Concepts Without Boundaries[1]

0

Philosophers have been interested in vagueness for centuries. One reason is the fascination, and threat, posed by the so-called sorites paradoxes. If someone is not bald, then he does not become bald by losing just a single hair. But then, it seems, however many hairs he loses, he can never become bald. No hair's loss marks the transition; so, it seems, there can be no transition. We know that the conclusion is false. The problem is to say how it can be avoided.

Vagueness is of interest independently of the paradoxes. It seems to be an extremely pervasive phenomenon, invading almost every area of thought, and banished from scientific work, if at all, only by constant vigilance. What is its origin? Does it correspond to a feature of the world? Or is it we, perhaps through our deficiencies, who are responsible? And is it obvious that it is a Bad Thing, given the extent to which the throbbing centres of our lives appear to be describable only in vague terms?

A more preliminary question is: what is vagueness? The standard definition is that a vague word is one which admits borderline cases. I agree that if a word is vague, then there are or could be borderline cases; but I deny the converse: non-vague expressions, too, can have borderline cases, so we do not yet have a grasp of the essence of vagueness. That essence is to be found in the idea that vague concepts are concepts without boundaries.

I

Some concepts classify by setting boundaries but some do not. In the philosophical tradition, the former have received all the attention, and have lent a distinctive character to

[1] Inaugural lecture delivered at King's College London, 1990. Reprinted in Rosanna Keefe and Peter Smith (eds) *Vagueness: A Reader*, MIT Press, Cambridge MA, 1996, p. 251–264.

attempts to study classificatory concepts and their linguistic correlates. Within what I shall call the "classical picture", a picture which dominates most thinking about thought and language, there is no room for the thesis I wish to put forward: that concepts can classify without setting boundaries.

According to this classical picture, the job of classificatory concepts is to sort or segregate things into *classes* by providing a system of pigeon-holes, by placing a grid over reality, by demarcating areas of logical space. Boundaries are what count, for a concept must use a boundary to segregate the things which fall under it from the things which do not. This intuitive view receives definitive expression in classical semantic theories. A predicate, linguistic vehicle of a concept, is thought of as having a meaning which fixes its extension, the set of things of which it is true. A semantic theory will provide an at least partial characterization of a predicate's meaning by specifying this set in some appropriately revealing fashion. In the light of such specifications, one can model the logical features of the language in which the predicate occurs by generalizing over the sets which predicates do, or can, determine.

Thus the classical picture, informed by a connection between concepts and sets present in the very word "classify", sees the theoretical resources of set theory as the proper instruments for describing language and thought.

Classes, and sets, have sharp boundaries. Hence, at least *prima facie,* a description of concepts or predicates in terms of what sets they determine is a description of them as boundary-drawers. This mode of description, and the picture which underlies it, thus makes no room for concepts without boundaries, those which are not boundary-drawers. We have a choice: we could take the classical picture as exhausting the ways classificatory concepts can be, and conclude that there are no concepts without boundaries; or else, convinced that there are such concepts, we could reject the comprehensiveness of the classical picture. What I suggest is

that almost all concepts lack boundaries, so that the classical picture is of very little use to us.

The concepts which classify without setting boundaries include the ones traditionally counted as vague: *red, heap, child, bald,* to take some famous examples. The first stage in showing that these are boundaryless concepts involves showing that there is no set of which they are true: they do not classify at all, if the only way to classify is to assign things to classes or sets.

<div align="center">II</div>

Sets have sharp boundaries, or, if you prefer, are sharp objects: for any set, and any object, either the object quite definitely belongs to the set or else it quite definitely does not Suppose there were a set of things of which "red" is true: it would be the set of red things. However, "red" is vague: there are objects of which it is neither the case that "red" is (definitely) true nor the case that "red" is (definitely) not true. Such an object would neither definitely belong to the set of red things nor definitely fail to belong to this set. But this is impossible, by the very nature of sets. Hence there is no set of red things.

This seems to me as certain as anything in philosophy, yet it can often be a bitter pill to swallow, by non-philosophers and by philosophers alike. In some debates about abortion, one can feel a real sense of shock at the realization that there is no set of persons: the concept *person* is vague at just the relevant point. The difficulty is that moral concepts are often boundary-drawing (especially so the more naive the morality), and legal concepts typically have to be. Trying to tie the application of a boundary-drawing concept (as *may legitimately be aborted* is supposed to be) with a boundaryless one like *is a person* poses a problem which is simply not soluble in the straightforward terms in which it is often posed. A quite general reflection on boundaries, their absence and presence, should reshape what one could expect to emerge from a discussion of abortion and the law.

