R. M. Sainsbury, Paradoxes, 3rd. edition, Cambridge 2009: Cambridge UP

3 Vagueness: the paradox of the heap

3.1 Sorites paradoxes: preliminaries

Suppose two people differ in height by one-tenth of an inch (0.1″). We are inclined to believe that either both or neither are tall. If one is 6′ 6″ and the other is 0.1″ shorter than this, then both are tall. If one is 4′ 6″ and the other is 0.1″ taller, then neither is tall. This apparently obvious and uncontroversial supposition appears to lead to the paradoxical conclusion that everyone is tall. Consider a series of heights starting with 6′ 6″ and descending by steps of 0.1″. A person of 6′ 6″ is tall. By our supposition, so must be a person of 6′ 5.9″. However, if a person of this height is tall, so must a person one-tenth of an inch smaller; and so on, without limit, until we find ourselves forced to say, absurdly, that a 4′ 6″ person is tall (see question 3.1), indeed, that everyone is tall (see question 3.2).

3.1

How should one respond to the objection that someone 4′ 6″ tall may be tall for a pygmy?

3.2

How might one use similar reasoning on behalf of the conclusion that no one is tall?

In ancient times, a similar paradox was told in terms of a heap, and a Greek word for "heap" – *soros* – has given rise to the use of the word "sorites" for all paradoxes of this kind. Suppose you have a heap of sand. If you take away one grain of sand, what remains is still a heap: removing a single grain cannot turn a heap into something that is not a heap. If two collections of grains of sand differ in number by just one grain, then both or neither are heaps. This apparently obvious and uncontroversial

supposition appears to lead to the paradoxical conclusion that all collections of grains of sand, even one-membered collections, are heaps.

Suppose you are looking at a spectrum of colors painted on a wall through a device with a split window which divides the section of the spectrum you can see into two equal adjacent areas. Suppose the spectrum is so broad and the windows so narrow that the colors in the two visible windows are indistinguishable, no matter where in the spectrum the device is positioned. The device is first placed at the red end of the wall, and then moved gradually rightward to the blue end. It is moved in such a way that the area that was visible in the right-hand window in the previous position is now visible in the left. At the beginning you will unhesitatingly judge that both areas are red. At each point, the newly visible area will appear indistinguishable from an area that you have already judged to be red and that is still visible. One feels bound by the principle that if two colored patches are indistinguishable in color, then both or neither are red; yet clearly there must come a time when neither of the visible areas is red. This looks like a contradiction: on the one hand, no two adjacent areas differ in color and the first is certainly red; on the other hand, the first area differs in color from some subsequent area (see question 3.3).

3.3

How would one construct a parallel argument with the paradoxical conclusion that no men are bald?

What do these paradoxical arguments have in common? In each case, the key word is *vague*: "tall," "heap," "red." A vague word admits of borderline cases – cases in which we do not know whether to apply the word or not, even though we have all the kinds of information that we would normally regard as sufficient to settle the matter. We may see how tall a man is, or even know his height to a millimeter, yet we may be unable to decide whether he counts as tall or not. We may see a collection of grains of sand, and even know exactly how many grains the collection contains, yet not know whether it should be called a heap or not. We may see a color under the most favorable conditions imaginable, yet not know whether it should be called red or orange. Our ignorance does not manifest failure to understand our language, and we may not be ignorant of any of the matters that one would expect to resolve the problematic issue: how tall someone is, how big a collection of grains is, how something looks.

Vagueness gives rise to *borderline cases*, ones in which we do not know what to say, despite having all the information that would normally fix the correct verdict. But what is vagueness itself? Is it a property of language?

Or of the world? Or of ourselves? On these questions there have been many opinions, illustrated in these claims:

1. Vagueness is *absence of fact*. When it is vague whether someone is tall, *there is no fact of the matter* whether or not he is tall. The reason we do not know what to say in borderline cases is that *there is nothing to know*.
2. Vagueness is *absence of definite truth*: a person is borderline for "tall" just on condition it is neither definitely true nor definitely not true that she is tall.
3. Vagueness is *absence of a sharp boundary*. E.g. in a series of men of closely similar but steadily diminishing height there is no last (definitely) tall man, and no first (definitely) non-tall man.
4. Vagueness is *incompleteness of meaning*. A vague expression is a bit like a partial function in math. Words whose meaning is fully specific and complete are not vague.
5. Vagueness is *indecision*: "The reason it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word 'outback.' Vagueness is semantic indecision" (Lewis 1986, p. 212).
6. Vagueness is *a feature of the world*: some things, like mountains, are vague, because it is vague what their spatial extent is; others, like properties, are vague because it is vague what things they apply to.
7. Vagueness is *ignorance*: there are sharp boundaries (facts of the matter, definite truth or falsehood, etc.), but we cannot know where they fall.

These claims are probably not all independent (see question 3.4), and some pairs are inconsistent. For example, if vagueness is a feature of the meaning of expressions, it seems unlikely it could be a feature of the world (see questions 3.5, 3.6).

3.4

Choose a pair of views which seem close, and see if you can argue that if one is correct, so is the other, and vice versa. If this fails, perhaps you can suggest a scenario on which we have vagueness on one but not the other of the views.

3.5

Provide an example of an inconsistent pair of views, other than the pair (4) and (6).

3.6

How could one who holds that vagueness is not a feature of reality (but only of our descriptions of reality) respond to the following argument?

> Mountains are part of reality, but they are vague. They have no sharp boundaries: it is vague where the mountain ends and the plain begins. So it is easy to see that vagueness is a feature of reality, and not just of our thought and talk.

On any reasonable view, vagueness needs to be distinguished from relativity and from ambiguity. Consider the property of *being above average in height*. Assuming that there is no problem about assigning numbers to people as measures of their height, this is not a vague property. A person is above average in height just on condition that the number that measures his or her height is greater than the number that measures the average height, and this is a completely precise condition. However, being greater than the average height, though precise, is *relative* to a given population. Being above average in height for a Swede involves being taller than being above average in height for an Eskimo, since the average height of Swedes is greater than the average height of Eskimos.

As the example shows, relativity does not entail vagueness, since it can hold of precise expressions, such as "is above average in height." Many vague predicates are also relative, but their relativity must not be confused with their vagueness. For example, "tall," unlike "above average in height," is vague and also relative. You can see that the vagueness does not entail the relativity by seeing that you could eliminate the relativity but leave the vagueness. If instead of "tall" we were to say "tall for a Swede," we would have eliminated some relativity, but the vagueness would remain (see question 3.7). Most people think that the difference of 0.1″ cannot make the difference between being tall for a Swede and not. This means that the argument of the opening paragraph will work as well for "tall for a Swede" as it did for "tall" (see question 3.8).

3.7

Does the restriction to Swedes eliminate all relativity?

3.8

Is "heap" relative in the way that "tall" is?

