentations, whether depicted by the imagination, which we admit frequently to take place, or in slumber or in under the influence of wine or of insanity: we shall declare that all presentations of this nature are devoid of perspicuity, to which we are bound to cling tooth and nail. . . .

The structure of the sorites argument and its cousins has been neatly encapsulated by Dummett (1975, 303) in what he refers to as "Wang's Paradox":

> 0 is small;
> If n is small, $n + 1$ is small:
> Therefore, every number is small[1]

A paradox, in the sense the expression is used here, is an argument with what seems to be impeccable structure and faultless procedure—but an absurd conclusion. The first premiss of the argument above seems irreproachable—if any number is small, surely that is zero. The second premiss seems not to be awry, either; adding one to a number, we suppose, cannot possibly make it *not* be small. Yet the conclusion is unacceptable: how can every number—a googol, for instance—be small?

The problem seems not to be with the argument, so it must, we suppose, be with "small." But it is hard to see what is wrong with the word. It conveys a fact about the thing of which it is predicated, some information that could even be useful on occasion: the small thing is, in particular, not large. These expressions are both relative to the context in which they are uttered, and they are both vague—in Fine's terms, deficient in meaning. They do not tell the hearer how massive the thing of which they are predicated is, or how tall, or how wide, and so

on. And they do not even tell what size the thing is with reference to some measurement system. But there is *some* meaning in them, and perhaps just the meaning needed at the time they are uttered.

So it seems there are vague predicates in natural languages—sometimes troublesomely so. It is one problem to discover the character and kinds (if more than one) of vagueness, but another, and prior, one to determine which, of the myriad predicates in a natural language, the vague ones are.

When I assert that there are five billiard balls on the table, for instance, one would at first suppose that only the most elaborate or implausible story could explain my being mistaken about such a matter: invisibility of some objects or something such would need to be invoked. So one might conclude, then, that "five," when correctly predicated of the number of billiard balls on the table, is *not* vague. But numerals, as Thomas (1995, 43–44) points out,

> are just words—descriptors of attribute value of the attribute of cardinality (in the first instance—where the so-called "whole" numbers are concerned), which may be used to characterize discrete populations. For example, in describing a flock of sheep in a field, we may—in an obvious ideogram—describe its cardinality as "{•••}" if there are "three" sheep in the flock. This notion of cardinality is an abstraction we may as easily apply to a fleet of ships or people in a crowd, or whatever. And descriptors of attribute-value for this attribute of cardinality may be designated as "{•}," "{••}," "{•••}," or "one," "two," "three," or "1," "2," "3," etc.

Thomas contrasts the exactness of numbers as descriptors of cardinality with the linear continuum. Before the last two centuries,

> there existed no convention by which every point in a continuum could in principle be described in a way that would distinguish it from every other point in the continuum.

But it is now possible to understand how the real numbers can be constructed starting from the whole numbers, and that there is a one-to-one correspondence between the points of a linear continuum and the real numbers. Thomas then suggests:

> It is by virtue of this correspondence that we could claim to be able—in principle—to distinguish every point in a continuum. But ... this ability in principle does not quite extend to an ability in practice.

Being *countable* (to be discrete and infinite (have the attribute of the cardinality of the whole numbers)) and being *measurable* (with the attribute of the cardinality of the linear continuum) appear to be two different features of collections of things in the world. When I determine the number of billiard balls on a table, I imagine I can say with certainty, for instance, as in the above example, "five." If I am asked, by contrast, to say how many stars are visible to the naked human eye at night under good viewing conditions, though I might momentarily be daunted I can at least think of a way to come up with an approximate whole number (stars are discrete things): assume that the visible stars are more-or-less evenly distributed throughout the sky; count the stars in one patch of sky, then multiply the number of stars counted by the inverse of the fraction that the patch is of the total area of the sky. (The ancient and beautiful phrase "as uncountable as the stars in the sky" turns out, on this procedure, to be much less impressive than it at first appears: a human observer with normal eyesight and under optimum conditions can see approximately six thousand stars—hardly an "uncountable" number—three thousand of them at any one time from the surface of Earth.) An important assumption has, of course, to be made in such cases: uniformity in both the sample, which is taken to be representative, and in the distribution of the things to be counted.

Consider again, for a moment, the billiard balls. What if I am not immediately next to the table, but on the other side of the floor of the large building which contains the table. Can I

still count the five balls? Or what if I am across the street from the building, looking at the table through tinted glass? Though it is in principle possible to count the balls, it may not be possible in practice. And it is this gap between principle and practice which might account for the vagueness of even such a thing as numbers:

> Vagueness in our knowledge is, I believe, merely a particular case of a general law of physics, namely, the law that what may be called the appearances of a thing at different places are less and less differentiated as we get further away from the thing. (Russell 1923, 91)

Another source of difficulty might be the nature of things one is trying to assess (count or measure). There just may be no answer to the question "How many?" with respect to such things.

Consider mountains. If one is trying to count mountains one must first determine what is to count (in another sense) *as* a mountain. In the topography of northern India, where mountains erupt from the surface of an alluvial plain, it is not very hard to count them in any particular area. But in the Rocky Mountains, by contrast, it is not necessarily a straightforward matter deciding which prominences above other prominences shall also be labelled "mountain."

In such a case it is, perhaps, possible to adopt some kind of arbitrary definition, calling a mountain any prominence above the surrounding topography that fulfills certain criteria of both relative and absolute size, but one can easily imagine such criteria being disputed.

The case of clouds is even thornier. It seems to be extremely difficult to determine what a *cloud* is, especially as, whatever they are, clouds can readily divide into more clouds and fuse into fewer. Observers of a sparsely clouded sky might agree, at a given moment, how many clouds there are, and a few moments later be unable to agree at all upon the number—"There

are now three clouds where there was one before," one observer might insist, while another claims there are now just two.

And it is not at all clear what it can mean to count "things" like thoughts or beliefs. Believing six impossible things before breakfast with the White Queen, for instance, might seem possible at first blush, but on more careful reflection beliefs seem just not to be the sort of thing that can be counted—it is not obvious what can count as the boundary between one thought or belief and another.

The notion of counting, then, seems to require that what are counted be regarded as being determinate, discreet things—that is, countable things must have discernible boundaries around them, and the boundaries must be sufficiently stable for long enough for the counting to take place. When things have been correctly counted, there is no remainder. Measuring, by contrast, seems at first to be a less narrow notion: all that things need to be measurable is a feature that can be discriminated by our senses or extensions of them. Measurement embraces the notion of remainders, amounts (not numbers) of a feature *not* taken into account by the means of measurement. In measurement there are intervals, degrees of accuracy, but in counting there is only correctness or incorrectness, unless one is speaking roughly.

I have deliberately avoided using the more philosophically loaded "object" and "property" in the above discussion since I do not know what either of them means in any deeply ontological sense. For my purposes, "thing" need only mean discriminable by some means from a background. Nor does this definition preclude a thing's being made up of, composed of, other things—a table is made up of atoms, which are themselves things, and those atoms are made up of yet other things, and so on. Thing-hood, on this broad definition, thus becomes rather less concrete than one might suppose is necessary;Q: the number of things there are in any situation under discussion might vary enormously from moment to moment, depending in part on the tenor of the discussion. This seems to me, however, the preferable alternative to there being some fixed number of things that an observer can, at least in princi-

ple, correctly determine—this project is doomed to failure, in any circumstance except perhaps mathematics, by (if nothing else) the limitations of human senses or their extensions and the limitations of our conceptual powers.

"Feature" is an equally difficult notion. Whatever features are, they seem to be able to be "removed"—if only in the imagination—from a thing without that thing's ceasing to exist. The standard example is once again colour: a blue ball does not cease to be a ball if it is painted yellow or observed under conditions that make it appear yellow. It does, if painted, cease to be a blue ball. Blue-ness seems to be a feature, in this example, of the ball. I do not know what "property" of a thing its colour is, and whether colours exist apart from things that are coloured, but for my purposes here this characterization of features will suffice, and does not, I think, run counter to even quite sophisticated notions of what a property might be.

The application, then, of words for colours, and of words indicating size, duration, shape, density, and other features of things or processes in the world, can all disputed at one time or another by competent language users. It is not just in their being *disputable* that the vagueness of such predicates consists—everything is disputable by, for instance, the determined skeptic. It is that there seems to be about such features of things no determinable fact of the matter. We cannot imagine a test that would determine once for all that a particular shade is, indisputably, blue, though a test for the wavelength of the light reflected by an object coloured that shade can be carried out to any degree of precision up to the limits of the currently available technology. Unless it turns out that even space itself is "granular," any length, a wavelength included, will never be measurable to complete precision. Features like length, and colour, appear to be continua. That notwithstanding, we do not think that measuring the wavelength to *any* precision will settle the matter: what colour an object is, is *not* just a matter of the wavelength of the light it reflects.[2]

The discussion of vagueness thus far has focused on the problems presented by the predicates in a sorites-type argument, where the logic seems right and it seems there is something

awry with the predicate itself, and of just distinguishing points on a linear continuum, as with colours like green and blue. It looks like sorites-type arguments may be reducible to the distinguishing problem, especially if Russell is correct.

