The Problems of Vagueness

Mat Carmody

Introduction

It has just turned seven o'clock on a cold and wet November morning. You know it has just turned seven o'clock because your alarm clock has sounded. You have to get up to get ready for college because you can't pretend to have fallen ill so soon after your last bout of sickness. You find yourself thinking that, what with it being only seven o'clock, there's plenty of time and that it can't hurt to spend just another couple of minutes in bed. How can two paltry minutes matter? Who could deprive anyone of little more than a hundred seconds of extra warmth and comfort on a morning like this?

I'd be prepared to bet that you've reasoned in this way at least once or twice in your life. We are all very willing to tell ourselves that an extra minute won't matter, that an extra few pence is of no importance and that an extra few chips will make no difference to our waistlines. We know that people have been thinking in this way for a very long time. The biblical patriarch Abraham tries to use the insignificance of small differences when arguing with nothing less than God. Yet it is with an ancient Greek philosopher called Eubulides that many people find the start of a deep philosophical puzzle that such reasoning can produce. Eubulides was famous for having seven riddles, which sound like brain-teasers, but which when analysed prove to be of more than coffee-break interest. Two of these are known as the paradox of the heap and the paradox of the bald man. We shall take a look at the first and return to the second in a moment.

The Paradox of the Heap

I hope you will agree that if I arrange them in the right way, a million and one grains of sand can make up a heap of grains of sand.[1] I hope you will also agree that a single grain of sand does not make up a heap. I put a heap-shaped one million and one grain before you and a single grain behind you. Your task is to remove a grain at random and throw it onto the sand behind you. You are to repeat the task one million times until you have a single grain before you.

It's tempting to think that the addition or subtraction of a single grain can't make a difference to whether something is a heap or not in the same way that the addition of just two minutes to your morning lie-in doesn't really matter. Heaps and lie-ins, we may say, *tolerate* small alterations. We shall call this tempting thought the *tolerance thought.*

Let's now present matters formally. An argument is a collection of sentences we call the premises alongside a sentence we call the conclusion. The premises we assume are true. An argument is a good or valid argument if there is a logical path from the premises to the conclusion. Let's put this thought and the two we opened with as the premises of an argument.

(1) 1,000,001 grains make up a heap

(2) 1 grain does not make up a heap

(3) The removal of 1 grain does not turn a heap into something that is not a heap (or vice versa).

What conclusion can we draw? Each time you remove a grain from the heap, you preserve the heap because of (3). Each time you add a grain to the 'non-heap' collection behind you, you still have a non-heap, also because of (3). Even after a million subtractions to the heap before you therefore, when just one grain remains, you still have a heap:

(4) 1 grain makes up a heap.

Similarly, after a million additions to the non-heap behind you, you have a 1,000,001 non-heap:

(5) 1,000,001 grains don't make up a heap.

We can therefore deduce, by combining (2) with (4) and (1) with (5):

(6) 1 grain makes up a heap and 1 grain doesn't make up a heap.

(7) 1,000,001 grains make up a heap and 1,000,001 grains don't make up a heap.

Two sentences present a contradiction if they both cannot be true together and they both cannot be false together. For example, the sentences 'John is at home' and 'John is not at home' present a contradiction, if we understand them literally as talking about the same John. A paradox is an argument whose conclusion presents a contradiction. (6) and (7) both present contradictions. Our argument is a paradox: the paradox of the heap. You should be able to see that the same argument can prove that 2 grains do and do not make up a heap, that 3 grains do and do not...that any number indeed do and do not. This is clearly unacceptable.

Our conclusion is the product of the premises and the logical reasoning we used to move forward from them. If we can't accept the conclusion of an argument, it is with either or both of these factors that we must find fault.

Reasoning

It often happens that we are tripped up by an argument that looks and feels logically valid but which turns out to be more complex than we think and conceals logical blunders. Yet this argument really just uses one logical principle.

The first principle we rely on is called modus ponendo ponens or (usually) just modus ponens for short. The principle says:

(MP) From [A] and [if A then B] deduce [B]

For example: from [it is raining] and [if it is raining, then Bernard will be at home] we can deduce [Bernard will be at home].

Modus ponens is the cornerstone of reasoning. It captures the idea of conditional judgement or reasoning or action, without which neither we, nor any other creature capable of the most basic of thoughts, could do anything. It therefore seems we can't reject it here.