Vagueness must also be distinguished from *ambiguity*. Consider the word "bank." It can mean the edge of a river or a financial institution. As a result, we do not know how to answer the question "Did he go to the bank this morning?" So far, there is a parallel with vagueness. We do not know how to answer the question "Is he bald?" if the person in question is a borderline case. However, there is a difference. In the case of ambiguity, a single sentence can be used to say, or ask, more than one thing. Before communication can proceed, the audience needs to determine *which* thing is being said or asked. In the bank example, there is no one answer because there is no one question. With vagueness it is different. If someone asks of a borderline case "Is he a child?," it is not that our problem in answering is the problem of knowing *which* question has been asked; there is only one possible question involved here. The problem is quite different: if the person of whom the question is asked is a borderline case, neither Yes nor No is a clearly correct answer. This question has a single vague meaning, and that is quite different from having two or more meanings.

Vagueness is a widespread feature of our thought. Consider the following list: "child," "book," "toy," "happy," "clever," "few," "cloudy," "pearl," "moustache," "game," "husband," "table" (see question 3.9). Is this feature ineradicable? Could we replace our vague expressions by precise ones? We will answer this question differently, indeed understand it differently, depending on what we take vagueness to be.

3.9

Show that each of the words in the list is vague by briefly sketching a borderline case. Can you think of any words that are not vague?

Arguably, it is important to the role of the concept of childhood in our lives that the word "child" be vague. For example, we take ourselves to have special duties to children that we do not have to adults, but these duties are not relinquished overnight; they gradually fade away, just as childhood itself does. There is no sharp boundary to the end of our duties, just as there is no sharp boundary to the end of childhood. If we were to replace "child" by some more precise term, say "minor," as we have to for certain legal purposes, it would no longer be possible to express the obligations we feel. Our duties to people as

children may end before or persist after their legal majority, depending on what age the law selects for coming of age, and depending on whether they are, as we say, "old for their age" or "young for their age."

Consider a second example of the same general kind of argument. It is important to the role of the concept of redness in our lives that the word "red" be vague. In particular, the vagueness is essential to the observational character of "red," to the fact that we can under favorable conditions tell just by looking whether something is red or not. If we replaced "red" by an expression having sharp boundaries, say one defined in terms of wavelengths, then the right way to approach applying it would be by an exact determination of the wavelength: in losing the vagueness, we lose the observational character.

These arguments are unsatisfactory twice over. First, they seem simply to presuppose a particular view of what vagueness is, without supplying a justification. In the argument about "child," it was explicit that vagueness was to be regarded as absence of sharp boundaries. The argument about "red" at least excluded some views of vagueness, e.g. (7) (vagueness is ignorance). Second, the arguments are fallacious. The uncontroversial premises in the case of "red" are these:

Under some conditions, and for some objects, we can confidently apply or deny "red" simply on the basis of observation.

For some objects, we cannot confidently either apply or deny "red."

These properties could hold of an expression which draws sharp boundaries, for example "is more than 6′ tall." Under some conditions, and for some objects, we could be confident, on the basis of observation alone, whether to apply or deny the predicate. This is consistent with the fact that, at least unless we have a ruler, there will be objects about which we are in doubt: those close in height to 6′. The argument erred in supposing that if a color word were defined in terms of exact wavelengths, we would have to use an exact wavelength-measurer to apply it. The argument failed to show that an absence of sharp boundaries was required for, or helped explain, the observational character of some vague expressions (see question 3.10).

3.10

Assess whether the following consideration might help reinstate an interesting connection between vagueness and observationality:

> In the case of "is more than 6′ tall" one could in principle decide all cases, if one had the help of a ruler. By contrast, in the case of "red," nothing could help one decide borderline cases.

The argument about childhood is also defective. To say that we have special duties of care to children is not in itself to say that how demanding those duties are is inversely proportional to how old a person is. If we do want to say the latter, then we seem to be able to say it straightforwardly, as I have just done, and in a way which does not require "child" to fail to have sharp boundaries. One can also accept that childhood fades gradually away without accepting that "child" draws no sharp boundaries. To show this by analogy: if one is under twenty-one, the time remaining until the moment of one's twenty-first birthday gradually diminishes, but not in a way that requires lack of sharp boundaries.

Whatever the explanation, there can be no doubt that vagueness is a very widespread phenomenon. Hence there are plenty of opportunities for constructing paradoxical sorites arguments. We must look at these more closely, and see what responses are possible.

3.2 Sorites paradoxes: some options

In this section, I will make more explicit one form of soritical reasoning, and will lay out some possible responses.

Sorites reasoning depends upon the supposition that vague expressions are "tolerant": small changes do not affect the applicability of the word. If someone is tall, so is a person a millimeter shorter; if a collection is a heap, so is one otherwise similar but with just one grain less. The paradox arises because big changes, ones which obviously do affect whether or not the word applies, can be constructed out of small changes.

Here is one explicit version of a sorites argument, showing the impact of the supposed tolerance of vague words. First premise:

(1) A 10,000-grained collection is a heap.

Second premise, manifesting the supposed tolerance of "heap":

(2) If a 10,000-grained collection is a heap, then so is a 9,999-grained collection.

Tolerance underwrites numerous further premises of this form, for example:

(3) If a 9,999-grained collection is a heap, then so is a 9,998-grained collection.

And so on. Let us call the first premise the *categorical* premise and the others the *conditional* premises. (A conditional statement is one of the form "If ..., then ...")

We are just as inclined to hold to these conditional premises when they concern small numbers as when they concern large numbers, and just as

inclined to hold to them for cases in which there is genuine doubt about whether a collection is heap-sized as when there is no such doubt. For example:

(10,000) If a 2-grained collection is a heap, so is a 1-grained collection.

We are firmly convinced that neither a 1-grained nor a 2-grained collection is a heap; but this does not stop us holding that *if* a 2-grained collection is a heap, then so is a 1-grained collection. This reflects our general conviction that taking away a grain cannot turn something from a heap into a non-heap. This also attaches to borderline cases, for example:

(9,925) If a 77-grained collection is a heap, so is a 76-grained collection.

We may be in genuine doubt about whether either collection deserves to be called a heap. Yet the conditional reflects our confidence that both or neither are heaps, and this is simply another way of putting the tolerance principle: a difference of a grain cannot be the difference between a heap and a non-heap.

We do not yet have a paradox. To get it, we apply to these premises a general principle of reasoning: given p, and a conditional of the form "if p, then q," we can derive q. This principle is still called by the Latin name it was given in the middle ages: modus ponendo ponens, or modus ponens, for short. Applying it to our first and second premises yields:

A 9,999-grained collection is a heap.

Applying the principle again to the above and premise (3) yields:

A 9,998-grained collection is a heap.

Continuing in the same way, we finally end up with the result that a 1-grained (or, for that matter, even a 0-grained) collection is a heap, and this is the absurdity.

As with any paradox there are three possible responses to consider:

(a) Accept the conclusion of the argument.
(b) Reject the reasoning as faulty.
(c) Reject one or more premises.