But Pelletier and Berkeley (1995, 828) suggest that there is another kind of vagueness, which they call *family resemblance* vagueness. Examples with this sort of vagueness, they suggest, are "religion" and "game" (or in terms of predicates, "is a religion" and "is a game"—another example: "is a friend") About the first example, they remark:

> there are a number of factors relevant to determining whether a social practice is a religion. Having none of these properties guarantees failing to be a religion, and having all of them guarantees being one. However, there is no continuum of the sorites variety here; for example, it is easy to distinguish possessing four from possessing five of the properties, unlike the sorites case where such a change is imperceptible. . . .

I am not sure it is *easy* to distinguish having the sorts of properties (I would prefer "features") that could be said to make a social practice a religion from not having them, or that such features do not have degrees. Bearing in mind the discussion above about the relation between the designators for cardinality, it would seem that it is at least a possibility that all features can be possessed in some degree. One might suppose that there are some features the possession of which cannot be a matter of degree: being square, for instance. There is certainly no *conceptual* vagueness about something's being a square—one can know perfectly clearly what the expression means—but what is not obvious is whether or not anything *is* square. Things can be *square enough* for certain purposes—close enough for jazz, so to speak. Even notably un-square things—like a hastily sketched, four-sided figure—can be understood to be *standing for* a square, but the sketch is not square.

If no feature can ever be had by something perfectly, that is, must always be had to some degree, there might then well be only unidimensional vagueness (of the sorites or colour-continuum type) and multi-dimensional vagueness (with a dimension for each of features that make up the social practice called "religion," say). The difference would then be just in the number of dimensions having to be considered, not in the *kind* of vagueness.

However many kinds of vagueness there are, our difficulty with vague predicates arises because we cannot understand how several of our ideas about the world (including human language) can be correct together—taken together they are inconsistent.

1. Things and features either *exist* or *do not exist*—not both of these nor anything "in between" them—and exactly one of these, at least in some cases, independent of our conceptions of them.

2. Every sentence attributing (or denying) a feature to a thing is either *true* or *false*, true when the entity has (or does not have) that feature, false when the thing does not have (or has) that feature.

3. Many features of things seem to be possessed by them (or are at least evident to us) in *degrees*—these features can be either continuous (colour spectra, for instance) or discreet (numbers of sesame seeds).

1 and 2 are all right together, but 3—which is undeniable (prima facie, at any rate)—causes us to be irresolute about 1, and hence about 2.

There are several approaches that can be taken to the problem of vague predicates.

We can suppose that we can just get along with the vague predicates we have, and designate/negotiate our way out of difficulties as they arise. But recall the intuition discussed above that at least some features of the world are not correctly dealt with in this way, and in some instances, it is culpable ignorance—no matter what language games or forms of life we may be embedded in—to proceed in this fashion.

We might take the problem to be an empirical one—1 and 2 really are the case; it is 3 that causes the difficulty. Our faculties of perception (and perhaps of conception) are limited. The solution to our problem, then, will be to get better apparatus, investigate further, until one discovers (and no longer needs merely to designate) the boundaries between predicates, so to speak. But there are features of the world about which, it would seem, *no* amount of investigating, with no conceivable apparatus, can determine the boundary of predicates referring to them.

We could suppose that the problem lies with our natural language, that is, with the two features, true and false, of some utterances in a language—as opposed to being or not being the case, which is a feature of the world.

There are several approaches that can be grouped under this rubric.

- We could dispense with or modify (at least one of) two important principles of logic: the law of excluded middle (that for any statement *p*, either *p* or not-*p*) or bivalence (there are only the two *truth values*: true and false).

These two, it turns out, are not necessarily the same: that is, under some interpretations of the meanings of these expressions, excluded middle does *not* imply bivalence. Dispensing with bivalence means, in general, to allow (or is it to discover?) there to be three or more truth values. But then it will have to be sorted out what *that* means—in particular, how one is to reconcile this revised notion of truth with our unsophisticated notions about what "true" and "false" mean.

- One could be pragmatical (not necessarily practical!).

This last seems, perhaps, merely more a *psychological* than a philosophical solution; that is, it is only about what it is we happen to do when we resolve our problems about vague predicates. But that might be the best we *can* do, given our intellectual and conceptual resources.

- One could demonstrate that there can be *no* solution to the problem of vague predicates.

If achieved and believed, this would show that the pragmatical approach is just that: an approach, not a solution. We just have to muddle through somehow, so to speak, in the absence of any rational approach to the difficulties raised by vague predicates—which difficulties must just be treated as disagreements about the use of words settled, if at all, by mutual agreement as they arise. I agreed, for instance, in the first example above to call the marble "green" when I saw that not doing so would bring the game to a grinding and acrimonious halt. But consider the sentence

It's good.

Suppose I utter that sentence in answer to the question, "What do you think about capital punishment?" Or "Any thoughts about all of the death and destruction that resulted from the earthquake?" Or "How's your golf game?" There are some predicates, like "good" and "bad," that we might suppose it is important—imperative, even—to apply correctly. But if such predicates are vague—and many might well be—then we have the problem that we are not, and maybe even *cannot* be, sure when we are applying them correctly.

To motivate the enquiry even more strongly, consider this artificial example. Suppose you have normal eyesight. In particular, you are not colour blind. You somehow have become the prisoner of an always truthful and resolute fiend who confines you, without clothing or

any other accoutrement, in a locked, bare room and says to you, from behind an obviously unbreakable pane of glass in one wall of the room:

> Look at that other wall. The patch to the left of that ordinary, unlocked door is, as you can clearly see, green. The one to the right of the door is, equally clearly, blue. The door is, I'm sure you can tell, some colour intermediate between the colours of the patches. I assure you the door is either green or blue, not aqua or ultramarine or grue or bleen or anything else. Behind the door is a table on which there is a single syringe containing a remarkable drug. If you open the door, I will take out the syringe and inject its contents into you.
>
> If the door is *green*, the drug in the syringe will cause your life to be long, your physical health robust, your mood cheerful—in particular, you will be completely free of thoughts about difficult and vexatious philosophical problems—and, at your life's end, the drug's final effect will be to make your death serene, unregretful, quick, and painless.
>
> If the door is *blue*, the drug in the syringe will also assure you of a long life, but that life will be miserable: you'll have to suffer constant, intense physical pain and mental fear, worry and torment—by, among other things, philosophical problems you believe, perhaps rightly, you cannot solve—yet you will be unable even to conceive of ending your own life or arranging for someone else to do so to curtail your torment. You will never become insensitive to the pain and anguish; they will increase steadily up until your death.
>
> If you do *not* open the door at the end of the hour, I will give you an injection from the syringe I am holding in my hand—its contents will put you to sleep in just a few seconds, and then kill you.
>
> Ask me any questions you like, including about the colours of things other than the patches and the door. I *won't* answer "What colour is the door?" I'll answer truthfully—including, of course, "I don't know" if *that*'s true.
>
> You cannot negotiate or bargain your way out of the choice I have given you. And there is no escape from this room. Your hour starts now.

> Oh, by the way, here's a snippet of information that might interest you: the wavelength of the light reflected by the pigment on the door is, as nearly as the best available measurements have determined, mid-way between the wavelengths of the light reflected by the pigments on the two patches.

Your only choice, in this fanciful but nightmarish situation, is whether or not to open the door—but you probably want that choice to be informed by knowing which drug is in the syringe behind the door. That, in turn, can be known by what colour the door is. But what colour *is* it? You know that it is mid-way between the colours of the patches, but the crucial question is: what does the *fiend* call the colour of the door? And can you find out in an hour by getting (truthful, honest) answers to any questions you are permitted to ask?

There is the way we happen to resolve the problem, when it arises, of vague predicates, but there might also be a (different) way we *ought* to resolve the problem. (In the scenario above, which would you be concerned to do? Might you, for instance, be able to persuade the fiend to agree with whatever colour you say the door is? Ought you to try to do that?) Philosophically, one searches not, at least not necessarily, for a theory of how conversation happens to be conducted, but how it *ought* to be conducted—that is, what the truth or truths of the matter are, if there are any and they can be known.

TWO

We might look hopefully to the state of the formal logic of a subject of interest in order to assess our degree of understanding of that subject. In their logical investigations, philosophers, linguists and mathematicians try, among other things, to isolate and study particular features of, say, natural languages, and by so doing endeavour to make explicit and clear what is otherwise hidden or confused in our knowledge of the workings of such languages.

But we should not be over-confident that logic will be the last, or the best, word on a subject. As the late Hao Wang reminds us ([1955] 1978, 6), to be rigorous is not necessarily to be reliable. This is especially true since it was discovered in mathematics and then in logic that a formal system, once abstracted from the subject or situation in the world (if any at all) that gives rise to it, can leave that subject/situation behind and explore what might be called purely formal relations.[1] And it is salutary to remember that "... exact formalizations [of ordinary concepts] almost always distort our ordinary language at one place or another" (Wang [1955] 1978, 4). Since it is with vague predicates in ordinary, natural languages we are here concerned, this caveat should always be borne in mind as we canvass the efforts at the understanding of the logic of vagueness.

