

THE PHILOSOPHICAL PROBLEM OF VAGUENESS

DOROTHY EDGINGTON

University College, Oxford

I.

Think of the color spectrum, spread out before you. You can identify the different colors with ease. But if you are asked to indicate the point at which one color ends and the next begins, you are at a loss. “There is no such point,” is a natural thought: One color just shades gradually into the next.

This is an illustration of the vagueness that infects a great deal of our vocabulary: old, bald, tall, intelligent, wicked, a case for such-and-such legal verdict or medical treatment. . . . There are cases where these predicates clearly apply or clearly fail to apply. But there are (it seems) no sharp cutoff points. There are borderline cases in which they neither clearly apply nor clearly fail to apply. And it is not even clear where the clear cases end and the borderline cases begin.

This causes problems. You are casting for a play and need a tall man for a part. If x will do for height and y is only a millimetre shorter, then y will do for height. Remove one grain of sand from a heap and you still have a heap. A man who is not bald and loses just one hair, is still not bald. These judgments are plausible. But they lead to paradox. The “sorites paradox,” literally the “paradox of the heap,”¹ is now the name of the whole family of such paradoxes. Apply the above judgments often enough, and you can conclude that a midget is tall, that one grain of sand makes a heap, that a man with no hairs is not bald.

I heard someone say on the radio recently, in a discussion about whether it ought to be illegal for parents to smack their children: “There’s all the difference in the world between a cautionary slap on the wrist and a vicious beating.” So (granting a little hyperbole) there is. But a series of actions, each insignificantly different from its predecessor, can take us from one to the other. Where to draw the line between the acceptable and the unacceptable? Social workers, and the legal authorities behind them, come in for severe criticism for removing children from their parents’ care unnecessarily and for not removing children from their parents’ care when it would have been better had they done so. Vagueness arises in two dimensions:

1. “Sorites” comes from the Greek word for heap, “soros.”

How much evidence do you need? And even if you had perfect information, at what point is interference justified?

We all face problems of this structure. If the pressure is too high, or the temperature is too high, the machine is in a dangerous state and should be switched off. How high is too high? Exact values will be arbitrary: A system does not normally switch from safe to dangerous as a result of a tiny change in the relevant variables. One can be criticized for being too cautious and for being too rash, but it's often not easy to see how to get it right. There do not appear to be algorithmic answers. There appears to be something we call "good judgment" which we all have to some extent and which some have acquired to a higher degree in certain areas of expertise, such as medicine or politics or engineering.

We can be led astray by sequences of the sort that make up a sorites. Hans Kamp did some experiments, I was told, with color patches, running from red to orange in imperceptibly small stages, seen in order, one at a time. Participants just had to say "red" or "orange." People did eventually switch. But there were patches which, if approached from the red end, got the verdict "red," and if approached from the orange end got the verdict "orange." We are more readily led further astray when our desires are involved: One more chocolate won't make any significant difference to my weight; this unwelcome task has to be done, but I can always convince myself that it is inconsequential whether it is done today or tomorrow. We are being *seduced* when we are led astray by sorites sequences, whether the seducer is ourselves or another. One technique for corrupting a politician (I was told by an expert on the Mafia) is to begin with an innocuous gift such as a box of chocolates on his birthday and move in very small steps from there, until you have him in the palm of your hand.

The phenomenon of vagueness can make life difficult, then. But I don't think it strikes ordinary folks as paradoxical. If action is required at some point on the sliding scale, though there is no unique point at which it is required, so be it. When thinking clearly, we realize we can be misled if we give undue weight to the virtual identity between a given state and its neighbors, at the expense of the wider picture. We competently perform countless tasks based on judgments involving vague predicates. "Pick the raspberries when they are red enough." "Walk a little more slowly." "There's not enough coffee for a full pot. Fill it about three-quarters the way up." Despite the lack of sharp cutoff points for what the situation demands, we usually have no difficulty in complying.

There ought to be something positive and constructive to say about reasoning with vague terms. We learn many concepts, both in everyday life and in particular areas of expertise, by familiarity with paradigms—"clear textbook cases" for a given treatment or verdict, say. We go on to face situations relevantly different, to some extent, from these clear textbook cases. The judgment that this is at least close to a clear case for a given treatment or verdict should be of some value to us. Our question, then, is

how to reason in vague languages—how to avoid the pitfall of the sorites, and yet put vague judgments to some legitimate use.

II.