Philosophers, too, are attracted by the classical picture, even when not engaged upon formal semantic projects. To take one example from a million, when Peter Strawson some time ago considered whether "Socrates is human" is equivalent to "Socrates belongs to the set of humans" his many interesting observations did not include the point that the equivalence must fail since there is no set of humans.

If vague predicates and vague concepts do not have "extensions" — sets of things of which they are true — they do not draw boundaries, at least not in any simple sense. For a boundary should divide things into two sets, those which fall on one side and those which fall on the other. So if vague predicates do not effect a division into sets, they draw no boundaries.

But, it may be objected, this is too simple. For may not a vague predicate draw an unsharp boundary? May not an unsharp boundary work without dividing into two? May there not be things whose status is unresolved by the boundary?

When one says that a vague predicate does not draw *sharp* boundaries, "sharp", I believe, does no work, for there is only one kind of boundary. Hence we cannot regard vague predicates as drawers of boundaries, but ones which are unsharp. I shall establish this by the following route. Anything worthy of the name *boundary* will effect set theoretically describable divisions, even if more complex ones than the simple twofold division envisaged just now. But any such division, however complex, will misdescribe the functioning of a vague predicate.

III

I shall start with a general reason for thinking that there is no adequate set-theoretic description of vague predicates. It is that a set-theoretic description of a language typically ends up identifying a set of truths. (More exactly, it specifies for each sentence a condition upon which the sentence will belong to this set, or to the set of truths-upon-S, for a relativization to some structure, S.) But the argument which showed that there

was no set of red things shows also that there is no set of truths in a language which contains the predicate "red". For consider a sentence ascribing *red* to a borderline case. This sentence will neither (definitely) belong nor (definitely) fail to belong to the set of such truths; which is another way of saying that there is no such set.

But surely (one may object) this simple thought does not do justice to all the cunning twists and turns which set-theorists have made in their attempts to describe vagueness. What about fuzzy logic? What about supervaluation theory?

I believe that the simple minded point of a moment back - that there is no set of truths - does ultimately carry the day against such theories. But I recognize that to convince an opponent involves more detail.

The preliminary claim I need to make is that one cannot do justice to the phenomena of vagueness, in particular to phenomena called "higher order vagueness", simply by increasing the number of sets of individuals associated with a predicate. A set theoretic description might start by associating a vague predicate not with two sets but with three: the set of things of which the predicate is (definitely) true, the set of things of which it is (definitely) false, amd the remainder, the set of borderline cases. So far, so good: a sharp predicate has two extensions, a positive extension, and a negative one (its complement within the domain) whereas a vague predicate has three, a positive one, a negative one, and a penumbral one (the complement within the domain of the union of the other two).

But a predicate which effects such a threefold partition is not vague. This fact, which shows why one cannot characterize vagueness merely in terms of borderline cases, follows from the fact that the partition does identify a set of truths, which we have seen to be inconsistent with vagueness. Once again more detail will reinforce the point. Consider a child developing into adulthood. We cannot associate "child" with a twofold division because there is no set of children (at a time). There is no such set because of the borderline cases. Thus far,

some encouragement for a threefold partition, which explicitly allows for borderlines. However, essentially the very point which scuppers a twofold partition also scuppers a threefold one: the latter posits a set of children, or at least a set of definite children. Yet it would be as absurd to suppose that a heartbeat could make the difference between membership of this set, and consignment to the set of borderline cases, as it would be to suppose that a heartbeat could make the difference between belonging to the set of children and belonging to the set of non-children. Childhood, even definite childhood, fades gradually away, and does not come to a sudden end. This is not to deny that we can by convention stipulate a sharp boundary; it is only to say that our concept *child* does not supply one.

A proponent of set theoretic divisions might seek to meet this point by making more divisions. He might say that a predicate is sharp, that is, not vague at all, or as I shall say is vague_0, just on condition that it draws a single boundary, thus partitioning the domain into two sets; that a predicate possesses the lowest level of vagueness, is vague_1, just on condition that it draws two boundaries, partitioning the domain into three sets. If it is unacceptable to suppose that "child" or "red" does this (on account of the unacceptability of supposing that there is a last heartbeat of one's childhood, or of one's definite childhood), then the correct description of such a predicate must look further up the hierarchy: perhaps it draws four or forty boundaries, making correspondingly many partitions of the domain.

