

---

# Chapter XVIII: Fuzzy Logics in Theories of Vagueness

NICHOLAS J.J. SMITH

---

## 1 Introduction

Vagueness is centrally a property of *predicates*—for example ‘bald’, ‘tall’, ‘heavy’ and ‘interesting’. Predicates that are not vague are said to be *precise*—for example ‘under 1.4m in height’, ‘weighs at least 500 grams’ and ‘north of the equator’.<sup>1</sup> However we cannot properly delineate vagueness just by presenting examples—for the examples generally exhibit various other properties as well, in addition to vagueness. So to delineate our topic properly, we shall do two things: first (Section 1.1) give three positive identifying characteristics of vague predicates and second (Section 1.2) distinguish some other phenomena which often accompany vagueness in natural language but are not our central concern here.

The chapter then proceeds as follows. In Section 2 we discuss what a theory of vagueness should do and introduce the main non-fuzzy theories of vagueness. In Section 3 we introduce the central topic of this entry: fuzzy theories of vagueness. In Section 4 we consider some major arguments in favour of such theories and in Section 5 we present—and reply to—some major objections to fuzzy theories of vagueness.

### 1.1 Characteristics of vague predicates

Vague predicates are generally picked out as those possessing three characteristics: *blurry boundaries*, *borderline cases* and *sorites susceptibility*.

**Blurry boundaries:** The *extension* of a predicate is the set of things to which the predicate applies. The boundaries of the extensions of precise predicates are sharp, whereas the boundaries of the extensions of vague predicates are blurry. If you place a pin on a map, you can draw a sharp circle around the points that are within one kilometre of the point hit by the pin; you cannot draw a sharp circle around the points that are *near* it. If you take a large crowd of persons—say, an audience at a concert—you can sharply separate the seat numbers of persons whose height is at least six feet from the seat numbers of all other persons; not so for the seat numbers of the tall persons. In both

---

<sup>1</sup> Discussions of vagueness are not always limited to predicates: sometimes expressions other than predicates are held to be vague [106, §3.4.6] and [120, p. 124]—and sometimes objects, properties and other constituents of the world itself (as opposed to expressions in language) are held to be vague [95], [106, p. 158] and [4]. Furthermore, some treatments of vagueness in these other areas appeal to fuzzy logics [104]. Nevertheless, the bulk of discussions of vagueness is focussed on predicates—and predicates will be our focus here.

cases ('at least six feet in height' and 'tall') some persons are definitely in the extension and some are definitely not in it—but in the former case the line between them is sharp while in the latter case it is blurry.

**Borderline cases:** In general, predicates have clear positive cases—things of which we confidently assert that they fall under the predicate—and clear negative cases—things of which we confidently deny that they fall under the predicate. Vague predicates, unlike precise ones, also have (persistent) borderline cases: things of which we will neither confidently assert nor confidently deny that they fall under the predicate. Consider again the predicates 'at least six feet in height' and 'tall'. Simply on the basis of looking at them, you might not be able to separate a group of people into those who are at least six feet in height and those who are under six feet in height—but once you know everyone's height, you can do so (without remainder). The same cannot be said for 'tall': even once you know everyone's height, there will still be borderline cases.

**Sorites susceptibility:** A *sorites series* for a predicate  $F$  is a series of objects with the following characteristics:

1.  $F$  definitely applies to the first object in the series.
2.  $F$  definitely does not apply to the last object in the series.
3. Each object in the series<sup>2</sup> is extremely similar to the object after it in all respects relevant to the application of  $F$ .

For example, a series of men ranging in height from seven feet to four feet in increments of a thousandth of an inch is a sorites series for 'tall' and for 'at least six feet in height'; a series of points one millimetre apart on a straight line from a point  $p$  to a point  $q$  one thousand kilometres away is a sorites series for 'far from  $q$ ' and for 'more than one kilometre from  $q$ '; and so on. Given a sorites series, we can generate an associated *sorites argument*:<sup>3</sup>

1.  $x_1$  is  $F$
2. For every  $x$ , if  $x$  is  $F$  then  $x'$  is  $F$ .  
(Or: There is no  $x$  such that  $x$  is  $F$  and  $x'$  isn't  $F$ . Or all of the following: If  $x_1$  is  $F$  then  $x_2$  is  $F$ , ..., If  $x_{n-1}$  is  $F$  then  $x_n$  is  $F$ . Or all of the following: It is not the case both that  $x_1$  is  $F$  and  $x_2$  isn't  $F$ , ..., It is not the case both that  $x_{n-1}$  is  $F$  and  $x_n$  isn't  $F$ .)
3.  $\therefore x_n$  is  $F$ .