The first thing to be clear about when discussing the relation of the logic of a language to the use of that language is that users of a language may be able to do so competently without knowing in detail *how* they are able to do so—indeed, this is the situation for every user of a natural language, since there is universal agreement that no one has yet succeeded in giving an exhaustive logical account of how a natural language works. But this does not, of course, preclude the possibility that a logic might capture *some* of the features of the working of a language. One could strive to create or discover a logic that is not only useful in or *for* a language, but by exploring the logical features of a language one can hope to discern more clearly how a language *ought* to be used.

Some utterances of natural languages, the statements or propositions, have the feature that they are either true or false. A statement is said to be *true* when it reports states of affairs or processes—I purposely leave these notions undefined—in some specified domain, and the states of affairs or processes it reports do obtain, or are the case. A statement or proposition is *false* when the states of affairs or processes which the statement reports do not obtain, are not the case.

How statements get to be true or false is a matter of seemingly endless discussion and exploration. The theories about how this comes about, in what truth and falsity consist, are many. Whatever sort of theory one adopts of truth and falsity—a correspondence theory, a coherence theory, or some more sophisticated variant or combination of these—on the traditional view at least these things that can be asserted *about* the statements of a language:

1) they must have only one of the value "true" or "false" at any one instant—*tertium non datur*, there is no truth values other than these,

2) for every statement, either it or the negation of that statement has the value "true" (and similarly for "false," of course),

3) a statement's value of "true" or "false" remains constant as long as the states of affairs or processes which it report remains unaltered.

The first is called the principle of bivalence, and the second is the law of excluded middle. For convenience, call the third assertion the principle of persistence, which is a reminder of at least what we hope is correct of the statements of our languages: that whatever relation it is that they bear to the states of affairs or processes which they are statements about that make them true or false, they do so in relation to those states of affairs or processes, varying as they do.

Logic is concerned with capturing certain features of collections of statements of languages called *arguments*, in particular the features of *consequence* and the *preservation of truth*.

Whenever the premisses of a valid argument are true, so should its conclusion be. This feature of valid arguments is often expressed by saying that the conditional sentence whose antecedent is the conjunction of the premisses of a particular valid argument and whose consequent is the conclusion of that same argument is a logical truth. Since in a valid argument its conclusion is true whenever all its premisses are, the argument preserves truth, in the sense that using such arguments one should never be able to pass from the collection of statements—the premisses of the argument—each of which is true, to a statement—the conclusion of the argument—which is false.

Suppose there is a binary relation R which is transitive, and let a, b, c... each name a different thing in a domain of discourse. Then this is, by that very transitivity, a valid argument:

$$\begin{array}{l} aRb \\ bRc \\ cRd \\ \quad . \\ \quad . \\ \quad . \\ xRy \\ yRz \\ \hline aRz \end{array}$$

Clearly, as long as the relation R is indeed transitive, it does not matter how many premisses of the indicated sort there are: the relation R will hold between the things named by the first and last names in the series. If all of the premisses of this argument are true, its conclusion will be, too.

Now consider another argument, whose premisses have the same structure but contain a different binary relation, S:

aSb
bSc
cSd
.
.
.
xSy
ySz

---

not-*aSz*

If this is a valid argument, then however the relation S is characterized, it is nontransitive, since the truth of the premisses does not "carry over" or "pass through" on the pattern of the first argument above, to the conclusion.

A relation can always be characterized extensionally just by listing the ordered *n*-tuples (here, duples) consisting of the things in the domain which are related by it. In formal terms, that is just what a relation is, but the relations for which there are words in natural languages typically have a feature that makes them interesting and importantly different from just any relation that can be defined in this fashion: relations expressed in natural languages are often intended to hold everywhere in the universe of discourse—in this case, the actual universe. That is, one can use *variables* rather than *names* in the definition of the relation.

There are binary relations indicated by English expressions that are transitive, of course: "is more massive than," "is implied by," and the like. But there are also binary relations expressible in English that are nontransitive, like S in the second argument above: "is a sufficient condition for" might be one such. That is, *a* might be a sufficient condition for *b*, and *b* sufficient for *c*, yet we ought not automatically to assume that *a* is a sufficient condition for *c*.

A relation that is, notoriously, nontransitive is expressed by the predicate "is similar to." Other relations, referring back to the examples presented in Chapter One, include "is as blue as," "is the same colour as," "can be labelled 'heap' as legitimately as."

So here, pursuing the discussion of logic and its ability to capture features of collections of statements of natural languages, we again encounter the problem of vagueness, this time in another guise.

A relation's being transitive is very desirable for at least this purpose: it gives us the assurance that we can extend our knowledge outward with it, so to speak. When we discover that something new in our experience of our world bears a certain relation to something else, and we confidently believe that relation is transitive, we then can link that new thing, as it were, to all the other things that bear that relation to each other. From the assumption—borne out by our experience and concepts of the regularity of patterns in the world (one important instance of which is cause and effect) persisting long enough for them to be reliably discerned and then sometimes even predicted—that such relations exist in the world and are expressible in our natural language, we can extend our knowledge of the world.

If we discover that a relation expressed by a predicate is *non* transitive, we lose our assurance that we can extend our knowledge of the world in this way.

Let us suppose for a moment that our problem with the vagueness of at least some of our predicates arises not from the world, which carries on doing whatever it is doing without being detained by our troubles in talking or reasoning about it, but with the relation our ideas, our concepts, and those predicates bear to the world.

It is remarkable, but perhaps indicative of the difficulty of the problem, that though vagueness in language was recognized as a problem from antiquity (see the passages quoted in Chapter One), it seems to have been only in the first part of this century that researchers began seriously to explore the logic of vagueness. With the independent development by Jan

Łukasiewicz and Emil Post of truth tables and semantic methods around 1920 (see, for instance, Post 1921) the first step toward trying to reform logic to take vagueness into account was made. In a sense, Łukasiewicz made a move similar to what Bolyai and Lobachevsky and Riemann did in geometry: he questioned an "intuitively obvious" assumption, namely, that there are only two "truth values." While investigating propositions about the future, Łukasiewicz proposed that they might have a truth value such as "possible" which was neither true nor false (Łukasiewicz 1920). In geometry, the parallel postulate was questioned; in logic, the principle of bivalence.

Łukasiewicz generalized his results to allow any finite number of values (Łukasiewicz 1922/23, 1930, Łukasiewicz and Tarski 1930). His results in three-valued logic were axiomatized by Wajsberg (1931); other logicians, like Post with his *m*-valued calculus (1921) and Kleene (1938), explored and further extended work in multi-valued logics.

Fuzzy set theory—initiated by Zadeh in his oft-cited 1965 paper—and the fuzzy logic developed from it suppose that predicates may be true to *some degree* or other, rather than just true or false only. Usually the designation of these degrees of truth is made with reference to the closed interval on the real line from 0 to 1, with 0 being the value "completely, utterly untrue/false" and 1 being "completely, utterly true." Since the goal is to model some heretofore unmodeled features of natural languages, it is supposed that the new calculus should retain, as much as possible, the features of the received (classical) structure while capturing features of natural-language reasoning that are not modeled in the older calculus.

In particular, natural languages have certain logical operators (connectives) which are truth functional—that is, knowing the truth values of statements joined by such connectives immediately yields knowledge of the truth value of the compound statement. The almost trivial case is negation: if a statement is true, its negation is false, and vice versa. A conjunction is true just when both its conjuncts are true, and false otherwise. And so on for disjunction, material implication, and biconditionality. (Though a truth-functionally complete

propositional calculus can be made with fewer than these connectives, this set reflects its members' origins in the logical connectives of natural languages.)

But how do we reason—correctly, we hope—with vague statements? What is the truth value of "He is tall," to use one of Zadeh's examples? "Tall," like "blue," is a vague predicate. When applied to male human beings of Caucasian descent, for instance, a person who is six feet in height is probably reckoned by most competent speakers of English to be tall. And a person who is five-foot-six is *not tall*, or *not very tall*. (Here the situation becomes more complex as qualifying words like "not very," called *hedges*, are introduced.) But what about a man who is, say, five-foot-nine? Is he "sort of tall"? In the case of this predicate, it is clear that for persons of the last-mentioned height, it does not do to give oneself only the options of "He is tall" or "He isn't tall." And "not tall" and "short" may not mean the same thing, since these two predicates might not cover the whole of the range of human heights—there are, we might suppose, heights between "short" and "not tall."

What, then, is the truth value, when referring to John and Robert, of a statement like "John is tall and Robert is not very tall" or "If John is tall, then Robert is not very tall"? Zadeh sought to find a way to calculate, given his continuum of truth values, the value of such grammatically complex statements.