The generalization of this set-theoretic approach is that a predicate is vague_n iff it draws 2^n boundaries, thus partitioning the domain into 2^n+1 sets. A predicate is sharp iff it is vague_0; is vague iff it is vague_n for some positive n; is higher order vague iff it is vague_n for some $n>1$, and is radically vague iff it is vague_n for all n. The phenomena mentioned a moment back - such facts, now agreed by the set-theorist, as that there is no sharp division between children and borderline cases of

children - are to be described, on the envisaged approach, by going higher in the hierarchy of higher order vagueness. If the set-theoretic description assigns a cut-off where none can be found in the actual use of the predicates, then the description has not set the level of vagueness high enough. The hope is that the unlimited upwards mobility is enough to enable one to get on top of all the phenomena.

This hope, however, is groundless. Indeed, its very structure should be unappealing: you do not improve a bad idea by iterating it. In more detail, suppose we have a finished account of a predicate, associating it with some possibly infinite number of boundaries, and some possibly infinite number of sets. Given the aims of the description, we must be able to organize the sets in the following threefold way: one of them is the set supposedly corresponding to the things of which the predicate is absolutely definitely and unimpugnably true, the things to which the predicate's application is untainted by the shadow of vagueness; one of them is the set supposedly corresponding to the things of which the predicate is absolutely definitely and unimpugnably false, the things to which the predicate's non-application is untainted by the shadow of vagueness; the union of the remaining sets would supposedly correspond to one or another kind of borderline case. So the old problem re-emerges: no sharp cut-off to the shadow of vagueness is marked in our linguistic practice, so to attribute it to the predicate is to misdescribe it.

IV

In effect, this same point is what scuppers the set-theoretic descriptions of vague languages offered by fuzzy logicians and by supervaluations theorists.

The fuzzy logic I envisage associates with each predicate not a set whose only members are individuals, the individuals of which the predicate is true, but what Goguen calls a "J-set": a function from each object in the domain of discourse to a real number in the open interval 0,1. The number, as value to the function, represents the degree to which the predicate is

true of the object which is argument to the function. The real numbers are continuous, so surely now we have a mode of description apt for vagueness, one which does not necessitate sharp boundaries? Yet a fuzzy set is a genuine set, a completely sharp object.

The reason for thinking this hope groundless is essentially the same as that given earlier: the fuzzy logician, too, will (whether he likes it or not) be committed to a threefold partition: the sentences which are true to degree 1, those true to degree 0, and the remainder. But to what in our actual use of language does this division correspond? It looks as if, as before, it should correspond to the sentences true beyond the shadow of vagueness, those in some kind of borderline position, and those false beyond the shadow. But, as several times noted, we do not know, cannot know, and do not need to know these supposed boundaries to use language correctly. Hence they cannot be included in a correct description of our language.

Fuzzy logic does, indeed, describe a feature of our use of many vague predicates: that, and how, they are associated with a dimension of comparison. Our use of "red" is properly regarded as regulated by such principles as that anything redder than a red thing is red. The "redder than" relation is tracked by fuzzy logic's numerical ordering, which in turn bears straightforwardly on the applicability of "red". For many predicates, things are more complex: many relations are relevant to applicability. Childhood is affected by age, but also by many other factors as well, so that two individuals of the same age may differ in point of how much of a child they are. Yet, plainly, fuzzy logic could be supplied with the resources to describe the nature and weights of a whole complex of applicability-determining relations. Could we not somehow reap these benefits without succumbing to fuzzy logic's threefold partitioning?

Here is one possible line of thought. One can expect the empirical data for fuzzy logic to be quite messy. One might run trials in which one asked people to do two kinds of thing: order

all the objects in terms of their strength as candidates for application of the predicate; and identify the definite cases and the definite non-cases The results would be variable, both intra- and inter-personally. Let us say that an *admissible* J-set for a predicate is one which matches some trial for that predicate both in point of order and in point of definite cases. That is to say, if in the trial the subject treats α as a better case of the predicate than β, then the J-value for α must exceed that for β; if the subject identifies an object as a definite case for the predicate, the J-value for that object must be 1; if the subject identifies an object as a definite non-case for the predicate, the J-value for that object must be 0. When it comes to specifying truth, the theorist could adopt a vague definition, for example he could say that an atom $\phi\alpha$ is true iff the object denoted by α has J-value 1 for almost all J-sets admissible for ϕ.