I turn to the problems vagueness causes for logicians and philosophers of language. Logic in the modern period, and the kind of philosophy of language that underpins it, was originated by Frege, adopted by Russell, the early Wittgenstein, and the logical positivists, and eventually became orthodox. Frege's and Russell's primary aim was to develop a logic adequate for reasoning about mathematics. Mathematics is *the* area of precision. We are not troubled by borderline cases of mathematical concepts.

From this perspective, it is one of the defects of natural language that for many predicates there are objects to which they neither clearly apply nor clearly fail to apply but rather hover on the borderline of applicability. In the type of ideal language which Frege constructed, whose semantic structure is crystal clear and in terms of which the notion of a valid argument is made perspicuous, this does not happen. The extension of a predicate—the set of things to which it applies—is perfectly precise. Hence the truth or falsity of a simple sentence is always a precise matter. This is still our standard idealized model of how logic and language work. The idealization is fine when we are not troubled by borderline cases, but it doesn't help us to understand how vague language works—what we can and cannot do with it.

Should we really try to reform language so that it conforms to the ideal of precision? How would we go about it? We could try to stipulate sharp boundaries between our predicates—between red and orange, and so on. But this conflicts with another ideal: A language should not make arbitrary, pointless distinctions. The difference between truth and falsity should be a difference that *matters*. Yet, on this proposal, there can be objects that are, let us say, qualitatively identical except in color, and so alike in color for it to be a matter of indifference which we choose for any purpose; and yet of one it is true that it is red, of the other it is false that it is red. The difference between truth and falsity is deprived of its point—the line serves no purpose beyond a fetish for precision.

Admittedly, for some specific, localized purposes, an arbitrary line is better than no clear line. There's a competition for the fastest time climbing two mountains over a certain height in a given area, starting and ending on a public road. In searching for the optimum peaks to tackle, you come across some unclear cases: Two summits, but so close together that perhaps this will count as only one mountain with a small dip near the top. The judges had better decide, however arbitrarily, what is to count, for the purpose of the competition, as climbing two mountains.

Similarly for the law, there is sometimes arguably a case for stipulating sharp boundaries. But there is a cost, even if it is sometimes the lesser evil.

There is a feeling of arbitrariness and injustice about the resultant fact that almost indistinguishable cases may be treated in a radically different manner.

Would an ideal language simply not have these vague words—not speak of people being tall or short but instead speak of them being six foot two inches in height? While the precise vocabulary is indeed better for some purposes, the converse is also true. It takes more time, effort, equipment, and opportunity to judge that someone is six foot two inches than it does to judge that they are tall. The latter judgment is easier to remember and to communicate. And often we simply don't have the corresponding precise vocabulary. What would we substitute for the judgment that someone is angry, or left-wing?

When the more precise vocabulary is available, vagueness can be reduced, but strictly speaking, it is not eliminated. "Tall" is vague. "Six foot two inches" is considerably less vague but, understood in a way in which we are capable of applying it, it is understood as applying within a certain degree of approximation. And there's no sharp boundary to the applicability of "approximately six foot two inches." It's doubtful whether there is such a thing as being *exactly* six foot tall, even at an instant. What about loose hairs, loose skin, the exact orientation of the spine? And at the microscopic level, there are molecules detaching themselves: Are they still part of the body or just close by? It's hard to think of any empirical predicate which we are capable of applying, for which it is not at least conceivable that it admits of borderline cases. (Though of course, what we think of as precise terms are precise enough, for practical purposes, for their potential vagueness to be ignored.)

III.

Vagueness is here to stay. It is not a defect of natural languages. It is extremely doubtful whether it is possible to eliminate it, and even if we could, we would be impoverished without it. We need a theory of reasoning in vague languages—one that explains what goes wrong in the sorites and tells us what use we can make of information which may not be clearly true.

The classical model is unhelpful when vagueness is a problem—when there are statements which are neither clearly true nor clearly false. Should we drop the "clearly" and claim that as well as true statements and false statements, there are borderline statements which are neither true nor false? Although there is some temptation to do so, I think it should be resisted. Consider an island divided into two countries, North and South, the border between which is not perfectly precise: There are spots such that it is not clear whether they are in North or in South. It doesn't follow that they are not in North and not in South but somewhere else instead. On the contrary, everyone can agree that there are just two possibilities, but it is unclear which obtains. Similarly, that statements hover on the borderline

between true and false does not imply that they are not true and not false but something else instead. It would not be definitely wrong to call a borderline statement true nor definitely wrong to call it false. That's quite different from saying it's not true and not false.

In any case, a threefold classification—true, false, neither—or, as I prefer, clearly true, clearly false, neither clearly true nor clearly false—is not much better than the classical twofold classification for explaining how vague language works. For we are as much at a loss about where to draw the lines between these three categories as we were with two.