Modus ponens allows us to deduce [1,000,000 grains of sand make up a heap] from [If 1,000,001 grains make up a heap then 1,000,000 grains make up a heap] and [1,000,001 grains make up a heap]. The conditional there is just an instance of our tolerance thought. By repeating the reasoning, we reach (4): 1 grain makes up a heap.

(2) and (4) present a contradiction. Formally, we can join them together to come up with the sentence: 1 grain of sand doesn't make up a heap and 1 grain of sand makes up a heap. In so doing, we use another logical principle, called *conjunction introduction.* A conjunction is just a sentence of the form [A and B]. The principle says:

> (Conj I) From [A] and [B], deduce [A and B]

This should seem too obvious to merit either denial or comment. We shall therefore conclude that the reasoning is not at fault.

Nihilism

It must therefore be one or more of the premises. When we are faced with premises all of which seem true, we need to think of what it would mean if they each weren't true and which of these situations would be the least unacceptable. If (1) were not true, then 1,000,001 grains would not make up a heap. I chose this number thinking that that many grains could be formed into a heap. I hope you'll think that, if I am wrong, then some large number of grains could be made into a heap. If we deny (1), therefore, we are in effect denying that any number of grains could make up a heap - or, more simply, that there are heaps. Similarly, if (2) were not true, we'd be in effect denying that any number of grains could fail to make up a heap.

Either denial seems impossible. It is however as close to an axiom in philosophy as anything is that there is no position on a problem, however crazy, that no-one has not chosen to adopt. It so happens that a couple of philosophers have defended the thought that (1) is false. The idea is that the paradox shows our vague words to be so defective that they don't truly apply to anything. Nor could they apply to any possible thing. It is rather as if 'heap' simply isn't fit to be applied to things in the same way that 'red' isn't fit to be applied numbers. In saying that 'three isn't red', we aren't saying that it is some other colour. Numbers don't have colours. Similarly, collections of material objects can't be made into heaps.

Most people regard this view as both nonsensical and desperate. Even if we allowed that nothing really is a heap, then plenty of things *look like heaps or are heaps according to how we use that word.* You should be able to see that these expressions are just as vague as 'heap' is. It therefore follows that nothing even looks like a heap.

Now we may be wrong about whether someone is tall because appearances are misleading. Yet whether someone looks tall just depends on appearances - indeed, how things appear to us individually. Generalising, if we allow that 'looks X' doesn't apply to anything, where X is replaced with any vague expression, it seems that we can't even describe how things seem. Yet we can surely do this at least. Even Descartes allowed that we couldn't be fooled about our most basic appearances.

Almost all philosophers think that the problem lies with the third premise. In order to explain why, it'll be worth looking at the puzzle from a different angle

Vagueness, Boundaries and Borderline Cases

It is often said that not everything is black and white. An action may not be clearly good or clearly bad but be good in some respects and bad in others. An accusation may neither be wholly true nor wholly false. An expression may neither clearly apply nor clearly fail to apply. For example, you can probably think of someone who you wouldn't call tall but of whom you couldn't really say that he isn't tall. He's somewhere in the middle. He's what we call a *borderline* case of tall.

If you took someone who was tall and gradually shrank them, then you'd find yourself with a borderline case of tall before finding yourself sometime later with someone not tall. Similarly, if you took something that was a heap and removed grains one at a time, you'd be faced with borderline cases of heaps before non-heaps. *Heap* and *tall* are alike in being *vague.* They allow for borderline cases. They invite the tolerance thought. Although we haven't formally spelled out the paradox for *tall*, it is not hard to do. Do you think a difference of a millimetre matters to whether someone is tall or not? If not, then by shrinking a person two metres in height by a millimetre at a time, you can end up with a person one metre tall that is still tall.

An expression such as *eighteen years of age or more* is one that we can regard as *sharp*. You are either eighteen years of age or more or you are not. The age of eighteen is a sharp boundary that separates these two possibilities. The thought is that *heap* and *tall* are vague just because there is no such sharp boundary. If there were a sharp boundary for *tall*, then there would be some exact height above which people were tall. If there were, then someone a millimetre below this height and therefore not tall would become tall by growing a millimetre. This contradicts the tolerance thought. The fact that we are tempted by tolerance shows that we don't like this idea of a sharp boundary. Indeed, we think that there are instead borderline cases, which replace this sharp line with a middle region.