The general idea behind the strategy is to take the set-theoretic description as a kind of basis, and exploit it in a vague way to deliver an account of whatever one takes to be the central semantic notion.

The idea is not unattractive, but it does not fall within the scope of the the approach I wish to attack. For what is envisaged is that the real work of describing the functioning of a predicate is done not by fuzzy logic itself, but in terms of some *vaguely* specified semantic notion. The proposal, then, ends up as not one in which a predicate is described by being associated with an extension or with boundaries.

A similar series of moves can be made in connection with supervaluations. The supervaluational theory, too, will end up making a threefold partition: the set of sentences true-upon-all-sharpenings, the set of sentences false-upon-all-sharpenings, and the remainder. An attempt to do justice to "higher order" vagueness by acknowledging vagueness in the notion of a sharpening would force one outside the set theoretic language in which the theory is supposed to be couched, and would mean that the real work of semantic

description was being done in a vague language, rather than in a set-assigning one.

∨

A vague concept is boundaryless in that no boundary marks the things which fall under it from the things which do not, and no boundary marks the things which definitely fall under it from those which do not definitely do so; and so on. Manifestations are the unwillingness of knowing subjects to draw any such boundaries, the cognitive impossibility of identifying such boundaries, and the needlessness and even disutility of such boundaries.

To characterize a vague concept as boundaryless is an improvement on characterizing it as one which permits borderline cases, since a non-vague concept may admit borderline cases. If "child*" is defined as true of just those people whose hearts have beat less than a million times, false of those whose hearts have beaten more than a million and fify times, and borderline with respect to the remaining people, it has borderline cases but behaves quite unlike our paradigms of vagueness.

A boundaryless concept cannot be described in set theoretic terms. How can it have a classificatory role? How is it to be described, either semantically or in terms of cognitive processing? How are the paradoxes of vagueness to be avoided?

Scepticism about whether boundaryless classification is possible can be set to rest, I believe, by contemplating a very familiar case: the colour spectrum, as displayed, for example, in an illustration in a book on colour. Looking carefully, we can discern no boundaries between the different colours: they stand out as clearly different, yet there are no sharp divisions. There are bands, but no bounds. This does nothing to impede the classificatory process: the spectrum is a paradigm of classification.

The image of pigeon-holes is powerful. Is there a comparable one which would represent how boundaryless concepts classify? We could, perhaps, think of such concepts

(10)

as like magnetic poles exerting various degrees of influence: some objects cluster firmly to one pole, some to another, and some, though sensitive to the forces, join no cluster.

At least one aspect of this image deserves more literal statement. Boundaryless concepts tend to come in systems of contraries: opposed pairs like child/adult, hot/cold, weak/strong, true/false, and the more complex systems exemplified by our colour terms. This is a natural upshot of boundarylessness, as we can see by reflecting on what is involved in grasping a concept.

Such a grasp, it must be agreed on all sides, involves knowing how something would have to be for the concept to apply to it, and how something would have to be for the concept not to apply. A distinctive feature of the classical picture is that it takes this latter fact as primitive. Grasping what a concept excludes is part of grasping the concept, and is achieved through the mediation of no other non-logical concept. Hence it is very natural to see the division between what a concept includes and what it excludes in terms of a boundary. Certainly, perception of a boundary would be enough; but the proponent of boundarylessness will insist that it is not the only way.

On the alternative picture, what a concept excludes is graspable in a positive way, mediated by other contrary concepts. A grasp of *red* attains grasp of what is not red at a derivative level, via a grasp of *yellow, green, blue* and so on. A system of such concepts is grasped as a whole, as can be seen in the way paradigms are used in learning. There are paradigms of red, but nothing is non-derivatively classifiable as a paradigm of not-red. Any paradigm of another colour will serve as a paradigm of how not to be red, but only in virtue of its positive classification as another colour.

Not just any clear case of the non-applicability of a concept will serve to help a learner see what the concept excludes. Television sets, mountains and French horns are all absolutely definite cases of non-children; but only the contrast with *adult* will help the learner grasp what *child* excludes. So it is no

accident that boundaryless concepts come in groups of contraries. Correlatively, the image of attracting poles, replacing the classical image of pigeon-holes, is not without value.