For some vague words, (a) may seem a totally unpromising possibility to explore. Could anything reconcile us to the suggestion that everyone is tall, or that a heap cannot be demolished grain by grain so that no heap is left, or that all colors are red? Some philosophers have used sorites reasoning to argue for conclusions of just this kind. I give one example in the next section.

3.3 Accepting the conclusion: Unger's view

Consider the following argument:

(1) Given just one gram of wood, you cannot make a table.
(2) If you cannot make a table out of n grams of wood, then you cannot make a table out of $n + 1$ grams.
(3) Hence you cannot make a table, no matter how many grams of wood you have.

Crucial to the argument is the idea that a gram cannot make the difference between not enough wood to make a table and enough wood. (If you feel doubts about this, try reducing the amount to a millionth of a gram.)

There is a close analogy between this reasoning and the reasoning of the previous section concerning heaps. That reasoning seemed to establish that heaps are indestructible: however many grains you take away from a heap, it remains a heap. The present reasoning runs in the reverse direction, and seems to establish that tables are uncreatable: however many bits you add, you still do not have a table. This conclusion amounts to the claim that there are no tables.

Nothing essentially depends upon the notion of literally "making" a table. We could use similar reasoning to show that there are no stones. If a region contains only one gram of solid material, it does not contain a stone. If a region contains n grams of solid material but not a stone, then a region as similar as possible except that it contains $n + 1$ grams of solid material also does not contain a stone. If we start with a region which we take, intuitively, to be occupied by a stone, and focus on a one gram part of it, and then successively expand our considerations outward by including, one at a time, a further adjacent region occupied by a gram of matter, we appear to be able to conclude, contrary to the original intuition, that however far we expand we will not get to a region containing a stone. Generalizing, the conclusion amounts to the claim that there are no stones.

These conclusions seem strange, but they have been advanced with at least apparent seriousness by Peter Unger (1979a). We can make them seem less mad by putting them in a certain perspective. Perhaps the conclusion we should draw from the existence of the sorites paradoxes is that vague concepts are deeply flawed: they commit us to absurdities. A flawed concept is one under which nothing can fall. So even though the world may contain all sorts of stuff, it is wrong to say that we can divide it up using concepts flawed by vagueness. Rather, we have to say that nothing matches these concepts: there are no tables, no stones, etc. (see question 3.11).

3.11

*How should someone who adopts Unger's viewpoint respond to the argument which concludes that even a 1-grained collection would make a heap?

One cannot dismiss this point of view as mad. However, it would not be wrong to regard it as something of a last resort, and Unger himself accepts that if there were some way of defusing sorites paradoxes, his own solution would be unattractive. We need to look at alternative responses.

Sorites reasoning appears to be extremely simple and to use only the fundamental logical principle of modus ponens, so the prospect of rejecting it as invalid (response (b) above) is not initially tempting.

The most obvious response to explore is (c): reject one or more premises. The next two sections explore two ways in which this might be motivated.

3.4 Rejecting the premises: the epistemic theory

The epistemic theory of vagueness (7 on p. 42) holds that vagueness is nothing but ignorance. So far as their semantics go, vague words are just like precise ones: they draw sharp boundaries. The epistemic theorist will see sorites paradoxical arguments as proving his view. When the arguments take the form displayed in the previous section, the epistemic theorist sees them as establishing by *reductio ad absurdum* the falsity of at least one premise. (An argument by *reductio ad absurdum* is one which starts by deducing an absurdity from some premises, in order to demonstrate that one of the premises must be false.) The conclusion (e.g. that there are no heaps) is plainly false, the reasoning is valid, so one of the premises must be false. The first premise (e.g. the one that says that a collection of 10,000 grains can make a heap) cannot be disputed. So it is the generalization exemplified by the second premise that must be false, the "principle of tolerance":

If an n-grained collection can make a heap, then so can an $n - 1$-grained collection.

On the epistemic view, there are heaps and also non-heaps, and these are separated by a sharp line. In other words, for some number n, a collection of grains with n members can make a heap but a collection of grains with $n - 1$ members cannot. The special twist the theory gives, from which it earns its name and also any plausibility it may have, is that we cannot, even in principle, know where this sharp boundary lies. That is why borderline cases are confounding: they ask for a verdict in a situation in which we cannot know what the right verdict is.

The epistemic theory is generally greeted with incredulity. How could there be a sharp boundary to the heaps, the bald men, the tall men, childhood, and so on? We can agree with epistemic theorists that, if there are such boundaries, we cannot even in principle know where they fall; but in contrast to epistemicists, we might take this as a reason for thinking that there are no such boundaries. If they really exist, what could prevent us telling where they lie, given enough time and effort?

Epistemicists have two things to say in response. First, they will insist that they do not accept the *verificationist theory of meaning*, or *verificationism* for short. This theory, characteristic of the logical positivism which flourished in the middle years of the twentieth century, has it that any sentence we can understand is one which, in principle, we could tell to be true, or false, as the case may be. Verificationism is inconsistent with the epistemic theory, which holds that if an object is a borderline case for a vague word, we have no means of determining the truth or falsehood of an application of the word to the object. However, at least in the crude form presented here, verificationism is now not widely accepted, so few would wish to attack the epistemic theory from this angle.

Epistemicists have a second response: an explanation of why we are incurably ignorant in borderline cases. On the epistemic view, it is because our cognitive mechanisms, for example our senses, require a margin for error if they are to deliver knowledge. This is so in ordinary cases, ones not involving vagueness. For example, even if I believe, glancing round the stadium, that 3,973 people are present, when in fact this is the actual number, my true belief cannot count as knowledge, for there being one person more or less would not have affected my estimate. Even though I am right, this shows that I am right by accident, and so do not know. The ignorance that, on this view, is distinctive of vagueness is of a special kind: it relates to the precise meanings of the concepts I use (at least according to some epistemicists, e.g. Williamson). I actually use the concept *red*. Invoking this concept, let us suppose that a certain patch p is the last red patch on the wall we described earlier, where the red patches become less and less red, and are followed by orange patches. Suppose, moreover, that I *believe* that p is the last red patch in the series. The epistemic theorist argues that this belief, though true, does not count as knowledge, for there is no margin for error. I could easily have come to possess, in place of my actual concept *red*, a slightly different concept, say *red**, which drew the boundary in a slightly different place, so that the last red* patch would have been to the right of p. Just as there being one person more or less in the stadium would have had no impact on what belief I formed, so my having this different concept would not have stopped me believing, falsely, that p is red*. Knowledge requires being non-accidentally right, and in cases like these, this in turn requires a margin for error. No margin for error, no knowledge (see question 3.12).

3.12

The version of epistemicism described here turns on a lack of margin for error concerning the concepts we use. Why should not an epistemic theorist instead rely on the margin for error needed by our perceptual systems? (We cannot distinguish the color of p from that of adjacent patches, so it was just luck we selected p as the last red.) First hint: compare the stadium case. Second hint: check out Williamson (1994: section 8.4).