There are at least initially intuitively plausible answers to these questions. The truth value of a negation, for instance, seems to be just unity minus the truth value of the negate, and vice versa. The truth value of a conjunction seems to be the truth value of the conjunct with the "lowest" value. This seems to make sense—a conjunction cannot be warranted with any greater strength than the strength of the "less true" conjunct. The truth value of a disjunction, then, is that of the "more" true of the two disjuncts. So, for instance, if the truth value of a statement A is 0.7, then the truth value of the negation of that statement will be 1 – 0.7, or 0.3. If statement B has the truth value 0.4, then the conjunction of A and B has also has that truth value, the

lower of the two statements' values. And the disjunction of A and B has the truth value 0.7, the higher of the two values.

But peculiar results emerge rather quickly. On this account (as Thomas 1995, 71 points out), the statement "He is tall and not tall" makes no sense, or at least very little—perhaps a story *could* be told that would explain or justify such an utterance. It should, we usually suppose, simply be reckoned a contradiction. Yet on the account just sketched, if either statement in this conjunction has a non-zero truth value, so will the conjunction—but surely, we suppose, a contradiction is true to no degree whatever; its truth value is zero. On the other hand, a statement containing the predicate "tall or not tall" can on this account easily have less than unity as a truth value, yet disjunctions consisting of a disjunct and its own negation/negate seem intuitively to be tautologies—always, completely true, with a truth value of 1.

There are good reasons that what we now call "classical" logic, explored and consolidated over many centuries, has taken the form it has: it turns out that any departure from it results in a profusion of complications. An obvious instance is the material conditional—every beginning logic student is initially puzzled by the material conditional, because it captures only part of what a conditional statement in a natural language conveys; the student must tease apart his or her intuitions about the meaning of conditional statements and understand which parts of them are captured by the material conditional. If one tries to capture more of that intuitive meaning formally, things become very complicated very fast—compare C. I. Lewis's "strict implication" (first developed in Lewis 1918, corrections in Lewis and Langford 1932) or the various forms of "relevant" and "relevance" logics (see, for instance, Anderson and Belnap 1975 and Anderson, Belnap and Dunn 1992)—especially if one wants to hang on to long-standing desiderata like consistency (but see below for a departure from even this logical canon).

Classical logic is very tidy—and, not incidentally I think, beautiful—but it does not work, at least unmodified, for vagueness; it is the very thing that seems to *cause* the problem in the

case of the sorites paradox. But the first major attempt to modify or augment classical logic to address vagueness, fuzzy set theory and logic, has its own problems. Perhaps they are remediable, though some critics think they are not, while others are not sure that fuzzy set theory and logic are not formally equivalent to the already highly developed probability calculus.[2]

A different modification of classical logic is suggested by Fine (1975) as one which can handle vague predicates. On Fine's analysis, referred to as the super-truth or supervaluation account, a vague sentence is "true if and only if it is true for all ways of making it completely precise" (Fine 1975, 265). The "making precise" is called *precisification*.

Consider the sentence "Friar Tuck is bald" (I am following the simplified account of super-truth in Pelletier and Berkeley 1995, 827–828). Choose a number—a precise number—of hairs to demarcate the bald/non-bald border. On this valuation, then, Tuck is either bald or not bald, and the sentence just mentioned is true or false. Now change the valuation: create an different bald/not-bald borderline, and so on. If the sentence above is true on *all* such valuations, then it is super-true, "really" true; if it is false in those valuations, it is super-false, "really" false. Sentences which do not fit these criteria are vague.

Recall that in fuzzy set theory the conjunction of a sentence and that sentence's negation could have a truth-value greater than zero—that is, it could be true to some (non-zero) degree. But the conjunction of a sentence and that sentence's negation on the super-truth analysis is false in every valuation, so it is "really" false, super-false. And the disjunction of any sentence and that sentence's negation is super-true. On other words, our intuitions about the contradictory and tautological character of certain sentences are preserved in the supervaluation account even when the predicates in those sentences are vague, but not in the (unrenovated) fuzzy set theoretical account.

Fine maintains that "[t]here is no point in withholding truth from a sentence that can be made true by improving any improvement in precision" (Fine 1975, 279). This also seems to conform to our intuitions about the investigation of sentences containing vague predicates

should go. But it is not intuitions, but *reasons*, that Fine claims guide his choice of the account he presents:

> These arguments [for preferring the super-truth view] are essentially claims of the following form: such and such theory is the only one to satisfy the reasonable conditions *X*, *Y* and *Z*. Such claims are of great importance, for they provide a point or rationale for the theory in question: if you want the conditions then you must accept the theory. All too often, truth-conditions for different languages have been constructed with insufficient regard for rationale. Their basis has often been a scanty set of intuitions. Thus a great advantage of the present approach is its possession of a uniquely determining rationale (Fine 1975, 279–280).

Fine appears to be keenly aware that "intuition, as we all know, is not a safe guide: it cannot be used safely as a criterion of either truth or fruitfulness in scientific explorations" (Nagel and Newman [1956] 1978, 16). As he himself points out twice in his paper, the working out of some of his results is very complicated (though not, of course, thereby necessarily either untrue or not useful), and it may not be obvious, or even clear after much work, that the results are correct.

In the end, Fine claims "to assert a vague sentence is to assert, generally, its precisifications" (Fine 1975, 282) (he contrasts this with the asserting of an ambiguous sentence, which is, he claims, "to assert, severally, each of its disambiguations" (ibid.)). He provides this image to help make his point about the relation of vague and ambiguous sentences:

> Ambiguity is like the super-imposition of several pictures, vagueness like an unfinished picture, with marginal notes for completion. One can say that a super-imposed picture is realistic if each of its disentanglements are; and one can say that an unfinished picture is realistic if each of its

> completions are. But even if disentanglements and completions match one for one, how we *see* the picture will be quite different (Fine 1975, 282–283).

There is yet another development of logic that might be a candidate for a correct logic of vagueness: paraconsistent logic (I follow the account in Priest and Routley 1989b, with changes to their notation as necessary).

A *theory* T—i.e., a set of statements closed under consequence—is *inconsistent* if, and only if, for some statement A, A and the negation of A are in T. Call T *trivial* if, and only if, every statement is an element of T. In classical logic, all theories that are inconsistent are trivial, since *every* statement is a consequence of an inconsistent set. But advocates of paraconsistent logic claim that there are inconsistent but *non* trivial theories.

One might be able, then, to develop theories which, though inconsistent, are not explosively so, and thus may have their inconsistency held in check, as it were, and thus still be very useful. "The important fact about paraconsistent logics is that they provide the basis for inconsistent but non-trivial theories. In other words, there are sets of statements closed under logical consequence which are inconsistent but non-trivial" (Priest and Routley 1989b, 151). There are in fact already a number of significant instances of inconsistent-but-not-trivial theories:

> A first example ... is naive set theory, the theory of sets based on the full abstraction axiom scheme For some y and for all x, x is an element of y if, and only if, A. This, together with extensionality, characterizes the intuitive conception of set. The theory is inconsistent since it generates the set theoretic paradoxes (e.g. where R is the Russell set, defined as x such that the negation of x is an element of x, standard paradox arguments show that R is an element of R and the negation of R is an element of R). Yet this is non-trivial because there are many claims about sets which the intuitive notion rightly rejects (e.g. that the set consisting of L is an element of L, where L is the null set) (Priest and Routley 1989b, 152).

Inconsistency is anathema to most logicians, to be shunned at all costs. But those costs are high—for instance, the complexities of the various axiomatizations of set theory necessary to avoid the paradoxes.

Intuitively, we might want to say that a vague predicate is just as true as it is false of a thing which is a *borderline* case of the predicate. If that thing does satisfy both such a predicate and its negation, then the situation is, on the above characterization, paraconsistent. Priest and Routley (1989a, 389) think there may actually be advantages to such situations:

> [C]onsider a colour transition from red to blue through magenta. At the borderline area between red and blue, it seems much more plausible to suppose the colour to be both red and blue, than neither red nor blue. An argument that the paraconsistent approach is better in general, is that truth-value-gap approaches characteristically produce a failure of the law of excluded middle at the borderline area. Yet as Dummett and others have pointed out, this is not so plausible. In a borderline case between orange and red, we would be inclined to say that the colour is either orange or red, and it follows from this that it is either orange or not orange.

Truth-value *gap* approaches are those in which some statements are thought to be neither true nor false; truth-value *glut* approaches, those in which some statements are thought be be both true and false. The paraconsistent approach, then, one of the latter, can be used to preserve the law of excluded middle, though of course at the price of the thoroughgoing consistency of theories that would otherwise appear so desirable.

Having sketched three approaches to the logic of vagueness, the question at this point is: can any of these solve the problem of vagueness in natural languages?

It is important here to distinguish describing a problem from solving one. One of these logics, or one yet to be developed, might give a good account of how vague predicates operate

in natural languages, but the logic itself might be fraught with difficulties, and thus not in an obvious sense *solve* the problem.

Suppose a correct logic of vagueness turned out to be similar to the computer-assisted proof of the Four Colour Theorem produced by Kenneth Appel and Wolfgang Haken in 1976, in the sense that it is so complicated (and long!)that although held to be correct, no human being has actually understood it in all of its details.