It also serves to record some empirical data. For example, subjects asked to classify a range of test objects using just "young" and "old" make different assignments to these words from those they make to them when asked to classify using, in addition, "middle-aged". The introduction of a third magnetic pole can attract some of the things only loosely attached to two existing ones, without diminishing the forces the existing ones exert.

VI

Let me now turn to the question about paradox, which might take a more aggressive form: does not the very notion of boundarylessness make the paradoxes unavoidable? For the absence of a boundary has been treated as the impossibility of very similar things differing in point of the applicability of a predicate. But then it seems that we can form a sorites series of objects, adjacent pairs being too similar to merit a difference in applicability, but remote pairs being sufficiently dissimilar to require a difference. Starting with a clear case of red, we assemble closely resembling patches in a series through which the colour shifts gradually towards yellow. The boundarylessness of red is supposed to ensure that there is no adjacent pair of which the first is red and the second not. The first member, by hypothesis, is red. Hence, by boundarylessness, the one adjacent to it is also red; and so on. So does not familiar reasoning lead inexorably to the intolerable conclusion that all members of the series are red, even the yellow ones?

This worry can take a form which can only be assuaged by a technical and formal semantic theory, and nothing of that kind is on offer here. However, let me a simple observation which should establish that the present picture of vagueness as

boundarylessness is no less well placed than any other to come to terms with sorites reasoning.

The classical picture has a totalitarian aspect: there is no difference between its being not mandatory to apply a concept and its being mandatory not to apply it. If the very nature of the concept *prime*, together with the nature of some number, say eight, does not require you to apply the concept to it, then the very nature of the concept, together with the nature of the number, requires you not to apply the concept to it. For a rational and fully informed thinker, there is no freedom.

By contrast, vagueness offers freedom. It can be permissible to draw a line even where it is not mandatory to do so. No one can criticize an art materials shop for organizing its tubes of paints on various shelves, including one labelled "red" and another "yellow", even though there is a barely detectable, or perhaps even in normal circumstances undetectable, difference between the reddest paint on the shelf marked "yellow" and the yellowest paint on the shelf marked "red". Hence one can consistently combine the following: *red* draws no boundaries, that is, there is no adjacent pair in the series of tubes of paint such that the nature of the concept, together with the colour in the tube, requires one to apply *red* to one member of the pair but withold it from the other; yet one can draw a boundary to the reds, that is, one may behave consistently with the nature of the concept in drawing a line between adjacent pairs.

The envisaged attack on boundarylessness can be set out as the following argument, which makes plain how the recent observation addresses it. A boundaryless concept is one which, for closely similar pairs, never makes it mandatory to apply the concept to one member of the pair, and withold it from the other; hence, the argument runs, a boundaryless concept is one which, for closely similar pairs, makes it mandatory never to apply the concept to one member of the pair, and withold it from the other. The inference depends upon the move from something being not mandatory to its

being forbidden; a move legitimate within the totalitarianism of boundary-drawing concepts, but not within the liberality of boundarylessness.

I do not suggest that this simple observation puts an end to the lure of sorites reasoning, which, like a virus, will tend to evolve a resistant strain. Must there not be an outer limit to the things to which it is mandatory to apply "red", and a first member of the sorites series with respect to which we have licence to withold? The answer is "No: 'mandatory', too, is boundaryless"; though I shall not now stop to show how this answer can be justified. It is enough for the moment to have shown that boundarylessness should give no special encouragement to paradoxical sorites-style reasoning.

VII

If standard set-theoretic descriptions are incorrect for boundaryless concepts, what kind of semantics are appropriate? A generalization of the considerations so far suggests that there is no precise description of vagueness. So what kind of description should be offered? More pointedly, I hear a certain kind of objector say: we can't even tell what boundarylessness is until you give us your semantics.

If driven in this way, I would urge an idea of Donald Davidson's. A semantic theory can quite legitimately be *homophonic,* that is, can reuse in the metalanguage the very expressions whose object-language behaviour it is attempting to characterize. Asked how a boundaryless predicate like "red" works, my first response would be: "red" is true of something iff that thing is red.