Let us now return to the third premise. It says that a small change of a grain can't turn a heap into a non-heap. If this were false, wouldn't it mean that a small change *could* turn a heap into a non-heap? It seems not. This would only be so if we think that there were two possibilities: heap and non-heap. No small change can turn a heap into a non-heap just because there are lots of borderline heaps in between. We can therefore say that the third premise isn't true without this committing us to a heap/non-heap sharp boundary.

What Vagueness Is and Isn't

Any expression that allows for borderline cases and the tolerance thought is vague. We've seen that *heap* and *tall* are vague. Eubulides invited also to think that by pulling one hair out of the head of a non-bald man, you couldn't make him bald. You should be able to construct a paradox using this tolerance thought along the above lines. Such paradoxes are called *sorites paradoxes* after the greek word for 'heap'.[2]

The ingredients needed to generate these puzzles are simple. You need to think of a dimension of change, such as numbers of grains or numbers of hairs. At one end of this dimension, we find things that are X. At the other, we find things that are NOT-X. All we add now is the tolerance thought: that small differences don't matter.

Here are some examples to get you thinking. Heights, weights, distances, times, temperatures are all dimensions of change. If a person is two metres in height they are tall and if one metre they are not tall. Yet can growing by a millimetre make someone tall who was not-tall? You should be able to generate similar puzzles for *short, heavy, light, near, far, hot, warm and cold*. Now think of any expression for which number is relevant. Alongside heap and *bald* we have *crowd* (number of people), *fleet* (number of ships) and *many, few, lots, some, plenty* (number of whatever is counted). Time is also a dimension of change.

It seems odd to say that the passing of a second can make something *old* that wasn't before, along with making a creature an *adult* of its species if it wasn't before. Finally, take any everyday material object, such as a *coffee cup*. It is made up of a large number of molecules. If I remove a very small number at random, I surely can't destroy the object by turning it into something no longer a *coffee cup*. The same is therefore true of any of the things that populate our environment - including ourselves. If *you* can tolerate the loss of a few molecules, then you do and don't exist!

Philosophers distinguish different types of vagueness. One type with which sorites vagueness is often confused is what I shall call 'comparison-class vagueness'. Suppose I tell you something has a height of ten metres and I ask you whether it is tall or not. You can't answer until you know what sort of thing it is. A ten-metre tall man is very tall. A ten-metre building is not tall. A ten-metre tree may be tall, depending on the species. It can therefore be vague whether something is tall just because it may be vague what sort of thing we're using as the basis for making our comparative judgements. This sort of vagueness is not at all the same thing as sorites vagueness. Once we have agreed that the thing is a beech tree and that it is indeed tall, we still generate a paradox for 'tall beech tree' via the tolerance thought that if a tall beech tree shrinks by a millimetre, it is still tall.

A third and similar form of vagueness is informational vagueness. You want to meet Anna for lunch. You ask me where she is. I tell you that she is in London. This information is very likely to be too vague to be of much use. What you wanted was her precise location. On the other hand, if you know that Anna only ever has lunch in one restaurant when she is in London, then my telling you that is in London rather than Paris will suffice. Information is vague or precise dependent on the use we intend to put it to.

Three Truth-Values

The sentence ‘Jupiter is a bigger planet than Mars’ is true. Another, more formal way of putting it, is that the sentence has the truth-value *true*. In the same way, we can say that the sentence ‘Mars is a bigger planet than Saturn’ has the truth-value *false*.

So far, we have spoken as if these are the only two values a sentence can have. We believe that it is true that 1,000,001 grains make a heap. What about 1,000,000 grains? Either the sentence ‘1,000,000 grains make a heap’ is true or false. If false, then a sharp boundary divides these two collections into *heap* and *non-heap*. Our tolerance thought says that the sentence is true, just because one grain can’t make a difference. But this generates a paradoxical conclusion.

We recently found a way through the dilemma: we admit borderline cases. At the level of truth, this translates into adding a third truth-value: *borderline* or *indefinite*. It may then be, for example, that ‘34,346 grains make a heap’ is *indefinite* just because they make up a borderline heap.

If we turn back to the formal argument, we are saying that (1) and (2) are true and that (3) isn’t. Should we say that (3) is indefinite or false?

> (3) The removal of 1 grain does not turn a heap into something that is not a heap (or vice versa).

One generalised and clear way of re-expressing (3) is as follows:

> (3’) If a collection of x grains makes up a heap, then a collection of x-1 grains makes up a heap.