Now, the difference between vague and precise predicates is as follows. Where  $F$  is a precise predicate, the sorites argument will be obviously mistaken: there is indeed some  $x$  such that  $x$  is  $F$  and  $x'$  is not  $F$ .<sup>4</sup> Where  $F$  is a vague predicate, on the other hand, the sorites argument will seem genuinely paradoxical: the premisses will all seem to be true, the reasoning will seem to be correct, and yet the conclusion will seem to be false. Thus vague predicates, unlike precise ones, generate sorites paradoxes.

<sup>2</sup> Except the last one, which has no object after it.

<sup>3</sup>  $x_1, \dots, x_n$  denote the objects in the series from first ( $x_1$ ) through to last ( $x_n$ ).  $x$  ranges over all objects in the sorites series except the last object and  $x'$  denotes the object immediately after  $x$  in the series.

<sup>4</sup> We don't need to know which object it is: the point is just that there is such an object.

## 1.2 Other phenomena

In natural language, vagueness tends to come packaged with other phenomena. So as to bring vagueness into clear relief, we distinguish four of these other phenomena now: *uncertainty*, *context sensitivity*, *ambiguity* and *generality*.

**Uncertainty:** Given an object  $x$  and a predicate  $P$ , a speaker may be uncertain whether or not  $x$  is in the extension of  $P$ . This is an epistemic phenomenon: the speaker lacks information; she does not know whether or not  $x$  is  $P$ . In this case,  $x$  is a kind of borderline case—and a predicate  $P$  may even have persistent borderline cases of this kind. For example, suppose we define the predicate ‘bearfast’ to apply to all and only objects that are moving faster than any polar bear moved on 11th January 1904. We know that some objects are bearfast (e.g. the jet plane flying overhead) and that some are not (e.g. the parked car across the street) but there are many things of which we will neither confidently assert nor confidently deny that they are bearfast—and it seems that we will never be able to gain the information required to classify these cases one way or the other. Such uncertainty also brings with it a kind of blurry boundary. In itself, of course, the boundary of the extension of ‘bearfast’ is perfectly sharp: but it is shrouded by an epistemic blur. Nevertheless, uncertainty is distinct from vagueness: ‘bearfast’ is not a vague predicate.<sup>5</sup> In particular, it does not generate sorites paradoxes. We can set up a sorites series for ‘bearfast’, beginning with an object moving at great speed and progressing by tiny increments to a stationary object—but the associated sorites argument will be obviously mistaken: there is indeed some object  $x$  in the series such that  $x$  is bearfast and  $x'$  is not. Of course we do not—cannot—know which object it is: but this will not make us think that the second premise of the sorites argument is *true*.<sup>6</sup>

**Context sensitivity:** A predicate can have different extensions in different contexts. For example, in a discussion of professional jockeys, a certain individual might count as ‘tall’, while in a discussion of professional basketball players, the same individual might not count as ‘tall’. Many standard examples of vague predicates—‘tall’, ‘loud’, ‘heavy’, ‘small’ etc.—are also context sensitive. Nevertheless, vagueness and context sensitivity are distinct phenomena—for vagueness arises even on a single occasion of use. Consider a vague predicate—say ‘tall’—as uttered on some occasion, in some particular context. Let the predicate  $P$  have as its extension the extension that ‘tall’ had on that single occasion of use. Then  $P$  has blurry boundaries and borderline cases—and it generates a sorites paradox.<sup>7</sup>

<sup>5</sup> It is uncontroversial that the existence of an object and a speaker, such that the speaker is uncertain whether the object is in the extension of  $P$ , does not render  $P$  vague. The aim of the present discussion of ‘bearfast’ is to establish a stronger claim: the existence of a whole *class* of objects such that *no* speaker *can* know whether these objects are in the extension of  $P$  does not in and of itself render  $P$  vague. The aim is not to rule out in advance a view that we shall encounter in Section 2.2: epistemicism about vagueness, according to which vagueness essentially involves a certain kind of ignorance or uncertainty.