Whether or not vagueness is at stake, this Davidsonian idea has met with resistance. No one could claim that such remarks are untrue, but they have been held to be trivial or unilluminating. Many such objections are confused, or are based on a misunderstanding of what, by Davidson's lights, a semantic theory should aim to achieve. In his view it should enable us to understand how some expression conspires with others to fix a truth condition, an understanding which would

answer the question: "what are these familiar words doing here?"; and it should supply an account of what it would be enough for a speaker to know, in order to understand a language. Homophony impedes neither aim. On the contrary, it can provide a check upon their successful accomplishment. But a homophonic semantics for vague expressions could lead to two more specific objections, one misguided, but one sound and important.

The misguided one is that homophonic semantics will fail to make explicit which predicates are vague. In fuzzy logic, for example, the distinction will emerge in the structure of J-sets. Those associated with vague predicates will, for some or many objects as arguments, deliver numbers intermediate between 0 and 1 as values; whereas a sharp predicate's J-set will have 0 or 1 as the only values, whatever objects are arguments. In a homophonic theory the information is in a way present, but is inexplicit. It is present, because in specifying the applicability of, say, "bald" homophonically, you specify it vaguely and so *as* vague. But it is inexplicit, since there is no rule, usable by one who did not yet know which predicates are vague, on the basis of which, together with the semantic theory, he could pick out the vague ones.

However I know of no reason why this fact in itself counts against the homophonic approach. First, I know no reason for thinking that all such information must be made explicit, if the mentioned aims of semantics are to be achieved. Secondly, even if the mentioned aims, or others, did require that the information be made explicit, I know of no argument to establish that this cannot be achieved simply by means of a list, and thus consistently with homophony.

The sound and important objection is that the homophonic approach fails us in connection with logic. Davidson himself envisages a first-order metalanguage, and thus a metalanguage of which classical model theory is true, and thus a metalanguage in which predicates are associated with sets as their extensions. Thus envisaged, the project succeeds in fixing a logic for the object language, namely, classical first

order logic; but it fixes it while at the same time mischaracterizing the semantics of the object language. For, despite syntactic homophony, if the metalanguage is first order, its predicates will be boundary-drawing, and so will misrepresent the object language predicates as also being boundary-drawing.

This first order feature of Davidson's proposals is, as he says, inessential. Abandoning it makes possible serious homophony: an account of the object language predicates in which they are not merely reused in point of their sounds or shapes, but also in point of their meaning. The problem then, however, is to say something worthwhile about the logic of the object language. There are two obstacles. First, we do not know what our actual logic, which would be reapplied homophonically, is. We do not know, for example, whether every instance of *P or not-P* is counted true in our language and thought, and one pertinant reason for this doubt stems from vagueness. Secondly, even if we knew what our actual logic is, we could not uncritically reuse it in a semantic project, for the existence of sorites reasoning casts doubt upon whether we are right to subscribe to the logic to which we actually subscribe.

The logic of vagueness, characterized as boundarylessness, thus remains to be described. I believe that the way forward involves taking the notion of a vague object as basic; but this is a suggestion I shall not pursue here.

VIII

If the semantic description of a vague concept is to have all the thinness of homophony, can we not achieve a richer description in other ways, perhaps in terms of the psychological mechanisms whereby a vague concept is acquired or applied?

For example, it seems that very often a boundaryless concept is acquired on the basis of paradigms. We acquire the concept from the inside, working outwards from central cases, and locating the central cases of contrary concepts, rather

than from the outside, identifying boundaries and moving inwards. Can this thought be used to say anything illuminating about the nature of boundaryless concepts?

If it can, I am not sure how. Perhaps we should try to specify a boundaryless concept's relation to the world in terms of a paradigm - an object, α, to which it quite definitely applies, and which might therefore be an appropriate example to use in a teaching situation - together with a relation of similarity. Then we might say, non-homophonically, that the concept is true of something iff that thing is sufficiently similar to α.

But the suggestion has many vices. For one, it presupposes, without any justification, that every boundaryless concept must be instantiated. For another, the condition will not state anything which users of the concept have to know, since there may be more than one paradigm, and so one could master the concept just as well without knowing anything of α. Thirdly, anything which might have worried one about boundarylessness, for example any problems about the sorites paradox, naturally remain just as they were, though now attaching to the similarity relation.

Surely, however, I should take note of, and perhaps make use of, psychologists who have, it might be thought, investigated essentially this phenomenon under the name of prototype theory. Eleanor Rosch, for example, has suggested that the notion of a prototype helps us to understand vagueness since prototypicality is a property of degree, and vague predicates are asssociated with such properties. However, it turns out that prototypicality, in this sense, is orthogonal to vagueness, as demonstrated by the fact that an absolutely definite case may have low prototypicaltiy (as penguins do relative to their classification as birds). Indeed, even boundary-drawing concepts induce prototypicality scales. Thus 2 is highly prototypical for *even number,* but it is no surprise to learn that there are plenty of even numbers which have a very low prototypicality rating for this concept: many

even numbers are extremely unlikely to be chosen in teaching or exemplifying the concept.