We might think as follows. (3)’ says something untrue, that is, false, namely that ‘heap’ continues to apply, no matter how few grains remain. This doesn’t mean that we have a sharp borderline. We interpose the borderline cases.

Unfortunately, we can’t say that - at least, not so quickly. It is not possible to explain here in detail why. You will remember that earlier we looked at two principles (MP) and (Conj I). These seem to capture some of the meaning of our everyday logical words ‘if’ and ‘and’. It happens that, if we say that (3) or (3’) must be false, we end up forced to challenge the logical principles we hold for these words and others like ‘or’, ‘not’, ‘all’ and ‘some’. It turns out that making (3) or (3’) false is too high a price to pay.

So we must hold that (3) and (3’) are indefinite. What does this mean? One thought would be: it’s indefinite because it depends on the number of grains. But this still implies that, for at least one x, if a collection of x grains makes up a heap, then a collection of x-1 grains makes up a heap, which we don’t want. A second thought would be: it’s indefinite because we can’t tell. But this makes the problem one of knowledge: it may be that the removal of a grain can turn a heap into a nonheap but we can’t be aware of this. This isn’t any more tempting.

We can in fact defend the thought that it is indefinite from a logical point of view and that this doesn’t translate into anything ‘intuitively’ obvious in everyday English. Yet you might think that (3) and (3') are better described as not indefinite, not true nor false, but nearly true. Can we make sense of this?

An Infinity Of Truth-Values

Some philosophers think that there are far more than three truth-values. Alan and Bill are 1.67m and 1.68m tall respectively. They are both borderline tall. It is indefinite whether they are tall or not. Yet Bill is taller than Alan by a small amount and so it should be possible to say that it is truer that Bill is tall than Alan is tall.

We can distinguish many degrees of tallness, of baldness, of redness, and so on. Indeed, every vague expression X generates a meaningful expression '...is more/less X than...'. We should therefore have as many degrees of truth as we can have degrees to which these expressions apply. Between the numbers 1.67 and 1.68 are infinitely many numbers. We should therefore admit infinitely many truth-values. Even if there aren't infinitely many degrees of being a heap, it will be better to have too many truth-values which we can group together than to have too few.

As far as the paradox is concerned, a very neat reply becomes possible. I shall leave out the technical details and try to convey the essence. It is true that 1,000,001 grains make a heap. It is slightly less true that 1,000,000 grains make a heap. It is slightly less true again that 999,999 grains make a heap. It is false that 1 grain makes a heap. Our premise (3) tells us that it is true that if x has n grains and is a heap, then x' with n-1 is a heap. We can say that this is nearly true but not quite true. What is quite true is that if if x has n grains and is a heap, then x' with n-1 is a heap to a slightly lower degree. We confuse these two thoughts. If we admit that (3) is almost but not quite true, it becomes possible once again to defuse the paradox. This is better than saying it is simply 'indefinite'.

Epistemicism

Those who defend alternative logics consider vagueness to be a semantic problem. They think that words like *heap* and *tall* have meanings that are incomplete. *Heap* is defined in such a way that some things fall under the definition - the clear cases of heaps. The definition equally excludes some things - the clear cases of non-heaps. The borderline status of a borderline case reflects a real and profound absence of truth or falsity stemming from the definition not covering such cases.

An alternative and radical proposal is that vagueness does not arise from incomplete meanings but from our knowledge of meaning. It is not semantic (to do with meaning) but epistemic (to do with knowledge). The Epistemicist tells us that *every* vague expression is semantically precise. In other words, there is a sharp boundary dividing heaps from non-heaps, tall people from non-tall people, and so on. We simply don't know where this boundary is. Furthermore, we can't know.

In order to explain why, let us examine the concept of knowledge. We say things like 'Bernard knows that we are meeting this evening at 8pm'. We can break this down into three parts: a subject (Bernard), a *propositional attitude* (knowing that) and a proposition (we are meeting this evening). For our purposes, we can think of a proposition as like a sentence about which it makes sense to ask: is it true or false? By way of comparison, if I say 'Bernard knows how to drive a car', we have an attitude - knowing how - but not a propositional one, as it does not make sense to ask whether 'to drive a car' is true or false.