Epistemic theorists may thus explain our ignorance, but they still have to establish the claim that there is anything to be ignorant of: that vague predicates draw sharp boundaries. Taking up this challenge in full would involve showing that alternative responses to sorites paradoxes are inadequate (see question 3.13).

3.13

Evaluate the following argument against the epistemic theory, as applied to color predicates:

> If a predicate stands for a manifest property (one which under some conditions detectably obtains), then under optimal conditions for manifestation, if there is a fact of the matter whether or not it obtains, that fact is detectable. We can view a region of the colored wall (one borderline for "red") under optimal conditions without being able to detect the presence or absence of redness; so there is no fact concerning whether that region is or is not red.

3.5 Rejecting the premises: supervaluations

One can reject one of the conditional premises without adopting the epistemic theory. One way is the *theory of supervaluations*. The underlying thought is that vague words are incomplete or defective in their meaning (see diagnoses 4 and 5 on p. 42), and this is manifest in borderline cases. For these cases, the meaning of the expression is inadequate to determine whether or not it is true of the object. One can specify the meanings of such incomplete expressions by specifying the various ways in which the incompleteness could have been remedied, that is, the various ways in which the expression could have been made precise.

Many vague words have a meaning which depends upon some underlying ordering, one which induces principles like this: if anything is a heap, then anything otherwise similar except bigger is a heap too; and if anything is a non-heap, anything otherwise similar but smaller is also a non-heap. If a vague word like "heap" had a complete meaning, then for every candidate for being a heap, the word would either definitely apply or definitely not apply. The upshot would be a sharp cut-off: a smallest heap (a least point in the underlying ordering). On supervaluationist views, vagueness is lack of a complete meaning. The incompleteness shows up in borderline cases, to which the word neither definitely applies nor definitely fails to apply. We can envisage various ways in which the meaning of the word could be completed, various ways (consistent with the underlying ordering) that it could have been made precise. By indicating which these are, we can reveal a vague word's incomplete meaning, rather as we could reveal what Jack's unfinished house is like by showing all the various ways in which he could complete it. What has actually been constructed is what all these possible completions have in common. Likewise, what there is of the meaning of a vague word is what all the various ways of completing it agree on.

For supervaluationists, the crucial tool is the notion of a sharpening (sometimes called precisification). A sharpening, $s(w)$, of a vague predicate, w (a word like "heap" which can be true or false of objects), meets the following conditions:

- If w is definitely true of something, then $s(w)$ is true of it.
- If w is definitely false of something, then $s(w)$ is false of it.
- For each object, $s(w)$ is either true of it or false of it.
- $s(w)$ respects the underlying ordering (if there is one). For example, if s("tall") is true of someone 6′ tall, it is also true of someone 6′ 1″.

The meaning of a vague predicate is then to be described in terms of *all* the sharpenings which meet these conditions (analogously to *all* the ways Jack could complete his house). Where the sharpenings disagree, we have vagueness; where they agree, we do not (just as where the house models agree, we have finished construction, and where they disagree we do not).

Truth, for a supervaluationist, is identified with *truth on all sharpenings* and falsehood with *falsehood on all sharpenings*. What matters is what is really present, the incomplete meaning (likewise, on the analogy, the incomplete construction). A sentence which ascribes a vague predicate to a borderline case will be neither true nor false. This is the critical result which the supervaluationist uses to reject a premise of the sorites reasoning. The

first (categorical) premise, for example that with 10,000 grains one can make a heap, cannot be rejected. But many of the conditionals fail of truth, within supervaluationist frameworks. Let us see exactly how.

Suppose α and β are borderline for "heap," and that α has one more grain than β. (Perhaps α has 76 grains and β has 75.) On the basis of what has been said so far, "α is a heap" is neither true nor false: since α is a borderline case for "bald," s("bald") will be true of α for some but not all sharpenings s; so the sentence will be true on some sharpenings (and so not false) and false on some sharpenings (and so not true). Likewise for "β is a heap." Consider the conditional

If α is a heap, then β is a heap.

For this to be true, it needs to be true on every sharpening. Given the general conditions on sharpenings, there is a sharpening of "heap" on which "α is a heap" is true and "β is a heap" is false. This is a sharpening which makes a 76-grained collection the smallest heap. The conditional is false on this sharpening (a conditional with a true antecedent and false consequent is false). Hence the conditional is not true on all sharpenings and so is not true. (Neither is it false, as analogous reasoning would show (see question 3.14).) At least one such conditional is sure to feature somewhere in a sorites argument constructed in the style envisaged in section 3.2. Hence, according to supervaluation theory, not all the premises of this kind of sorites argument are true. This is enough to offer a resolution to the paradox. The theorist need not specify a particular non-true premise.

3.14

How could it be shown that the conditional is not false?

It is not that, on this account, *any* sentence in which a vague word is applied to an object in its penumbra is neither true nor false. For example, the sentence

α is either a heap or not a heap

is true on all sharpenings, and so true, even if α is a borderline case. Every sharpening draws the line somewhere. Wherever it draws it, "heap" will either be true of α on that sharpening, or false of α on that sharpening. Hence, for any sharpening, either "α is a heap" or "α is not a heap" is true on that sharpening. Therefore " α is a heap or α is not a heap" is true on every sharpening – that is, on the account, *true* (see question 3.15).

3.15

Does the supervaluational account entail the following: either "α is a heap" is true or "α is not a heap" is true?

What interesting point about the supervaluational treatment of "or" does your answer highlight?

To sum up, supervaluational accounts purport to dissolve the paradox by showing that not all the premises of the paradoxical argument are true. In particular, the principle of tolerance does not hold: we cannot infer that if someone is tall, so is someone just a millimeter shorter. One merit of the account is that it preserves standard logic. For example, it brings all instances of "either *p* or not *p*" out true. On this account, we need have no truck with the response I labeled (b): giving up some logical principles.

I shall consider four problems for supervaluational accounts. The first is that preserving classical logic may not be such a good thing. Vagueness has often been supposed to throw doubt on the logical principle called the Law of Excluded Middle, exemplified by:

Either he is an adult or he is not.

It might be suggested that given that there are borderline cases of adults, we should be reluctant to affirm this instance of the law. Affirming it might be exploited in an argument which could reasonably be held to be flawed:

Either he is an adult or he is not. If he is an adult, then watching the hard-porn movie will do him no harm. If he is not an adult, then he will not understand it, so, in this case too, watching it will do him no harm. Either way, it will do him no harm to watch it.

We might object: the argument fails to take account of the person on the borderline between childhood and adulthood. For him, it is not right to say "Either he is an adult or he is not." It is precisely because he is between childhood and adulthood that seeing the movie might harm him.

The second problem for supervaluational theories is that they assign intuitively the wrong truth value to a sentence which is central to sorites paradoxes:

For some number n, an n-grained collection is a heap but an $(n-1)$-grained collection is not.