The proof (or "proof") of the Four Colour Theorem was accomplished with the aid of a computer's exhaustive analysis of what were determined, by human beings, to be all the relevant cases. It was held, reasonably, that if it could be proved that in every one of these cases a map in the plane could be coloured with no more than four colours, then the theorem itself would be proved. A proof by cases is certainly in general held to be legitimate. But no human being has checked, or probably even can check, that the proofs of each of the cases is correct. Has, then, the Four Colour Theorem been proven? (Though the Four Colour Theorem has been widely held to be true since the problem was raised ca. 1852—no one has ever made a map on the euclidean plane that could *not* be coloured with four or fewer colours—no one had been able rigorously to *demonstrate* that four colours would be enough for any map, either.)

A logic of vagueness like that, then, might be correct, but for any but the most abstruse purposes, useless.

Each of the logics canvassed above tackles the problem of vagueness while confounding, in one way or another, our intuitions about how language has to work if it is going to tell us—within what limits we do not know—what is true about the world. Such intuitions are resilient; this is explainable both by their general utility and by our reluctance to abandon conceptions and practices that work well enough, whatever that criterion may mean in given circumstances.

This is not, of course, to say that our intuitions cannot be improved: compare the beliefs many people have about physics which cause them to be wrong when predicting the direction of movement of some objects (a weight swung around the head and suddenly released, for instance) but quite right about others (enabling people who have never studied physics—or, indeed, people who lived before there was any formal physics—to juggle, for example). Even a simple explanation of inertia will remedy the weight-releasing problem; it need not be detailed or even expressed mathematically.

A theory is, among other things, our telling ourselves how something goes, an account of how to predict how things will be, or at least understand how they are and came to be that way. As Einstein wittily remarked, we want a theory of something to be as simple as possible—and no simpler. A good theory about some phenomenon (vagueness in natural languages, say, or the patterns of events at the sub-atomic level) might have to be very complex: the patterns of the phenemona it is meant to explain might themselves be exceedingly complex. Such a complex theory might still be very useful, in some circumstances. But if a theory of vagueness were going to be useful in resolving the diffulties that arise in everyday discourse due to the vagueness of predicates in natural languages, it would have to be employable on the fly, so to speak. Classical logic, though complicated enough, is still useful for sorting out at least some problems that arise in our day-to-day talking and reasoning about the world. It does not yet appear that developments of logic which attempt to account for vagueness will be usable in this way.

THREE

> All I have done, really, is to replace one family of metaphors and images with another .... It's just a war of metaphors, you say—but metaphors are not "just" metaphors; metaphors are the tools of thought. ... [I]t is important to equip yourself with the best set of tools available.
>
> – Daniel C. Dennett, *Consciousness Explained* (p. 455)

Vagueness is sometimes a problem for users of natural languages, to the extent that it prevents such a language's users from realizing their purposes through the use of that language (compare the problem faced by the hapless person at the end of Chapter One).

Suppose that any natural language must represent the world somehow.[1] Suppose further that the world has *depth*: that is, it is highly detailed and rich in its structure and the relationships between its parts.

The first supposition I will not argue for.

Evidence for the second is abundant—the paradigm example is physics, which now regularly has to modify, sometimes significantly, previous theories in light of results of, for example, experiments at higher and higher energies, enabling researchers to investigate smaller and smaller entities, or entities that are otherwise unobservable but can be made to "show themselves," as it were, for vanishingly small intervals at high energies. It is not obvious that there will be an end to this investigation, at least a conceptual end (all that there is to investigate will have been investigated). It seems, rather, that the limit to investigation will be a technological (and financial) one.

Suppose now that Russell's (1923, 89) account of accuracy of representation is correct:

> One system of terms related in various ways is an accurate representation of another system of terms related in various other ways if there is a one–

> one relation of the relations of the one to the relations of the other, such that, when two or more terms in the one system have a relation belonging to that system, the corresponding terms of the other system have the corresponding relation belonging to the other system.

Any formal system contrived by human beings for human beings will be limited in how complex or rich it is. This limit would seem to be below (far below?) the degree of richness of the world—at least the evidence up to now suggests this. Such a formal system must, then, ignore some features of the world; that is, its creator must *abstract* from the myriad of features possessed by the world and represent just certain important features (important for the user of the representation) the relations between which that the system will (it is hoped) exhibit clearly.

A *purpose* might be characterized as that which causes an entity to act/refrain from acting to cause/allow the world or some part of it to become/remain some particular way. Purposes in human beings are associated with feelings—the *wish* that something be or not be the case; they are what the feelings move us to do or refrain from doing. Systems for helping us understand and act in the world—either innate and inexplicitly understood or deliberately devised—are shaped by what we need to be able to accomplish in our world in order to fulfill our purposes.

No such system will capture—would be of use if it did capture—all the features of the world.

How can the user of such a system be sure it captures the *important* features of the world? By surviving, continuing to be able to use the system. (The user of the formal system can afford to get some things quite wrong through that system's use—they "do not matter"—and others not wrong at all, which is not thereby getting them completely *right*.)

Frege claims in his *Grundgesetze der Arithmetik* (vol. 2, § 56; Geach and Black 1980, 139):

> A definition of a concept (of a possible predicate) must be complete; it 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; though for us men, with our defective knowledge, the question may not always be decidable. We may express this metaphorically as follows: the concept must have a sharp boundary. If we represent concepts in extension by areas on a plane, this is admittedly a picture that may be used only with caution, but here it can do us good service. To a concept without sharp boundary there would correspond an area that had not a sharp boundary-line all round, but in places just vaguely faded away into the background. This would not really be an area at all; and likewise a concept that is not sharply defined is wrongly termed a concept. Such quasi-conceptual constructions cannot be recognized as concepts by logic; it is impossible to lay down precise laws for them. The law of excluded middle is really just another form of the requirement that the concept should have a sharp boundary. . . .

But it is just in their *being* vague that vague predicates have their use. A natural language in which the extension of every predicate was completely known (or at least definable) would be a language ill-suited to the world we live in, a world which is and could remain, at least in some important respects, mysterious. Such a language would be unusable: for some situations it would require huge vocabularies and protracted qualifying phrases. If I am asked by my physician how long ago I had a specific surgical procedure performed on me and I reply "about six years ago," that information—though vague—will suffice for most of the physician's purposes. The vaguenesses of natural languages enable their speakers to talk about what they do not, in some sense, know—though the operation must have taken place on a particular day, at a particular hour, I may well not recall in which month, or even in which season of the year, I had the operation. This is a useful thing to be able to do, since there are many

important things with which we must contend that we do not understand at all well, but which can be captured sufficiently well by an appropriate vague concept or predicate.

Contrast the metaphors used by Russell and Fine to describe vagueness. Fine's invites what it seems could be, in effect, never-ending precisification; Russell's, a recognition of where one's knowledge is, as it were, in the world. Russell (1923, 91) remarks:

> It would be a great mistake to suppose that vague knowledge must be false. On the contrary, a vague belief has a much better chance of being true than a precise one, because there are more possible facts that would verify it. If I believe that so-and-so is tall, I am more likely to be right than if I believe that his height is between 6 ft 2 in. and 6 ft 3 in.

Our beliefs about the world do not need to be true (super-true, absolutely true); they need only to be good enough—and thus useful enough—for our purposes.

We might wish to seek a solution in logic for the problem of vague predicates for various reasons. We might want to devise machines with which we can communicate easily, but in order to do this we would have to understand and construct algorithms—if computers are not in the offing that can do this extensively themselves—that could enable the machine to "understand" vague predicates. An expert system designed to take medical histories (and thus save medical personnel the time it otherwise would take them to do so) might ask me the same question about my last surgery and have to be able to make the proper use of the information provided by my "about six years ago." This it could only do if it could understand the significance of my vague response.

The history of human enquiry is strewn with claims that such-and-such was impossible, only for that such-and-such to be later accomplished. "Proofs" must always exists within the frame of what is known, and it is precisely what is *not* known that is to be discovered (or not). Whether or not human enquiry converges relentlessly (if painfully slowly) on the truth—

"Truth is in the future," as the pragmatist slogan has it—is itself an issue that can only be thought about within the framework of current conceptual apparatus.

The inadequacy of the logical apparatus available up to the first part of this century to capture the use of vague predicates is manifest: the principle of bivalence (PB) and the law of excluded middle (LEM) were heretofore taken both mutually to imply each other and to be obvious. It was, when it comes to considering how natural languages are actually used, a case of not seeing a rather large forest due to staring fixedly at some very smooth, and not very leafy, trees. As long as PB and LEM were "obvious" in the way, for instance, the parallel postulate had been in Euclidean geometry—that is, intuitively plausible but unproved within the system—logic could make no headway on the problem of vagueness.

In order to realize some of our purposes, we have to measure or count things in the world. But measurements (and even, perhaps, countings—see Chapter One) are vague. The world eludes the confines of any measurement.

The infinitesimal calculus does its work extraordinarily well—for instance, it gives the "right" answers for discovering the volume of casks. But this "rightness" is not about some absolute correspondence to reality, but about the relation between the reality outside human agents and the intention, the purpose, of those agents making the measurement—deciding how much to pour into a cask, in the mentioned example, or how much the cask's contents can be sold for. This reckoning can be arbitrarily close, but always one side or other of the "true" value. But as long as the cask accommodates the amount poured in, and the buyer does not dispute the number indicating the amount contained, the purpose of the measurement has been fulfilled.