Our 'attitude' to the proposition 'we are meeting this evening' may be one of knowledge but also one of belief, desire, fear or hope, amongst others. An attitude is something like a way a mind relates to information. Alongside knowledge, we have the closely-related attitude of belief. The essential differences between belief and knowledge are as follows. Firstly, you can believe something false but you can't know something false. It is possible to say 'Until I was eight, I believed New York was the capital of America' but not 'Until I was eight, I knew New York was the capital of America'. Secondly, knowledge is more 'robust' or reliable than a true belief. I ask you whether Bernard was at home yesterday and you reply, 'I believe he was'. If I ask, 'are you sure?' it would be fair for you to say, 'no'. You have some grounds to believe he was at home but you may not be supremely confident. If you reply, 'I know he was', then you advertise your confidence. When we want to find things out, we tend to search for people who know, rather than simply believe.

Consider now the following situation. I am at a football stadium in which there are 74,362 people, including me. I guess how many people there are and I guess that there are 74,362. Now you will agree that, even though I am right, I don't know that there are 74,362. I am right by chance. I haven't done anything, such as count the people, to make my lucky guess into a robust piece of knowledge. Indeed, had there been a few hundred more people, I could have just as easily guessed, wrongly, that there were 74,362 people because it would have looked no different to me. I may not be able to have exact knowledge of the number of people in the stadium but I can have inexact knowledge. I can know that there are at least 2 and not more than 1,000,000 people in the stadium just by looking because these two situations look sufficiently different from the actual one for someone with my capacities for discrimination. If I had a better capacity to judge crowd sizes, I would have less inexact knowledge. I might know, for example, that there are between 60,000 and 80,000 people.

The phenomenon of inexact knowledge is very general. You can know on the basis of appearances someone's height or age or weight or distance from you roughly, which is to say that you can say what heights, ages, weights and distances are clearly wrong and thus define a range of possible answers in the middle. The better a judge of stadium capacities, heights, ages, and so on, the smaller the middle range wherein lies the right but inaccessible answer.

The Epistemicist likens the exact number of people in the crowd to the exact number of grains he says is the sharp boundary between a heap and a non-heap. Let us call that exact number n. As a speaker of English, I understand the word 'heap' and, when presented with heaps, borderline heaps and non-heaps, will react appropriately. As with so many of our words, our understanding consists in being able to use 'heap' in a certain way rather than being able to voice a clear definition of heap. (Look around you and find a word to classify each object you can see. Can you define any of these words precisely? Do you feel this undermines your claim to be using these words correctly?)

The Epistemicist says that how we use ‘heap’ in fact determines n but that we don’t know how our collective use fixes n. It is not as if we can examine every possible situation each of us could be in to see how each of use would use ‘heap’. Since we can’t see this ‘total use’, we couldn’t discriminate it from the total use of a group of speakers of a language Schminglish who used ‘heap’ very much like us but not identically to us. Since it is use that determines the sharp boundary and the uses are very similar but distinct, ‘heap’ in Schminglish determines a sharp boundary of - let us say - n+1 grains.

By way of analogy, think of the speakers of English as a stadium with 74,362 people and the speakers of Schminglish as a stadium with 74,363 people. In the same way that we can’t distinguish their total number in either case, making them appear indiscriminable, we can’t distinguish how English and Schminglish speakers use ‘heap’. So since we can’t distinguish our use precisely, we can’t distinguish the number n precisely. We can only have inexact knowledge. Just as I can know that the stadium has more than 2 people and fewer than 1,000,000, I can know that 2 grains can’t make a heap and 1,000,000 can. Just as I can’t know that the stadium contains 74,362 people, I can’t know whether 74,362 grains make up a heap. It is a borderline case not because there is no right answer, as the semantic approaches claim, but because I can’t know the right answer.

So Who’s Right? Logic vs. The Incredulous Stare

The semantic approaches start from the simple observation that the sorites paradox supposes that every statement is true or false. This supposition is called the ‘principle of bivalence’. The common recommendation is then to reject this principle and allow for more truth-values. The common problem is that the logics that result don’t seem to work properly.

We saw earlier two logical rules (MP) and (Conj I). A ‘logic’ is a collection of such rules which tells you what you can and cannot deduce from your premises. In order to design a set of rules, you have to decide how many truth-values you are going to allow. For you are in effecting asking yourself: if A is true/indefinite/etc. and B is true/indefinite/etc., then can I deduce C? Classical logic is a system of rules with the principle that there are two truth-values, **true** and **false**. Non-classical logics are those that have different sets of rules and/or more truth-values.