<sup>6</sup> Recall that on some formulations of the sorites argument there is a single second premise that makes a general claim, while on other formulations this single premise is replaced by a multitude of premisses each of which makes a particular claim. For simplicity, I shall generally write of ‘the second premise’. Relative to formulations involving a multitude of premisses after the first one, talk of ‘the second premise’ being true (false) should be interpreted as talk of all (some) of them being true (false); talk of ‘the second premise’ being accepted should be interpreted as talk of all of them being accepted; and so on.

<sup>7</sup> The example involving ‘tall’ given in the text above illustrates a simple and familiar kind of context sensitivity: relativity to a comparison class. It is uncontroversial that vagueness and this kind of context

**Ambiguity:** An ambiguous expression is one that is susceptible of multiple interpretations that give rise to different truth conditions. A classic example is ‘bank’, which can mean (among other things) the edge of a river or a financial institution—and the truth conditions of ‘John went to the bank’ depend on which interpretation is given. Some vague expressions are also ambiguous—for example, ‘heavy’ can mean ‘of great weight’ or ‘very important or serious’—but we can see that vagueness is distinct from ambiguity by noting that vagueness can still be present after disambiguation. For example, even once we specify that we are using it in the sense of ‘of great weight’, ‘heavy’ is still vague: it has blurry boundaries and borderline cases and it generates a sorites paradox.

**Generality:** Colloquially, someone might describe a person as giving a vague response if, when asked his age, he replies ‘I was born last century’. This reply—in contrast to ‘I was born in 1926’—leaves open many possibilities. However ‘born last century’ is not vague in our sense:<sup>8</sup> for example, it does not generate a sorites paradox (it is obviously not true that for any persons born, say, one minute apart, if one was born last century then so was the other).

## 2 Non-fuzzy theories of vagueness

We shall gain a better understanding of theories of vagueness that employ fuzzy logics if we are familiar with the main alternatives. In this section, therefore, we look at prominent theories of vagueness that do not employ fuzzy logic.<sup>9</sup> First, in Section 2.1, we make some brief general remarks about what a theory of vagueness (fuzzy or non-fuzzy) is supposed to do.

### 2.1 Theory of vagueness

One fundamental question that a theory of vagueness should answer is: What is the meaning (semantic value) of a vague predicate? Another is: How should we reason in the presence of vagueness? Thus a theory of vagueness needs to cover two bases: semantics and logic. If (as is standard) we adopt a truth conditional approach to semantics and hold that correct reasoning must be truth preserving then the two goals are closely connected.

A more specific requirement on any theory of vagueness is that it solve the sorites paradox. There are two aspects to such a solution. One task is to locate the error in the sorites argument: the premise that isn’t true or the step of reasoning that is incorrect. This part of the solution should fall out of what a theory of vagueness says in general about semantics and logic. There is also a further task, that goes beyond pure logic and semantics. A satisfying solution to the sorites should explain why it is a paradox rather than a simple mistake. Thus, as well as locating the error in the argument, a

---

sensitivity are distinct phenomena (because ‘tall’ is still vague, even when we fix the comparison class—e.g. to NBA players in 2014). In Section 2.5 we shall encounter contextualism about vagueness: a view according to which vagueness essentially involves some other form of context sensitivity. Whether or not the same kind of argument also defeats contextualism about vagueness (i.e. vagueness cannot be any kind of context sensitivity because vagueness remains even once the entire context—not just a comparison class—is fixed) is a topic of debate; see e.g. [3].

<sup>8</sup> At least not if we suppose, for the sake of argument, that each person is born at a particular instant.

<sup>9</sup> As our main topic here is fuzzy theories of vagueness, we give only a brief presentation of these other views; furthermore, the choice of aspects to focus on in this presentation is partly guided by considerations of what needs to be on the table in order best to understand fuzzy theories when we get to them later. For a much more detailed presentation of these views see [106, ch. 2].

theory of vagueness must provide an explanation of why competent speakers find the argument compelling but not convincing: why they do not spot the error immediately and yet—even in the absence of a clear idea of what the error is—are not inclined to accept the conclusion.

## 2.2 Epistemicism

On this approach, the semantics and logic of vagueness are both entirely classical. Vague predicates—like precise ones—have crisp sets as their extensions. Classical reasoning is correct in the presence of vagueness just as it is in the precise realm of mathematics. This is not to say that there is no distinction at all between vague predicates and precise ones: while they are the same from the semantic and logical points of view, there is an epistemological difference between them. Although the extensions of all predicates are crisp sets, with vague predicates we cannot *know* where the borders lie. Thus, according to the epistemicist, the blurriness of the boundaries of vague predicates is of an epistemic sort: in themselves the boundaries are perfectly sharp but they are hidden behind a veil of ignorance. For all objects  $x$ ,  $Px$  is true or false; the borderline cases are the objects  $x$  for which we cannot *know* whether  $Px$  is true or false.