In my view, a more fine-grained model is needed to understand reasoning with vague terms. Go back to the color spectrum. There are clear cases of red; and clear cases of orange; and things of varying degrees of closeness to these two poles. This is the essence of vagueness. It is this fact that generates the sorites sequences and the lack of sharp cutoff points.

A useful idealized model of the phenomena represents the degree to which a statement is close to clearly true by a number between 1, for clearly true, and 0, for clearly false; the closer to 1, the closer to clearly true, with 0.5 when evenly balanced between these two poles. The idealization is of instrumental value for it gives us access to arithmetical operations in exhibiting the logical structure of statements that take these degrees. To insist that it is an idealization is to admit that there are no uniquely correct numbers to assign. It is a precise model of an imprecise phenomenon. Worthwhile results obtained from this framework must be robust enough to be independent of small numerical differences, which may be without real significance. (We do the same, to good effect, with degrees of closeness to epistemic certainty, degrees of preference, and many physical magnitudes.)

I grant that there is no exact point at which clear truth ends and something very close to it begins. But in the use we make of the structure, this doesn't matter: The difference between clear truth and almost clear truth—between 1 and 0.99—is an insignificant difference upon which, normally, nothing hangs.

I grant that comparisons cannot be made arbitrarily finely. Which is closer to clearly red, this orangeish patch or this purplish patch? Comparisons are more difficult with "intelligent," "wicked," "left-wing." We must make do with judgments of rough equality when no clear comparative judgments can be made.

What sort of structure do these degrees have? What sort of logic do they generate? We can say that the degree to which A is close to clearly true— $v(A)$, for short—is 1 minus the degree to which its negation, not A , is close to clearly true: If " x is red" gets 0.9, " x is not red" gets 0.1. Beyond that, there is disagreement amongst degree theories. Consider conjunctions. The most advertised rule says the value to be assigned to " A and B " is the minimum of $v(A)$ and $v(B)$. I think this is wrong. Let $v(y$ is red) be 0.5. And let x be a little less red than y —let $v(x$ is red) be 0.4. According to this rule, $v(x$ is red and y is red) is 0.4. That seems right. Now note that $v(y$ is not red) is 0.5. So

by this rule $v(x \text{ is red and } y \text{ is not red})$ is 0.4. But as y is redder than x , it is surely completely wrong to judge that x is red and y is not. This last conjunction should get 0.

In the above example, $v(y \text{ is red}) = v(y \text{ is not red})$, but $v(x \text{ is red and } y \text{ is red})$ does not equal $v(x \text{ is red and } y \text{ is not red})$. Thus we need a structure that does not require that the value to be assigned to a conjunction is fixed by the values to be assigned to its parts. There is one well-known degree-theoretic structure applied to propositions which has this character, and it is at home in application to matters of indeterminacy. It is probabilistic structure. The probabilities of heads, of tails, of not tails, may all be 0.5. But the probability of (heads and tails) is zero, and the probability of (heads and not tails) is 0.5. This is accommodated in probability theory by the rule that the probability of $(A \text{ and } B) = \text{the probability of } A \text{ multiplied by the probability of } B \text{ given } A$ (i.e., on the assumption that A). The probability of tails given heads is 0. The probability of not tails given heads is 1. Analogously for degrees of closeness to clear truth: $v(y \text{ is red given the hypothetical decision that } x \text{ is red})$ is 1; $v(y \text{ is not red given the hypothetical decision that } x \text{ is red})$ is 0. Hence the conjunctions differ.

Giving the degrees probabilistic structure gives us the analogue of a nice result. Take any classically valid argument. It is demonstrable that the improbability (one minus the probability) of the conclusion cannot exceed the sum of the improbabilities of the premises. First, consider plain epistemic uncertainty; there are good reasons for holding that, idealized as precise, "degree of closeness to certainty" has probabilistic structure. The above result (1) vindicates accepting the conclusions of valid arguments of whose premises you are nearly certain, provided that there are not too many such premises; if you have two premises, each at least 99% certain, your conclusion is at least 98% certain. And (2) it explains the pathological cases, like the lottery paradox: "Ticket 1 won't win; ticket 2 won't win; . . . ticket 100 won't win; so none of the tickets will win." We have a valid argument, a large number of premises, each of which is very close to certain, yet the conclusion is certainly false. For the small amounts of uncertainty in each premise can add up to a maximal amount of uncertainty.