Classical logic is highly regarded because its rules capture what we feel to be patterns of proper reasoning. It is therefore claimed by many to be logic that not merely our minds, but any intelligent mind, would employ.[3] It is therefore claimed that, because of this, we should accept that every sentence is indeed either true or false. This claim would be undermined, of course, if non-classical logics outperformed classical logic when it came to representing reasoning. The problem is that each non-classical logic comes in for heavy criticism on this score. It is alleged that they each have rules that permit deductions that are unacceptable. It is not possible to go into details here. I shall just say that one of the key planks in the defence of the Epistemic position is a demonstration of just how hopeless the opposing positions are on the matter of logic.

The Epistemic position needs to build a solid defence of itself because its central claim is *prima facie* so outrageous as to lead one to think that *any* other position is preferable. Its central claim is that every vague expression does in fact have a sharp boundary. Yet there is something disturbing about the thought that something we do determines a sharp boundary that we can't find. We should be happy to accept that there are scientific and mathematical facts that may escape our knowledge because we lack the equipment to measure them and the minds to understand them. But can it be true that there are facts about who is tall and who is bald that are beyond us? Each position has its advantages and drawbacks and sorites paradox remains unsolved. Let's finish by looking at two issues that follow from the initial puzzle and illustrate the breadth of the issues vagueness raises.

Higher-Order Vagueness

The Epistemicist asks us to believe in simplicity. There is a sharp but hidden boundary dividing the Xs from the non-Xs for all vague expressions X. The Semanticist rejects this. But is he merely shifting the problem?

One semantic approach said that there were three truth-values. Things are divided into X, borderline case of X and not-X. If so, does this not mean there is not one but *two* sharp divisions? Consider a line of people who differ gradually in height. The tallest is 2m tall, the smallest 1m50cm and each adjacent pair differ by a half-centimetre. We don't believe that there is a sharp boundary dividing them into tall and not-tall. Is it any more obvious where the boundary between the tall and the borderline tall is?

Our second semantic approach said that there were lots of truth-values. Things are divided into many different degrees of X, from 'full' to 'zero'. If so, does this not mean that there are lots and lots of sharp divisions? Here, one might just say: yes. Every difference in height is reflected by a difference in degree of being tall. We therefore argue as follows. A man 2m10cm in height and a man 2m in height are both simply tall. *Tall* covers a range of heights. The same is true of not-tall. In between those heights tall and not-tall are lots of heights and lots of different degrees of being tall. Let's lump them all together as reflecting different ways of being borderline tall. In this way, we resurrect the problem of the previous paragraph. It seems that at some point, we must cross a height marking the end of tall and the beginning of borderline tall. But where?

It has seemed to many philosophers that the problem doesn't arise because of *higher-order vagueness.* So far, we have said it can be vague whether someone is tall or not tall and that such a someone is borderline tall. But couldn't someone be a borderline case of borderline tall? If so, then there is no boundary between tall and borderline tall. There is a range of heights corresponding to borderline borderline tall.

You may have already guessed what the next problem is going to be. If borderline borderline tall is sandwiched between tall and borderline tall, then we have *even more* sharp boundaries. Where does tall cross over into borderline borderline tall? Of course, if borderline borderline tall can itself have borderline cases, then we can squeeze this new category into where a borderline was supposed to be. But the problem simply reappears at the next level up.

Many philosophers think that we have to admit these higher levels or orders of vagueness just because it is absurd to suppose that there are sharp boundaries between tall, borderline tall and not-tall. Yet it is not obvious that it is a price worth paying. It seems that if we start ‘going up’ we must go up indefinitely. This means that tall makes infinitely many fine divisions. Can this simple little word like *heap* hide so much complexity?

A plainer line of attack is that it is entirely spurious, at least as a means of avoiding sharp boundaries. We reason just as we did a moment ago. There are some things that are, simply, tall and some that are, simply not-tall. In between, let there be as many divisions into degrees of tall and ever-higher orders of borderline cases of tall as you wish. Collect them altogether and label them as n = ‘not clear cases of tall and not-tall’. Once again, we seem to have arrived back with three categories: tall, n, and not-tall. Once again, we seem to have arrived back with two sharp divisions, marking the edge of ‘tall’ (with everything else) and ‘not-tall’ (with everything else).