A more promising alternative source of understanding boundarylessness is the parallel distributed processing (PDP) model of the material basis of cognitive activities. Crucial to such a model is another notion of a prototype: an object which has played a causal role as a positive instance in so adjusting the weights of the hidden elements of the network as to help tune it to its recognitional task. This might turn out to correspond more closely to what, a moment ago, I called a paradigm.

Any attempt to describe boundarylessness in such psychological or neurophysiological terms will, however, miss the normative features. We might explain confident application of, say, "red" as the organism's response to a high level of activation of the output of the red-recognizing network, and a reduction in confidence by a combination of a reduction of the level of this output and an increase in the level of output from the yellow-recognizing network. There may also be explanations in neurophysiological terms of the tendency to include more colours in the reds if you start from reds and move gradually to the yellows than if you reverse the direction of application. But no such facts will begin to capture such aspects of the use of the word "red" as the mandatoriness of its application in some central cases, the freedom available for borderlines, and such rules as that anything at least as red as a red thing is not merely likely to be called "red" but ought to be so called.

The general point is that the vagueness of a vague expression need not, and perhaps should not, feature in a psychological account of how it is used. This account may describe the dispositions-at-a-time to use the expression in terms of a probability function, and may describe a more enduring state in terms of a range of such functions. But the psychologist's task would be made no harder if he resisted anything homophonic in describing the inputs to the functions. Thus he might describe the input to a person's function

relating to the ascription of "red" not in our colour vocabulary but in terms of the physical constitution of the light striking the eye. So while one should expect harmony between the semantic fact of boundarylessness and such psychological descriptions, the latter can never exhaust the former.

IX

Let us take stock. Vagueness should be characterized as boundarylessness, not merely in terms of borderlines. Boundarylessness cannot be described sharply, for example set-theoretically; so, whatever insight psychological descriptions may offer, the only semantic description which appears plausible is vague, for example homophonic. We must reject the classical picture of classification by pigeon-holes, and think in other terms: classifying can be, and often is, clustering round paradigms.

Just how widespread vagueness is can be underestimated. Let me draw attention to an area sometimes wrongly thought to be free of it: biological species. Even in quite recent philosophy, there is a tendency to suppose that species come in the "eternal and fixed forms" beloved, according to John Locke, of the Port Royal logicians. It may seem that *strawberry* draws boundaries, since there are no borderline cases. But this is just an accident. There could very well be, and no doubt with the advent of genetic engineering soon will be, a series of plants between strawberries and raspberries, many of them borderline for both concepts. Such concepts do not impose boundaries, but constitute one of the largest and most impressive systems of contrary boundaryless concepts. Locke was right to draw attention to the lack of boundaries by reminding us of boundary-defying "monsters".

One practical application of work on vagueness is in cognitive science, where a possible goal is to implement in machinery the vagueness of our concepts. Another application has already been mentioned. The law must rule a boundary between legitimate and illegitimate acts. Here, boundarylessness would be out of place. Yet such rulings must

often traverse territory spanned by a boundaryless concept, like that of being a person. Given the nature of boundarylessness, semantics give freedom. There is some number of minutes such that the nature of the concept of a person, together with the nature of the world, makes it neither mandatory nor impermissible to apply the concept to a foetus of that age in minutes. Hence arguments that use the vague concept to establish or overthrow a sharp ruling are alike inadequate. We can no more argue that aborting a foetus of this age is right because it is not a person than we can argue that it is wrong because it is a person, if *person* is vague at the crucial point. In general, only a pragmatic justification could be found for drawing a legal line in an area where there are no relevant boundaries.

I mention this merely as an example of a possible application whose details remain to be worked out. It proleptically exemplifies my hope that work in the philosophy of vagueness will enable us better to understand how the demands of law and morality should be tailored to the boundaryless fabric of most of our thought and talk. More generally, I hope that there will always be a way whereby philosophical research, however arcane, will feed into our aspiration to understand ourselves and our world, and that this understanding will, as Susan Stebbing hoped, enable us to improve both.