We intuitively tend to believe that this sentence is false (and likewise similar sentences for other vague words), for it seems to affirm the existence of a sharp boundary, and so fails to do justice to the vagueness of

"heap" (see question 3.16). On the supervaluational theory the sentence comes out as true, since it is true on every sharpening (see question 3.17).

3.16

Would the epistemic theorist agree with this judgment?

3.17

How is this demonstrated? (You might want to check out Sanford 1976.)

The third problem for supervaluational theories is that the underlying conception of vagueness as nothing but incompleteness of meaning is at a minimum insufficient. Predicates which, intuitively, are not vague may be incomplete in meaning. For example, we might define a minor by the following clauses:

(1) People who have not reached their seventeenth birthday are minors.
(2) People who have reached their eighteenth birthday are not minors.

This definition intuitively does not make "minor" vague. But its meaning is incomplete: it fails to speak to persons who are seventeen years old.

The clauses which define "minor" are entirely precise, and that is why we are right not to think of "minor" as vague. The problem relating to seventeen-year-olds is not vagueness or imprecision but incompleteness. We find the same phenomenon among familiar words. Perhaps "pearl" is true of anything made of a certain material and formed in an oyster and false of anything not of that material. This leaves undetermined the word's application to pearl-sized lumps of pearl-material synthesized outside of any oyster. I suggest we should see "pearl" not as vague, but as incomplete: there is a blank where there should be a rule. If this is right, supervaluation theory's guiding idea, that vagueness is a kind of incompleteness, is not enough to characterize vagueness. If vagueness is incompleteness of meaning, it must be a special kind of incompleteness.

The fourth alleged problem for the supervaluational account of vagueness presented here is that it fails to allow for "higher-order vagueness." The account presupposes that the notion of a sharpening is precise. Yet it seems that the notion is itself vague: might there not be a borderline case of a word being definitely true of something? "Is he *definitely bald*? I'm not sure; maybe so, maybe not." If there is vagueness here, then it would be wrong to assume that the notion of a sharpening is precise.

Let us suppose that is right: *sharpening* is vague. How does this amount to a problem for supervaluational accounts? One could not turn to supervaluationism to provide a precise description of vagueness; but that is consistent with its providing a correct description. Even if there is vagueness about which sentences are true on all sharpenings, it is not vague whether there are sentences in standard sorites arguments which are neither true on all sharpenings nor false on all sharpenings (see question 3.18). The supervaluationist's explanation of how standard sorites arguments should be resisted, by denying that all the premises are true, remains intact. The demand that vagueness be described in precise terms goes well beyond the demand for an adequate solution to sorites paradoxes. So perhaps there is no problem after all (see question 3.19).

3.18

This assumes that, even if "sharpening" is vague, each vague predicate has at least one clear case of a sharpening. What example could you give to illustrate this?

3.19

How would you assess the following objection to a version of the supervaluational account which allows that the notion of a sharpening is vague?

> If "sharpening" is vague, then no sentence can be definitely true. Truth involves appeal to *all* sharpenings; what is to count as a sharpening is vague, so it is not definitely true of any collection that these are all the sharpenings there are (with respect to a given predicate). However, it is absurd to suggest that "Yul Brynner was bald" is anything other than *definitely* true.

3.6 Rejecting the reasoning: degrees of truth

We envisaged three possible responses to the paradox:

(a) to accept the conclusion;
(b) to reject the reasoning; or
(c) to reject one or more premises.

We have taken (a) to be a last resort. We have explored two versions of response (c): the epistemic theory and the supervaluational theory. I now want to turn to what has sometimes been classified as a (b)-type response. (As we will see, this classification is open to question.)

We are strongly disinclined to allow that there could be anything wrong with modus ponens. Nevertheless, some theorists have tried to place the

blame on this principle of reasoning, and I will try to explain their grounds.

When asked to assess a claim to the effect that a sixteen-year-old is an adult, it is natural to say something like "That is *to some extent* true," or "There is *a certain amount of truth* in that"; likewise whenever a vague predicate is applied to a borderline case. The response to the paradoxical argument I now wish to consider takes this very seriously. The suggestion is that we introduce *degrees of truth*. If a predicate like "adult" definitely applies to an object we will say the application has the maximum degree of truth, conventionally 1. If a predicate definitely does not apply to an object we will say the application has the minimum degree of truth, conventionally 0. Borderline cases will be registered by intermediate degrees of truth. Ascribing "bald" to a man who nearly qualifies as bald will rate a degree of truth closer to 1 than applying it to a man who nearly qualifies as non-bald. A degree of truth theory thus takes very seriously the point that the meaning of a vague word says something about the borderline cases, by contrast with supervaluationists, who see borderlines as marking incompleteness of meaning. Degree theory seeks to represent *what* the meaning says by the various degrees of truth.

How can a degree theory dissolve the paradoxical argument? It must assign the highest degree of truth to the categorical premise of the argument and the lowest degree to the conclusion; how will it treat the conditional premises?

Suppose collections of grains of sand start becoming borderline for "heap" at around the 100 mark. What should a degree theorist say about the conditional

> If this 95-grained collection is a heap, so is this 94-grained collection?

The antecedent of the conditional is

> This 95-grained collection is a heap.

The consequent is

> This 94-grained collection is a heap.

According to the degree theory, the antecedent is nearly but not quite true. Perhaps it is assigned the degree of truth 0.96. The consequent is also nearly true, but not quite so nearly true as the antecedent. Perhaps it is assigned the degree of truth 0.95. What degree of truth should be assigned to the conditional itself?

There is room for variation in detail, but the general idea is that if the antecedent of a conditional is truer than its consequent, then the conditional cannot be wholly true; thus the conditional in question needs to be assigned a degree of truth less than 1. The justification for this lies in part with the analogy with the standard case in which degrees of truth are not taken into account: we say that a conditional whose antecedent is true and whose consequent is false cannot be true, because a conditional should not lead from truth to falsehood. Analogously, a conditional should not lead to a lower degree of truth. The greater the amount of truth lost in the passage from antecedent to consequent, the lower the degree of truth of the whole conditional.

So far, the degree theorist's response is of type (c): reject the premises. On this theory the conditional premises, though very nearly true, are not quite true. Hence we need not be fully committed to them. Hence the paradoxical argument does not commit us to the paradoxical conclusion. However, the account needs to go further. Even if we are not fully committed to the premises, they are very nearly true. The degree theorist has to explain how we can have premises that are all very nearly true, yet a conclusion that is wholly false.