We don’t appear to be any more knowledgeable about these boundaries. So where epistemicism asks us to believe in one hidden and inaccessible boundary, all semantic approaches ask us to believe in (at least) two hidden and inaccessible boundaries. It is this strange fact that Epistemicism asks us to believe that provides a lot of the drive to find an alternative. If semantic alternatives are in the same boat, however, then Epistemicism comes off in a better light. Or, you might conclude, no-one comes off in a good light at all. Vagueness seems to commit us to sharpness wherever we turn! There is therefore something very deeply wrong in how we understand the relation between minds, language and the world.

Onticism

Some say that vagueness is a problem with defective meanings. Others say that it is a problem of knowledge of those meanings. Others still say that vagueness is a feature of things.

This is a natural view when we are talking about things we can refer to with names, such as objects and places. Does a cloud occupy a precise space or is it a vague entity with fuzzy spatio-temporal boundaries? Does London have a sharp geographical boundary? Or is it vague whether some bits of land are part of London or not? Does London have a sharp chronological boundary? Did it come into being at some precise moment? Or was the transition from the first settlement to the final city a vague one?

Many philosophers think that reality itself is not vague. A man has a precise height but may be vaguely tall because of a problem with the word *tall*. Equally, there is a precise configuration of buildings and bits of land in the world but which parts fall under ‘London’ is vague through it being unclear what that word means as well.

One reason they think this is because of a famous little argument that appears to prove that it is incoherent to suggest that the world might really be vague. It goes as follows. Let us suppose it is vague whether London today ($London_{2004}$) is identical with London a thousand years ago ($London_{1004}$). It is surely not vague that $London_{2004}$ is identical with $London_{2004}$ - that’s trivial! So there is something $London_{2004}$ has that $London_{1004}$ lacks. This is the property of being non-vaguely identical to $London_{2004}$.

There is a widely-held view that if x and y are identical, then they have identical properties. It follows that if x and y do not have identical properties, they are not identical. We have just found a property that $London_{2004}$ has that $London_{1004}$ lacks. So they are not identical after all. This contradicts our assumption that it was vague whether they were identical. The idea that two things may be vaguely identical is therefore paradox-inducing and unacceptable.

If objects can be vague, then it will not always be clear whether one object is identical to another. So if x is a vague object, there will be an object y such that it is vague whether x = y. We have just proved that it is not possible to say that it is indefinite whether one object is identical to another. It follows that objects cannot be vague!

This might seem a bit fishy. Does this little piece of reasoning really prove that all clouds are precise entities? According to most philosophers, the answer is no. It is true that clouds are vague entities and it can be true to say that it is indefinite whether cloud x is identical to cloud y, just as common sense demands. What the argument shows is that we can't make sense of this by supposing the vagueness is somehow a feature of the clouds themselves. That leads to paradox. The right conclusion is that vagueness is a feature of how we refer to these things.

The semanticist says that when I use the word 'cloud' in 'that cloud is fluffy', I don't pick out a particular vague entity. I don't manage to pick out any one thing at all. It is vague what particular region of space I am referring to. If I wonder whether cloud x is identical to cloud y, this can be vague just because it is vague which two regions I am thinking about.

The Epistemicist says that I do pick out a precise region of space but I don't know which one. When I wonder whether cloud x is identical to cloud y, then either it is or it isn't. I may not be able to know the right answer. When this happens, it will be vague whether they are identical or not.

Conclusion

The vast majority of our words are vague. We are seduced by the thought that small differences don't matter when it comes to using them. Yet this leads to the unacceptable conclusion that they apply to everything and nothing. The problem of vagueness is to understand how vague expressions do have limits and how small differences can matter. A solution to this will be a solution to the sorites paradox. As with so many of the best philosophical problems, it is simple and fundamental and continues to vex philosophers over two millennia after it was first written down. Not *so* much so, however, as to get them out of bed any earlier on cold winter mornings.

Mat Carmody
Richmond-upon-Thames College

[1] We shall assume from now on that we are dealing with heaps of grains of sand. We shall also assume that collections of grains of sand have, where necessary, the right structures and shapes to be heaps.
[2] 'sorites' translates as 'heaper, one who heaps'.

[3] We are not trying to capture, in a logic, how we actually reason, because we often reason illogically. We are trying to discover the correct way to reason. In the same way, when we do mathematics, we're not investigating how we actually operate with numbers, because we often make mistakes. We are trying to discover the right answers - the answers we should get to if we reason correctly.