It seems that the account of vagueness in natural languages that is likeliest to be correct (for now) will not be in logic, but a pragmatic account.

A pragmatic account of vagueness notices, as logical accounts generally have not, two features of natural languages: their efficiency and their indexicality.

Expressions in natural languages have conventional meanings assigned to them; these conventions are of course understood—to the extent they are understood—in the course of learning the use of the expressions. But as Suchman (1987,58) points out,

> [t]he significance of a linguistic expression on some actual occasion ... lies in its relationship to circumstances that are presupposed or indicated by, but not actually captured in, the expression itself. Language takes its significance from the embedding world, in other words, even while it transforms the world into something that can be thought of and talked about.

There is an indefinite range of possibly relevant features in the situation in which an utterance is made. And we do not conceive of a list of the situation's relevant circumstances—each utterance has, as it were,

> an implicit et cetera clause. One consequence of this practice is that we always "mean more than we can say in just so many words"... (Suchman 1987,60).

The background of any assertion does not form part of the semantic content of a sentence as a sentence, but perhaps to the consternation of the logician, truth conditions for any assertion are always relative to that background (Suchman 1987, 60–61). Expressions are efficient in the sense that though they have conventionally assigned meanings, their significance also depends crucially on the circumstances of their utterance.

This reliance on the situation for significance is an expression's indexicality. That is,

> ... while one can state procedures for finding the expression's significance, or rules for its use, the expression's meaning can be specified only as the use of those procedures in some actual circumstances ... (Suchman 1987, 58–59).

Suchman borrows the example of the expression "That's a nice one," contrasting the difference in the expression's significance when uttered by a visitor referring to a photograph in her host's photo album, and when one person says it to another while they both stand in front of the lettuce bin at a grocery store. One can readily understand that the significance of "that" in the expression is indexical (being a paradigm example of linguistic indexicality, a relative pronoun). What is interesting, at first even surprising, is that the significance of "nice" is indexical, too. As Suchman notes about the photograph example,

> ... visitor and host will never establish in just so many words precisely what it is that the visitor intends and the host understands. Their interpretations of the term ["nice"] will remain partially unarticulated, located in their unique relationship to the photograph and the context of the remark. Yet the shared understanding that they do achieve will be perfectly adequate for purposes of their interaction. It is in this sense—that is, that expression and interpretation involve an active pointing to and searching the situation of talk—that language is a form of situated action (Suchman 1987, 59).

"Nice" *means* the same thing in both imagined situations, but its *significance* is very different—the prior intentions of the speaker and hearer of the remark in the circumstances of its utterance affect how it is understood, and the course of the communication after it is uttered.

"Nice" is also a vague predicate—there are photographs and cabbages to which we can say it does apply, those to which it does not, and those to which we would not be sure whether or not it would correctly apply. But in the context of the two utterances sketched above, this concern does not arise, except perhaps for a persistent philosopher. The word evokes in each of the speaker and hearer a number of indefinite, more-or-less clear, associations, and that is just the job it needs to do. It could be "precisified"—"Would the photo be nice if it were slightly lighter" or "Would the cabbage be nice if it were slightly lighter"—but it almost certainly does

not *need* to be made more precise in either of these scenarios for effective communication to take place.

In the passage from Frege quoted above, the notion of a *boundary* was introduced. The things to which a predicate applies were imagined within the boundary of that predicate, and the things to which it did not apply without the boundary. So far the classical view, but then there are the things that seem to be *on* the boundary—of which it cannot be distinguished, no matter how much one knows or can know about the situation, whether they are inside or outside the boundary.

It is a natural metaphor for describing such situations. Isaac Barrow (1700, 263), for instance, charmingly transfers the notion of a (physical) boundary between political states to notions of concern in ethics:

> As frequently between neighbouring States there do rise Dissentions and Contests about the just limits of their Territories; so doth it frequently happen between Vertue and Vice, Right and Wrong, Duty and Miscarriage in Practice; for although the extreme degrees, and even the middle regions of these things are very distant, yet the borders of them do lie very close together, and are in a manner contiguous; a certain ridge of separation running between them, which commonly (being very narrow, thin and obscure) it is not easie to discern. So it particularly falleth out in the matter before us, wherein our Text is concerned. Duty and Offence do nearly confine, and almost indiscernibly differ one from the other; for there are about this Case Precepts which seem to contradict; there are Duties appearing to thwart one another.

Two things are especially noteworthy in this passage. One is the qualification of a boundary as "very narrow, thin and obscure"—Barrow is not detained by concerns (quibbles?) about boundaries' being infinitely thin. The other is his observation that some things are well within the borders while others are very near to them. It is important not to lose sight of the fact that

even when a vague predicate is at issue, it may well still be possible to distinguish clear instances of things that have that feature (niceness, blueness) even if there are some for which it is more difficult or impossible.[2]

Returning to the notion of boundaries between states: it is sometimes agreed between countries that there will be a zone, in neither country, between them (this practice can take on ominous aspects when the two countries are hostile to one another, as recent history reminds us). Countries might agree to such an arrangement to solve a problem that otherwise would constantly arise, about when people or goods (or tanks) are in one country and when they are in the other. The agreement recognizes that the parties to it do not (perhaps cannot) agree below a certain threshold, and in order for them to carry on, relatively uninterrupted, with their business it has been agreed to allow a certain vagueness to persist in the meaning of "location of the border between our two countries."

It also becomes quite meaningful, in such circumstances, to speak of travelling *along* the border between the two countries, not just *across* it. That is, the boundary can itself be "thicker" than infinitely thin. This does not remove the vagueness of where each side of this agreed-upon region lies—to concern oneself with that is again to be back in the predicament of having to have *definite* boundaries somewhere, and that does not work, does not serve either party's purpose. Having established a vague boundary, the concerns about the vagueness of the edges of *that* boundary drop away.

One can also imagine that descendants of the original parties to the agreement about the zone no longer recall (if they were even taught) how the agreement came about. But they observe the conventions handed down to them, and thus are able to get along without engaging in serious wars.

Perhaps vague predicates are something like that: their users do not know (or need to know) exactly where their boundaries lie—though they can readily recognize an instance of something's being clearly within or without the extension of the predicate. And the conven-

tion of where the boundaries lie can be learned, through the examples used to teach the meaning of the predicates, closely enough that, most of the time, there is no confusion about the predicates' use (though there will be the occasional misuse or confusion).

Or perhaps it is something like this: Circe advises Odysseus (*Odyssey*, Book 12, around line 100) that he must make his way through a narrow strait[3] between the six-headed monster, Scylla, and the whirlpool created by Charybdis. Their proximity will force Odysseus to tack close to one or the other—"You shall then see the humbler rock, so close by/ That you could measure the distance with a dart." Circe recommends the Scylla side, since then Odysseus will lose only six of his crew, while he would lose his entire ship to Charybdis. Odysseus follows Circe's advice, with the predicted result (he prudently does not tell his crew what lies in store for them—not for nothing is his frequent epithet "crafty"). Thus the proverbial expression, "To shun Charybdis and fall into Scylla."

Now suppose, instead, that the strait was not so narrow, but very wide. Scylla and Charybdis ply their deadly trades, but Odysseus's ship can pass easily between them, with crew unharmed. What factors would then influence how Odysseus would plot his course through the strait? Perhaps the beauty of one cliff or the other, or concern for the smoothest passage to give his crew much-needed rest (or at least respite before resuming their trials).

For some vague predicates, it is very much in our interest to make them more precise, and not necessarily arbitrarily. To improve the precision of such a predicate is to improve our knowledge of, or communication about, the world in an important, useful way. Other vague predicates, however, can continue to allow for much variation in their significance: they are about features of the world that are not urgent for our purposes. In such cases it might be possible to make the predicate more precise, but it might also be impractical or undesirable to do so. One has correctly used a vague predicate when, in using it, one has stayed well within its boundaries on all sides.

We have many purposes when using our natural language. Some of them, especially when doing empirical science, are served by knowing the extension of our predicates as clearly as we can, and it can be worthwhile to make considerable effort to achieve the greatest clarity one can manage. But there are many other uses of language, and in some of these other uses it may not only be unhelpful to be as clear as one can about a predicate's extension, it might actually be a serious drawback. To have to specify the wavelengths of the light reflected by a surface to which the word "green" or "blue" might be applied in English would, in almost every case in which the expressions are used, just take up time uselessly. If in the early stages of negotiating a date for an important meeting I ask the participants if they are available "mid September," that can prepare them to make some time—two hours, for instance—available on a day in the middle of that month. As the time in which the meeting needs to take place for some purpose approaches, a specific day and time (September 12, at ten o'clock in the morning, say) can be agreed to. But the appointment need hardly be made more specific than that: the time does not need to be made more precise.

If there is a disagreement (as there was for the marble players described in Chapter One), then it can be negotiated on the spot. Or there can be an agreement to disagree, or one person can temporarily—for the duration of the exchange—adopt what appears to him the practice of the other in the interest of accomplishing together the purpose (a game, say) they might have in common.