One way to do this, though not the only one, is to deny that the conclusion follows from the premises, and thus deny the validity of modus ponens; this is a response of type (b). On the degree-theoretic account envisaged here, modus ponens does not preserve degree of truth: the conclusion of an argument of the form "If p, then q; p, therefore q" may have a lower degree of truth than any of the premises. The conditional mentioned earlier, about 95-and 94-grained collections, is extremely close to the whole truth; perhaps it has a truth degree of 0.99. The antecedent, we suggested, had a truth degree of 0.96. Yet applying modus ponens yields a conclusion with a truth degree of only 0.95, lower than the truth degree of either of the premises. Modus ponens is straightforwardly valid as applied to sentences with the extreme truth degrees, 0 or 1: that is, one cannot get a conclusion with degree less than 1 from premises of degree 1. However, in the intermediate degrees the application of modus ponens can lead to a "leakage" of truth. The leakage may be small for each application, but can be large if the number of applications is large, as in the case of the paradoxical argument; and it can be regarded as a sufficient condition for the invalidity of modus ponens.

We earlier thought that modus ponens was a principle that simply could not be abandoned. What the degree theorist suggests, however, may well be consistent with all we really believed about modus ponens. There are two reasons for this.

The first is that normally we have in mind only the cases in which modus ponens is applied to sentences that are (*completely*) true or (*completely*) false. For these cases, the degree theorist's view agrees with our intuitions.

The second is that it is arguably not correct to assume, as we have so far, that an argument's validity is properly understood in terms of preservation of degree of truth, in the sense that the conclusion of a valid argument must have a degree of truth no lower than that of the least true premise. Perhaps it is better (and it would certainly preserve tempting analogies with degrees of belief) to say that the conclusion of a valid argument cannot have a greater degree of falsehood than the sum of the degrees of falsehood of its premises (Edgington 1992). (If a number, n, measures a degree of truth, then $1-n$ measures a corresponding degree of falsehood.) Then degree theory need give us no reason to say that modus ponens is not valid (see question 3.20). If we adopt this position, we must reclassify the degree theory as one which rejects the premises of sorites arguments, or rather, does not fully accept them.

3.20

Why not?

A full defense of the degree of truth theory would require the consideration of a number of issues that I shall briefly mention. First, it is necessary to say something about what a degree of truth is. Second, some account must be given of the source and justification of the numbers that are to be assigned as degrees. Third, the full implications of the degree theory for logic must be set out and defended.

A degree theory treats vagueness as a semantic matter, a matter of meaning, and not as an epistemic one, a matter of knowledge or ignorance. The semantic theory of degrees of truth registers the semantics of a vague predicate as different in kind (since involving the intermediate degrees) from the semantics of a precise predicate. So the first thing a degree theorist would have to do is refute the epistemic theory. Let us suppose that this has somehow been accomplished. The remainder of the defense presupposes that vagueness involves there being no sharp boundaries.

A key property of truth is marked by the platitude that we aim to believe what is true. If we could show that degrees of truth had an analogous property, we would have gone some way toward explaining what a degree of truth is.

Suppose that you are fairly sure that Arkle won the Gold Cup in 1960. Your memory may fail you about some matters, but you are pretty reliable about the history of the turf. You reckon that you have a very much better than fifty-fifty chance of being right that it was Arkle. If you are at all attracted by gambling, it will be rational for you to bet on his having won, if you can get odds as good as even; for if you follow this policy generally, you will win more than fifty times out of a hundred. A policy that will result in your winning more than you lose is a policy that it is rational for you to pursue. It is rational to perform a particular action that is required by the pursuit of a rational policy.

We want to believe what is true, but we do not always know what is true. The greater the confidence we have in a proposition, the more it affects us as if we believed it to be true. If we are almost certain that our house will not burn down, we will not spend much money insuring it against fire. If we are almost certain that we shall be alive tomorrow, we do not waste much time today making arrangements for our death.

It is rational for our beliefs to vary in strength, reflecting variations in our confidence, and thus variations in our assessment of the quality of our information. We may be less than totally confident because we are less than fully informed. The less than total confidence mirrors our epistemic deficiencies.

Vagueness may also lead to less than total confidence. Suppose you know, having had it on impeccable authority, that all and only red mushrooms are poisonous. You wish to kill Jones. All other things being equal, you would prefer to poison him, and you would prefer to do so using mushrooms, so that it will look like an accident. However, the only mushroom you can find right now, though reddish, is not a clear case of red. Will you use it to try to poison Jones? It depends upon how important it is to you to succeed, how important it is to succeed at first attempt, and how soon Jones must die if his death is to be of service to you. How reasonable it is to use the mushroom depends upon the weights you assign to these other factors, and upon your degree of confidence in the redness of the mushroom. The more confident you are that this mushroom is really red, the more reasonable it is to use it; the less confident, the less reasonable. In the context, this confidence affects your action in the same way as would a lack of confidence springing from lack of information, from fear that your memory fails you, or whatever. From the point of view of action, it is rather as if you had less than total confidence in the statement "This mushroom will do the job."

There is also a sharp contrast. Less than total confidence springing from incomplete evidence or fear of unreliability mirrors our epistemic deficiencies; but according to degree theorists, less than complete confidence springing from an appreciation of vagueness does not. If the mushroom is a borderline case, it is not your fault that you are unsure whether it should be counted as red; indeed, you would be at fault if you took yourself to be required firmly to classify it either as red or as not red. Continuing our assumption that the epistemic theory has been rejected, we can go further: no matter how perfect your memory and senses, no matter how infallible your reasoning, the mushroom stays on the borderline. On the question of whether the mushroom is red, an omniscient being could do no better.

Where we have incomplete information, or unreliability, there is a chance of improvement: we can in theory raise our confidence by getting more information. Where we have vagueness, there may be no chance of improvement. Given your language and the way the world is, you can do no better than have partial confidence in "This mushroom is red." Truth is what we seek in belief. It is that than which we cannot do better. So where partial confidence is the best that is even theoretically available, we need a corresponding concept of partial truth or degree of truth. Where vagueness is at issue, we must aim at a degree of belief that matches the degree of truth, just as, where there is no vagueness, we must aim to believe just what is true.

The second part of a defense of a degree of truth theory is to explain and justify the origin of the numbers that are assigned as degrees of truth. Suppose there are two mushrooms, both borderline cases of red, but one redder than the other. If you want to commit the poisoning, and you have full confidence in the information that all and only red ones kill, you should choose the redder if you choose either. The redder one must be closer to being a definite case of red. This suggests how we could justify assigning degrees of truth: we have to assign a higher degree to a redder object; or, if we are dealing with "heap," we must assign a higher degree to borderline cases the more numerous they are, that is, the closer they come to being definite cases for "heap." In short, the source and justification of assignments of degrees of truth would lie in our comparative judgments involving borderline cases (see questions 3.21, 3.22).

3.21

How would you respond to the following objection?

It is one thing to say that the comparative form of "red," namely "redder than," is to be used as the basis for assigning degrees in connection with "red"; but it is quite another thing to apply this to "heap." The basis for the assignments would be comparisons involving "heaper than"; but this is nonsense.

3.22

How would you respond to the following objection?

I agree that there are degrees of redness, but I cannot see that this means that there are degrees of truth.