Similarly with vagueness. You *can* rely on classically valid arguments whose premises are at least close to clearly true, provided that there are not too many such premises. For instance: "Jones is a case for treatment X. If Jones is a case for treatment X, so is Smith. So Smith is a case for treatment X." If each premise gets at least 0.99, the conclusion gets at least 0.98. But iterating the argument many times can take you from a clearly true premise to a clearly false conclusion. A large number of insignificant departures from clear truth can lead, by valid reasoning, to a maximal departure from clear truth.

In claiming that degrees of closeness to clear truth have probabilistic structure and stressing a structural analogy between vagueness and epistemic uncertainty, I am not siding with those who say that vagueness *is* epistemic uncertainty: I am not saying that there is a sharp line dividing the

red from the orange and we merely don't know where it is. Mere epistemic uncertainty is not the only application of the calculus of probability. It also applies to cases of physical indeterminism. Where there is vagueness, there is indeterminacy in the relation between our words or concepts and the things to which they apply. And the best theoretical framework in which to understand that indeterminacy, in my view, has the formal structure of probability.

The above account can be related to another popular approach to vagueness known as "supervaluationism." This uses the idea of a permissible sharpening of a vague predicate: a sharp predicate that agrees with the vague predicate on the clear cases. A statement is then clearly true if it is true on all permissible sharpenings, clearly false if it is false on all permissible sharpenings, and neither clearly true nor clearly false if it is true on some but not all permissible sharpenings. This yields a threefold classification, and we have lost the valuable idea of statements being closer or further from clearly true. But, as some authors have pointed out, this feature can be reintroduced. A statement is close to clearly true if it is true on *almost all* sharpenings. For simplicity, assume that there are only finitely many sharpenings. Then we can think of the degree to which a statement is close to clearly true as the proportion of sharpenings on which it is true. Proportions satisfy probabilistic structure. We get the same results.

I think of this as a useful heuristic device. It can be easier to think through a complex case, in terms of what is true on all, almost all, about 90 percent, about half, and so on, of sharpenings. I do not accept supervaluationism as an account of the fundamental nature of vagueness. On the supervaluationist picture, all features of reality—all objects and properties—are sharp and precise. Vagueness arises because of semantic indecision: We simply haven't made up our minds to which, amongst various candidate sharply bounded objects and properties, our words and thoughts refer. There is no property of redness, or being old, but a battery of related sharp properties between which we show indecision in our use of "red" and "old." There is no unique mountain, Mount Everest, with fuzzy edges; no unique man, Fred Bloggs, whose temporal bounds are a little imprecise. There is, instead, a battery of overlapping sharply edged mountains, and we have left undecided to which we refer by "Mount Everest"; similarly for "Fred Bloggs." This is a strange ontology. And it represents vagueness as a very superficial phenomenon. It does not do justice to the world of experience in which objects really do have slightly imprecise edges, and the properties of objects worthy of our attention do allow for indeterminacy in their application.

Finally, let me comment on the boom industry, fuzzy logic. Washing machines, car braking systems, refrigeration units, and so forth all have improved controls if programmed with fuzzy logic. The term covers a variety of techniques for replacing on-off devices by "more-or-less" devices. A thermostat is more efficient if, instead of switching on and off, it merely adjusts the output to higher or lower.

If the term “fuzzy logic” is used with sufficient generality to cover any degree-theoretic treatment of vagueness, I am in favor of fuzzy logic. On the other hand, what are often put forward as the basic principles of the discipline are rules which I think are demonstrably unsound, like the “minimum” rule for conjunctions discussed above (p. 373). The unsound rules are simpler, though, and may well be good enough in limited areas of application—when designing washing machines, say. Probabilistic structure is computationally much more complex. Engineers are pragmatic folk—they will use whatever formula they can come up with that gives the greatest increase in efficiency per unit cost. The trade-off depends on the nature of the project and how high the stakes are. As I recall reading in a review of a popular book on fuzzy logic some years ago: If you’re sending a man to the moon and you need to design degree-sensitive systems, you had better use probability theory.

I have argued that we need more more-or-less thinking, and less on-off thinking, when we are reasoning in vague languages. This doesn’t always help. Sometimes, as in some of my early examples, an “on-off” sort of action is called for at some point on a sliding scale, but there is no unique point at which it is required. We possess the skill to handle such situations, unreflectively, most of the time. But—as with many of our skills—we don’t have a theory of how we do it.

REFERENCES

- The best collection of articles and the best full-length book on the philosophical problem of vagueness are as follows.
- Keefe, Rosanna and Peter Smith, eds. (1996) *Vagueness: A Reader* (Cambridge, MA: MIT Press).
- Williamson, Timothy (1994) *Vagueness* (London: Routledge).