Competent users of a natural language have, in virtue of that very competence, adequate means both to discover appropriate or "correct" uses for vague predicates and to negotiate successfully with other language users the significance of such predicates when the need for doing so arises. The problem of vague predicates is not currently susceptible to a solution in logic, but has usually been so in our everyday discourse.

FOUR

> The ideal, as we think of it, is unshakable. You can never get outside it; you must always turn back. There is no outside; outside you cannot breathe. —Where does this idea come from? It is like a pair of glasses on our nose through which we see whatever we look at. It never occurs to us to take them off.
>
> – Ludwig Wittgenstein, *Philosophical Investigations*, remark 103

> Doubt, usually, perhaps always, takes its rise from surprise, which supposes previous belief; and surprises come with novel environment. I will only add that though precise reasoning about precise experiential doubt could not entirely destroy doubt, any more than the action of finite conservative forces could leave a body in a continuous state of rest, yet vagueness, which is no more to be done away with in the world of logic than friction in mechanics, can have that effect.
>
> – The Pragmaticist to Doctor Y
> in C. S. Peirce's "Consequences of Critical Common-Sensism"

It is said that one understands a language well when one can appreciate humour in that language. A more prosaic criterion might be how well one understands and can apply a language's vaguest predicates.

Learning a language is a subtle and complex affair. There are words that can, at least at first, be understood by ostension or demonstration: "apple," "chair," "sit." When one is learning such words, it is expedient to remain, as it were, well inside the boundary of their extension. A tutor would be unhelpful and possibly outright confusing to point to a Granny Smith as the first example of an apple, or to the more exotic objects in an industrial museum—or to a dentist's chair or even a tree stump—when teaching the meaning of "chair," or to get down on the floor and place one's buttocks on her heels when demonstrating what sitting is.

Nevertheless, Granny Smiths *are* apples, and tree stumps (more controversially) might be, in some carefully delineated circumstances, chairs, and *seiza* is a form of sitting (in Japan, at any rate). Once rudimentary language skills are acquired, however, it can be explained to a learner that apples are fruits with certain features, one of which is usually the red colour (but there are green, and yellow, and variegated apples, too), and sitting is usually done on chairs, but in some cultures much more commonly on the floor or ground.[1]

The knowledge of the use of predicates like "heap" and "blue" is not especially difficult of acquisition. But neither is it exceptional to hear conversations like the one between the marble players reproduced in Chapter One. They are the kind of predicates the correct application of which will often be discussed, or even disputed, but most speakers will learn that people vary quite widely in the circumstances in which they will use such predicates, and adjust accordingly.

The puzzles invoked by patterns such as the sorites paradox are, in a sense, misunderstandings—not of the predicates per se, necessarily, but of their correct use, which is determined by many factors. Dummett's formulation of Wang's Paradox brings this out nicely. At the outset of our enquiry it may seem obvious what the meaning of "small" is, but what we really ought to notice from the puzzling conclusion of the sorites-type argument is not so much the logic (which, after all, seems fine), but that we are not as sure as we might have thought we were about the numerous, sometimes very thin, strands that have to be drawn together to give a good account of the use of "small." It is readily apparent, for instance, that the meaning of the word is contrasted with the meanings of "large," "tiny," "gargantuan," "minuscule," and the like. But none of these expressions has an absolute value—the meanings of the expressions seem to slide about, depending upon their application.

To search too determinedly for the logic of, or for, vague predicates is to miss an important aspect of what vague predicates are: they are about what is typical, usual, ordinary, familiar, contrasted with atypical, unusual, extraordinary, unfamiliar. None of these notions is, or

needs to be, clearly delineable. But some people, in their efforts to explore the wonders of language, wander into tiny and convoluted passages.

> For the sake of promoting clarity, philosophers often introduce distinctions that sunder what the contingent arrangements of this world (as we see them) have *in fact* conjoined. However, this produces not insight but problems. When we set facts aside, the concept at issue itself disintegrates in a destructive fission. And as we saw in the case of the *sorites* paradox, this disintegration generally manifests itself through an aporetic conflict of opposing arguments—all seemingly equally good but all, in the final analysis, equally unsatisfying. ... The language of everyday life is attuned to the prosaic, workaday practicalities of ordinary life and not to the requirements of theorizing philosophy. No matter how far we go in trying to explain, to specify, to qualify, we never manage to do more than to keep muddy waters astir. The processes of exposition never achieve the sought-for clarity. Problems, puzzles, and difficulties arise at every stage and, try as we will, they can never be eliminated altogether (Rescher 1994, 167–168).

This remark of Rescher's is intended to put forward his program of Standardism in philosophy, the view that good philosophy is done by eschewing the exotic and science-fiction example and trying to engage—often difficult enough as it is—the commonplace and typical. Rescher would look askance, I am sure, at the scenario set out at the end of Chapter One—it is difficult to imagine anyone actually facing such a bizarre choice. But Rescher's program goes awry in this respect: it is sometimes enquiring into the out-of-the-way, the oddball, that new insight about a problem is generated. We cannot know in advance which lines of enquiry will prove most productive. The farther a new idea diverges from the received view, the likelier it is to produce radical insight—if it turns out to be correct. At the same time, however, a wild conjecture is likelier to turn out to be nonsense. A claim by someone that he has trisected an

angle using only straight-edge and compasses is rightly received with skepticism—a huge amount of what we know about plane geometry tells us that that task is impossible—but if someone *could* show us how to trisect an angle, it would result in a radical and far-reaching change in how we conceive geometry.

The concepts we already have do not determine, but do affect, sometimes very strongly, the stretchings and developments of those concepts. Thus it is extremely difficult to relinquish the law of excluded middle or the principle of bivalence. They are very strong intuitions; we cannot imagine or feel how they could *not* be true: we think that everything we know would be called into question.

Radical changes in our network of concepts are almost impossible to achieve, at least quickly—there is an inherent, and probably evolutionarily prudent, conservatism to our ways of thinking about the world. An appropriate analogy here might be with the words of natural languages: over time, they change both their form and their meanings. To cast aspersions on someone used to be to bless him; now, the expression means to insult or denigrate a person or his views. We can give some account for how such a change might come about—in this case, due to dissatisfaction with some aspects of institutional religion, one of the rituals of which was taken to stand for the whole. It is possible, in a carefully circumscribed context, to use the expression in its first meaning, but the meaning now is, unless one is somehow put on notice to search for another, the second. This change could not have happened quickly—each speaker of English familiar with the expression at all would have had to have encountered the new, upended meaning by hearing or reading.

By contrast, it is the great advantage of formalized systems like logic that, in the relative safety of the imagination, one can try out just such radical assumptions. Just change an axiom, or remove or add one, and see what happens. That the results of such activity are occasionally beautiful structures is perhaps compensation for how often it they are useless in the everyday world.

Even though I conclude that logic as currently conceived does not have the way in to providing a formalized conception of vague predicates and a solution to the problems of their use, that is not to say that anyone should abandon the search—to advocate that is conservatism in the bad sense: stodginess and self-promotion or self-protection. Perhaps we can, like the mathematicians who are platonists on weekdays, formalists on Saturdays and not sure what they are on Sundays, be Rescheresque Standardists for some part of the week and relentless logicians the rest—with no one to say how much time we should spend being which. And if we are not certain that we have answered all, or even some, of our questions correctly, we can at least aspire to and hope that we ask the right questions.

Rescher is, in any event, correct in this assessment:

> Philosophical deliberations ultimately pivot on untidy concepts attuned to our practical dealings in a complex world where some degree of oversimplification is always required in the interest of manageability. In philosophy, we are constantly constrained to make rough-approximation statements—"promise-breaking is morally wrong," for example—indulging in generalizations that eventually need further qualification and amendment since what is claimed is not strictly and unexceptionally so (here, in certain cases of incapacity or of conflicts of duty), but will at best represent how matters stand in the normal course of things. Standardism accepts a relaxation demanded by the fact that exact and rigidly universal generalizations cannot meet the needs of the situation (Rescher 1994, 170–171).

Untidy our everyday concepts may be, but that makes them neither useless nor even necessarily difficult to use. We make distinctions when we have to—it is not that we have all the distinctions we will ever need ready-made in either our language's categories or in our experience of the world up to this moment. When we do have to make a distinctions, we do, not knowing whether we have carved the world at the joints. The expressions we use to indi-

cate those distinctions are used correctly when we use them within their boundaries—wherever they are; we do not know exactly where that is, but we can at least sometimes know that they are "over there," so to say, and our use of the expressions falls well within them. It is the best we can do within our perceptual and cognitive limitations.

NOTES

Notes to Chapter One

1. Dummett explains his name for this formulation thus: "The title [of Dummett 1975] relates to an article by Professor Hao Wang which I remember reading in an ephemeral Oxford publication many years ago. I should probably have abandoned it had I published the article sooner, since I never supposed that Professor Wang intended anything but to display the general form of a range of ancient paradoxes; but, since the name has gained some currency, I thought it better to leave it" (Dummett 1975, 324 note).
2. Note that how colour vision in human beings actually works is not known; it is, in any case, an extremely complicated process, as yet poorly understood.