The third part of the defense of a degree theory would involve justifying how one ascribes degrees of truth to logically complex sentences. Degree theories of the kind under consideration, in which the degree of truth of a complex sentence is determined by the degree of truth of its constituents, depart from ordinary, so-called classical, logic. Whereas classical logic has it that all sentences of the form "p and not-p" are false, and all sentences of the form "p or not-p" are true, some degree theorists demur. On their view, when p has only a medium degree of truth, "p and not-p" will not be completely false, and "p or not-p" will not be completely true. We have already seen, in the case of the argument about the harmful effects of watching hard-porn movies, that there is at least some case for holding that, if p is vague, "p or not-p" is not without qualification true. Furthermore, the naturalness of "It is and it isn't," as a response to the question whether a borderline-case mushroom is red, gives at least a preliminary indication that the degree theorist is right to recognize that not all instances of "p and not-p" are completely false.

However, there are problems. Standard forms of degree theory assign a conjunction a degree of truth equal to that of the least true conjunct. This means that, if x is a clear case of something red, y is a borderline case of red, and both x and y are, to the same degree, intermediate cases of something small, the conjunctions "x is red and x is small" and "y is red and y is small" will be assigned the same intermediate degree of truth, which is unintuitive (see question 3.23). Again, suppose Eve is definitely female and a borderline case for being an adult, so that "Eve is an adult" is assigned an intermediate degree of truth. Assume that it is definitely true that a woman is an adult female. Standard degree theory cannot distinguish between "Eve is an adult or Eve is a woman" and "Eve is an adult and Eve is not a woman," assigning both the same intermediate degree (Fine 1975).

3.23

Which conjunction do you think should have the higher degree of truth? Cf. Edgington (1992).

The degree theory, like the supervaluation theory, appears to be at risk from considerations relating to higher-order vagueness. On the colored wall discussed previously, red gradually gives way to orange from left to right. At the left end, a sentence "This patch is red" will presumably be assigned degree 1. However, which is the last patch relative to which the sentence is assigned degree 1? If we allow that there is a last such patch,

then it seems that we have introduced a borderline after all: between the definite reds and the others. If we say that there is no such patch, we seem to be committed to the absurdity that even when the sentence is applied to an orange patch it is still awarded the maximum degree of truth.

Part of the problem is that we tend to think of semantic theories, like supervaluational theory or degrees of truth theory, as themselves expressed in a sharp language (the metalanguage). This means that we tend to think that vagueness can be described in precise terms. The phenomenon of higher-order vagueness suggests (though it does not entail) that we are mistaken in this tendency.

If we try to describe a vague language by a vague one, we may not have conquered the problems concerning vagueness with which we began. Sorites paradoxes may threaten the language in which our semantics is expressed. If "has degree of truth 1" is precise, then it seems to draw a sharp boundary where no boundary should be. If it is vague, we will be tempted to suppose that if applying "red" to one of two indistinguishable patches merits degree of truth 1, so does applying it to the other, and we will be set on the familiar slippery sorites slope.

3.7 Vague objects?

Are there vague *objects*, or is vagueness something that arises, not from the way the world itself is, but rather from how we describe it? This question is answered by the epistemic theory: the home of vagueness lies in our cognitive faculties, not in the world. An answer is also (optionally) suggested by the degree theory: if the degrees are based on the responses of subjects to borderline cases, perhaps vagueness is a matter of human psychology rather than a feature of the world (Schiffer 2003). Those who believe that vagueness is properly to be described by saying that borderline cases are ones in which there is no fact of the matter might lay the blame on the meanings of the words; but they also might lay the blame on the world. Can we make sense of the idea that the world itself is vague?

We can start by reverting to an earlier question (3.6). The argument for discussion went like this:

> Mountains are part of reality, but they are vague. They have no sharp boundaries: it is vague where the mountain ends and the plain begins. So it is easy to see that vagueness is a feature of reality, and not just of our thought and talk.

Even if we like the conclusion, we should not accept this argument for it. Given our language, which contains words like "mountain," we can ask a vague question: does this spot belong to the mountain or to the plain?

However, if we could give a complete description of the world without making use of such a vague expression, we would have no inclination to infer from the vagueness of the question to the vagueness of the world. We would not seem to be greatly handicapped, in describing the world, if we lacked the word "mountain": we could just draw the contour lines on our maps. In many cases, for example "heap," our use of the vague word is guided by sharp underlying facts, for example, by (in part) how many grains the collection contains. Each collection has a definite number of members, and we could in principle give more information about a collection by the sharp fact of how many members it has than by the vague matter of whether or not it is a heap. So the displayed argument does not undermine the view that vagueness comes from our thought and talk, rather than being an objective feature of the world.

Let us consider an old story. Theseus had a ship. When a plank rotted, it was replaced, and thanks to the repair the ship remained in service. After a while, none of the original planks were left. Likewise for the other kinds of parts of the ship – masts, sails, and so forth. Did Theseus' ship survive? There was a ship in continuous service, and we incline to hold that this is indeed Theseus' ship, much repaired. But suppose that someone had kept the rotted planks and other parts and then reassembled these into a (possibly unseaworthy) ship. Does this have a better claim to be the original ship of Theseus? There is vagueness of some kind here. The question is: is the ship *itself* vague, or does the vagueness end with the word "ship," leaving the ship itself uncontaminated?

It seems to me that the second answer is the right one. In such a case, we can give an agreed and relatively precise account of the "facts of the matter." We know just what happened. It is a verbal question to which object, if any, we ought to apply the phrase "the ship of Theseus." So the vagueness comes from words.

This view could be supported by a two-stage argument. First, show that *identity* is not a vague relation; that is, show that questions of the form "Is this thing the same as that thing" have definite answers. If so there should be a definite answer to whether the long-serving ship that Theseus still sails is the same as the ship later reassembled from the discarded parts of Theseus' original. The suggestion is that, quite generally:

If β is α, then β is definitely α; and if β is not α, then β is definitely not α.

Suppose β is α. It seems indisputable that, for any object x,

x is definitely x.

So "is definitely α" is true of α. If β is α, anything true of α is true of β. So "is definitely α" is true of β. Hence:

β is definitely α (see question 3.24).

3.24

Can you establish the rest of what is needed for the non-vagueness of the identity relation, namely:

if β is not α, then β is definitely not α?

For help with this tricky question, consult Evans (1978) and Wiggins (1986).

Suppose we think that it is vague whether the ship Theseus still sails is the same as the one assembled from the discarded parts; that is, that they are not definitely identical. The above argument would show that we should conclude that the ships are distinct.

The second stage of the argument involves showing that if identity is not a vague relation, then objects are not vague. The idea is that if an object were vague, it would be a vague matter what object it is identical with. Since the first part of the argument has supposedly shown that identity is not vague, the conclusion is drawn that objects are not vague.

I close with four qualms. First, the second step of the argument is not cogent as stated. It is not unreasonable to suppose that, if there are any vague objects, collections with vague membership conditions are included. So if there are uniformly colored blocks of various colors and weights on my table, including various red ones and orange ones, the collection of red blocks on my table has a vague membership condition, and for some block, it may be a vague matter whether it belongs to the collection or not. The collection seems as good a candidate as any for being a vague object. Yet we might insist that the identity conditions for such collections are completely precise. If we ask whether the collection of red blocks is the same collection as the collection of heavy blocks, we might insist that an affirmative answer requires that the collections coincide in every respect: all the definite members of the red collection must be definite members of the heavy collection and vice versa; all the definite non-members of the red collection must be definite non-members of the heavy collection and vice versa; and so on through any further distinctions. There would then never be any vagueness about whether this collection is the same as or different from that collection. We thus have not excluded the possibility of vague objects without vague identity.

Second, denying that there are vague objects seems to presuppose that the "facts themselves" are precise. I said that, in the case of Theseus' ship, the facts of the matter are "relatively precise." They are precise relative to the vagueness of "ship," since they can be stated without using that word. However, other words, such as "plank," have to be used. This is just as vague as "ship." Can we be sure that there is a range of ultimate facts that can be described without using any vague expressions at all? Such a belief would demand careful justification.

The third qualm is this: identity over time, as discussed in the case of the ship, must be governed by principles such as: replacing some, but not too many, parts of an artifact does not destroy it, but leaves you with the very same artifact. Such principles are vague. How could the identity relation, which they determine, be precise (see question 3.25)?

3.25

How would you evaluate the following argument?

> If you exist at all, you are a vague object, for we believe that a molecule more or less cannot make the difference between whether you exist or do not. On this basis, we can construct a paradoxical argument: taking away one molecule will not make you cease to exist, taking away one more will not make you cease to exist, and so on; thus you can exist even if no molecules of you do. This shows that you are as vague as a heap. However, there are no vague objects, therefore you do not exist.

For essentially this argument, see Unger (1979b).

The final qualm is more technical: the best implementation I know of the semantic approach mentioned at the end of the previous section, in which a vague language is used to describe vagueness, posits the existence of vague objects, namely, vague sets. This might give one a starting point from which to extend belief in vague objects more widely.

Suggested reading

The *Stanford Encyclopedia* is likely to be the first port of call for those wishing to explore vagueness further. Two relevant entries are Sorensen's (2006a) on vagueness (plato.stanford.edu/entries/vagueness/) and Hyde's (2005) specifically on sorites paradoxes (plato.stanford.edu/entries/sorites-paradox/). Both these entries have good bibliographies. For overviews of vagueness, which also contain defenses of the authors' own views, see Williamson (1994, ultimately defending epistemicism) and Keefe (2000, ultimately defending supervaluationism). The present chapter does not mention contextualist approaches to vagueness and the sorites. Starting points would be Kamp (1981), Raffman (1994) and Graff-Fara (2000). For

a very comprehensive bibliography, see the Arche website at St. Andrews, in particular: www.st-andrews.ac.uk/~arche/projects/vagueness/bibliography.shtml.

Section 3.1

An introductory article is Black (1937); see also Dummett (1975). For a survey of treatments of sorites paradoxes, see Sainsbury and Williamson (1995). For an argument for the utility of vagueness, see Wright (1975). For his more recent views see Wright (1987) and (2001).

Section 3.3

For the view that there are no tables, see Unger (1979a).

Section 3.4

The epistemic theory can be traced back to Chrysippus (*c.* 280 – *c.* 207 BC). Its current main defenders are Sorensen (1988) and Williamson (1992a, 1992b, 1994) (they defend different versions of the theory). See also Cargile (1965). Among many critics, see Wright (1995) and Ludwig and Ray (2002).

Section 3.5

A classic text for the supervaluation theory is Fine (1975); see also Van Fraassen (1966); Kamp (1975). These papers involve some technicalities. For a less formal account of the underlying idea, see Dummett (1975, esp. pp. 256–7). For a brief history, a critical discussion, and a development of the theory to accommodate higher-order vagueness, see Williamson (1994); also Keefe (2000: ch. 7). For the relation between supervaluation theory and higher-order vagueness, see also Fine (1975, esp. §5).

The criticism of supervaluation theory given in the text, that it verifies the intuitively unacceptable "For some number n, an n-grained collection is a heap but an (n–1)-grained collection is not," has been attacked. Various writers have suggested that the quoted sentence is acceptable, as long as we realize that its truth does not require that there be a number n such that the following is true:

> An n-grained collection is a heap but an (n –1)-grained collection is not.

See, e.g. Dummett (1975, pp. 257–8); and, for opposition, Kamp (1981, pp. 237ff.). The claim depends upon two features: (1) a view of "there is" according to which it is like "or" (so that to say that there is a student who smokes is to say that either Sally smokes, or Michael smokes, or …, and so on through all the students); and (2) a view of "or" according to which a statement "*p* or *q*" can be (definitely) true even though neither "*p*" nor "*q*" is. A standard alleged example of the latter is "This is orange or red," said of a borderline case. See again Dummett (1975, p. 255). It is not at all clear whether the combination of (1) and (2) is less paradoxical than the paradox of the heap. It involves claiming that "there is a such-and-such" can be true, even if "that is a such-and-such" is false of each thing in the universe.

For a defense of supervaluationism from the objections fielded in the text see Keefe (2000).

For higher-order vagueness, see Wright (1992), the symposium between Hyde (1994) and Tye (1994b), and Williamson (1999).

Section 3.6

A classic source for degree theory is Zadeh (1965). I had mostly in mind the kind of theory advanced in Goguen (1969). For a version which does not treat the connectives as degree functional see Sanford (1975). For a more philosophical and less technical account of degree theory, see Peacocke (1981). For criticisms see Edgington (1992), Sanford (1976), and, on the resulting logic, Fine (1975), and Williamson (1994). For an account of validity and degrees according to which modus ponens remains valid, see Edgington (1992).

For the notion of partial belief, see Ramsey (1926) and Jeffrey (1965, chs. 3 and 4). A substantive question is whether a similar argument could be used to underwrite objective probabilities. If the answer is affirmative, as Mellor (1971) argues, then a question of crucial importance would be whether one can give a satisfactory account of why the arguments reach different destinations: degrees of truth in one case, objective probabilities in the other. Edgington (1992) suggests that less than full confidence induced by vagueness behaves differently from less than full confidence induced by lack of information. Recent discussions of some versions of degree theories can be found in Schiffer (2003) and MacFarlane (2007) (a hybrid view, combining degrees of truth with some aspects of epistemicism).

For a semantics with a vague metalanguage, see Tye (1994a).

Section 3.7

On vague objects, see Evans (1978), Lewis (1988), Copeland (1994), Salmon (1982, pp. 243ff.), Tye (1990), Wiggins (1986) and Williamson (2003).