Notes to Chapter Two

1. Perhaps all concepts, however abstract or formal, have their origin in human experience of the world, but in any event it seems eminently possible to modify them to such an extent that they have no obvious, or even tenuous, relation to the world from which they came.
2. See Gaines 1983 for the interesting suggestion that fuzzy logic and probability are different developments of a common base structure, and Thomas 1995 for the a recent attempt to renovate fuzzy set theory. Fine 1975 thinks the idea of continuum-many values on the closed real interval, the usual foundation of fuzzy logic, is a mistake (see especially p. 298, note 4).

Notes to Chapter Three

1. It is generally supposed that it must be necessary, in order for us to understand anything of the world at all, to have representations—"mappings" of some sort—of the things in it (and perhaps its features considered more abstractly) in our brains, or at least our minds (if one allows that the brain and mind might not be identical). It is not yet known how such representations are stored or manipulated, or whether there might be some other, quite different, way in which the brain/mind "understands" the world.
2. I owe this observation to M. F. Ereshefsky.
3. The location of the strait is not given by Homer, but in later, classical times it was identified with the Strait of Messina which separates Sicily and Italy. D. F. Hamilton has suggested to me in conversation that the actual model for the strait might have been in the Black Sea. The Strait of Messina is 24 miles long and varies between 2.5 miles and 14 miles in breadth; Scylla is supposed to have been on the Italian side.

Note to Chapter Four

1. For people of cultures that do not use chairs, sitting on one can be regarded as not really *sitting* at all, but more like *perching*, adopting a posture intermediate between sitting and standing.

BIBLIOGRAPHY

Anderson, Alan Ross and Nuel E. Belnap, Jr.

1975 *Entailment: The logic of relevance and necessity.* Vol. 1. Princeton: Princeton University Press.

Anderson, Alan Ross and Nuel E. Belnap, Jr. and J. Michael Dunn

1992 *Entailment: The logic of relevance and necessity.* Vol. 2. Princeton: Princeton University Press.

Barrow, Isaac

1700 "Sermon XXI: Of Quietness, and doing our own Business." In *The Works of the Learned Isaac Barrow, D.d., Late Master of Trinity-Colleg in Cambridge (Being All His English Works) In Three Volumes*, vol. 1, 263–272. London: Printed for Brabazon Aylmer, at the Three Pigeons against the Royal-Exchange in Cornhill.

Barwise, Jon

1989 *The Situation in Logic.* N.P.: Center for the Study of Language and Information. CSLI Lecture Notes Number 17.

Barwise, Jon and John Perry

1983 *Situations and Attitudes.* Cambridge, Massachusetts and London: The MIT Press.

Black, Max

1937 "Vagueness: An exercise in logical analysis." *Philosophy of Science* 4(4), October, 427–455.

Burns, Linda Claire

1991 *Vagueness: An investigation into natural languages and the sorites paradox.* Dordrecht, Boston and London: KluwerAcademic Publishers.

Cicero

[1933] 1956 *De Natura Deorum, Academica.* Translated by H. Rackham. London: William Heinemann Ltd and Cambridge, Massachusetts: Harvard University Press.

Dennett, Daniel C.

1991 *Consciousness Explained.* Boston: Little, Brown and Company.

Diogenes Laertius

[1925] 1972 *Lives of Eminent Philosophers*. Vol. 1. Translated by R. D. Hicks. Loeb Classical Library 184. Cambridge, Massachusetts, and London: Harvard University Press.

Dummett, Michael

1975 "Wang's Paradox." *Synthèse* 30(3/4), April/May, 301–324.

Fine, Kit

1975 "Vagueness, truth and logic." *Synthèse* 30(3/4), April/May, 265–300.

Gaines, Brian R.

1983 "Precise past—fuzzy future." *International Journal of Man-machine Studies* 19, 117–134.

Geach, Peter and Max Black (eds.)

1980 *Translations from the Philosophical Writings of Gottlob Frege*. Third edition. Totowa, New Jersey: Rowman & Littlefield.

Hisdal, Ellen

1985 "Reconciliation of the yes-no versus grade of membership dualism in human thinking." In M. M. Gupta, et al. (eds.), *Approximate Reasoning in Expert Systems*. N.P.: Elsevier Science Publishers B.V. (North-Holland).

Kleene, Stephen Cole

1938 "On notation for ordinal numbers." *Journal of Symbolic Logic* 3: 150–155.

Lee, H. D. P.

[1936] 1967 *Zeno of Elea: A Text, with Translation and Notes*. Amsterdam: Adolf M. Hakkert.

Lewis, Clarence Irving

1918 *A Survey of Symbolic Logic*. Berkeley: University of California Press.

Lewis, Clarence Irving and Cooper Harold Langford

1932 *Symbolic Logic*. New York: The Appleton-Century Company.

Łukasiewicz, Jan

1920 "O logice trójwartosœciowej" [On three-valued logic]. *Ruch Filozoficzny* 5: 169–171. [Reprinted in Storrs McCall (ed.), *Polish Logic 1920–1939*, Oxford: Oxford University Press, 1967.]

1922/23 "Interpretacja liczbowa teorii zdań" [Interpretation of the theory of propositions]. *Ruch Filozoficzny* 7: 92–93.

1930 "Philosophische Bemerkungen zu mehr wertigen Systemen des Aussagenkalküls." *Comptes rendus des séances de la Société des Sciences et des Lettres de Varsovie*, Classe III, 23: 51–77.

Łukasiewicz, Jan and Alfred Tarski

1930 "Untersuchungen über den Aussagenkalkül." *Comptes rendus des séances de la Société des Sciences et des Lettres de Varsovie* 23, Classe III, 30–50.

Nadin, Mihai

1983 "The logic of vagueness and the category of synechism." In Eugene Freeman (ed.), *The Relevance of Charles Peirce*, 154–166. La Salle, Illinois: The Hegeler Institute/Monist Library of Philosophy.

Nagel, Thomas

1986 *The View From Nowhere*. New York: Oxford University Press.

Nagel, Ernest and James R.Newman

[1956] 1978 Gödel's Proof." In Irving M. Copi and James A. Gould (eds.), *Contemporary Philosophical Logic*, 14–34. New York: St. Martin's Press.

Peirce, Charles Sanders

1960 *Collected Papers of Charles Sanders Peirce*, Vol. 5/6. Edited by Charles Hartshorne and Paul Weiss. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

Pelletier, Francis Jeffry and István Berkeley

1995 "Vagueness." In Robert Audi (ed.), *The Cambridge Dictionary of Philosophy*, 826–828. Cambridge: Cambridge University Press.

Post, Emil L.

1921 "Introduction to a general theory of elementary propositions." *American Journal of Mathematics* 43: 163–185.

Priest, Graham and Routley, Richard

1989a "Applications of paraconsistent logic." In *Paraconsistent Logic: Essays on the inconsistent*, 367–393. Edited by Graham Priest and Richard Routley and Jean Norman. München, Hamden and Wien: Philosophia Verlag.

1989b "Systems of paraconsistent logic." In *Paraconsistent Logic: Essays on the inconsistent*, 151–186. Edited by Graham Priest and Richard Routley and Jean Norman. München, Hamden and Wien: Philosophia Verlag.

Rescher, Nicholas

1994 *Philosophical Standardism: An empiricist approach to philosophical methodology*. Pittsburgh and London: University of Pittsburgh Press.

Russell, Bertrand

1905 "On denoting." *Mind* 14, 484–485.

1923 "Vagueness." *The Australasian Journal of Psychology and Philosophy* 1, 84–92.

Shainberg, Lawrence

1995 *Ambivalent Zen: A memoir*. New York: Pantheon Books.

Suchman, Lucy A.

1987 *Plans and Situated Actions: The Problem of human–machine communication*. Cambridge: Cambridge University Press.

Tanahashi, Kazuaki and Tensho David Schneider (eds.)

1994 *Essential Zen*. San Francisco: HarperCollins.

Thomas, S. F.

1995 *Fuzziness and Probability: An essay on the foundational nexus among semantics, measurement, uncertainty, and inductive and deductive inference, with application to decision analysis under uncertainty*. Wichita, Kansas: ACG Press.

Wajsberg, Mordechaj

1931 "Aksjomatyzacja trójwartościowego rachunka zdań" [Axiomatisierung des dreiwertigen Aussagenkalküls] [Mit deutsch. Zusammenfassung]. *Comptes rendus des séances de la Société des Sciences et des Lettres de Varsovie*, Classe III, 24: 126–148.

Wang, Hao

[1955] 1978 "On formalization." In Irving M. Copi and James A. Gould (eds.), *Contemporary Philosophical Logic*, 2–13. New York: St. Martin's Press.

Wittgenstein, Ludwig

[1958] 1992 *Philosophical Investigations*. Translated by G. E.M. Anscombe. Second edition. Oxford and Cambridge, Massachusetts: Blackwell.

Zadeh, Lotfi A.

1965 "Fuzzy Sets." *Information and Control* 8, 338–353.