The first part of the epistemicist solution to the sorites paradox—saying what is wrong with the argument—is straightforward. According to the epistemicist, there is a sharp cut-off in the sorites series between the last object that is  $P$  and the first object that isn't—so the second premise is false. The second part of the solution—saying why we are taken in—is trickier. According to the epistemicist, we cannot know where the cut-off is—and so we mistakenly think that there is no such cut-off. This is why we are inclined to accept the second premise even though it is in fact false.

Note that the second part of the solution involves a departure from the usual modus operandi in formal semantics, in which the semantic theory one develops is taken to be implicit in the ordinary usage of competent speakers [101]. In the case of the epistemic theory of vagueness, the explanation of ordinary competent speakers' reactions to the sorites argument turns on their being fundamentally mistaken about the semantics of the predicates they are using. For suppose that a speaker did realise that a predicate  $P$ —say, 'bearfast'—had sharp but unknowable boundaries. Then she would not think for a moment that the second premise of a sorites argument for  $P$  was actually *true*—even though she would indeed be unable to mark the cut-off point in the sorites series between the  $P$ 's and the non- $P$ 's.

A major problem for epistemicists is what I have elsewhere dubbed the *location problem* [106]. According to the epistemicist, in a sorites series for  $P$ , there is a last object  $x$  that is  $P$ —and the very next object  $x'$  in the series is not  $P$ . More generally, the epistemicist thinks that the boundaries of the extensions of vague predicates are (in themselves) perfectly sharp (even if they are hidden beneath a blur of ignorance). The problem is to explain how these boundaries get to be exactly where they are. Why is  $x$  the last object in the series that is  $P$ —why not some other object that is similar to  $x$ ? What is it that makes our vague term 'tall' (say) have *this* crisp set  $S$  as its extension rather than some other crisp set  $S'$  (which might differ from  $S$  in only very small ways—for example  $S$  includes Bill and excludes Ben, who is a nanometre shorter than Bill, while  $S'$  includes them both)?

It is generally accepted that language is a human artefact. The sounds we make mean what they do because of the kinds of situations in which we, and earlier speakers, have made those sounds (e.g. had we always used the word ‘dog’ where we in fact used ‘cat’ and vice versa, then ‘dog’ would have meant what ‘cat’ in fact means and vice versa). Hence there should be some sort of connection between *meaning* and *use*. More precisely, consider the following kinds of facts:

- All the facts as to what speakers actually say and write, including the circumstances in which these things are said and written, and any causal relations obtaining between speakers and their environments.
- All the facts as to what speakers are disposed to say and write in all kinds of possible circumstances.
- All the facts concerning the eligibility as referents of objects and sets.

There is widespread (not universal) agreement in the literature that if these facts are insufficient to determine (unique) meanings for some utterances, then those utterances have no (unique) meanings. That is: semantic facts are never primitive or brute—they are always determined by the meaning-determining facts; and the meaning-determining facts are the ones just itemised. This generates a constraint on any theory of vagueness: if the theory says that vague predicates have meanings of such-and-such a kind (e.g. crisp sets), then we must be able to satisfy ourselves that the meaning-determining facts itemised above could indeed determine such meanings for actual vague predicates. To the extent that the meaning-determining facts do *not* appear sufficient to determine meanings for vague predicates of the kind posited by some theory of vagueness, that theory is undermined. This is precisely where the epistemicist theory faces a problem.

To see this, let’s start with the following proposal concerning the connection between meaning and use (for the case of predicates):

(MU) The claim  $Pa$  is true if and only if most competent speakers would confidently assent if presented with  $a$  in normal conditions and asked whether it was  $P$ , and is false if and only if most competent speakers would confidently dissent if presented with  $a$  in normal conditions and asked whether it was  $P$ .

There are obvious counterexamples to (MU), for example natural kind terms, and technical terms subject to the division of linguistic labour [87]. However the principle is *prima facie* plausible for common, everyday, non-specialist vocabulary.

The epistemicist must reject (MU): for any borderline case  $x$  of baldness (say), ‘ $x$  is bald’ is either true, or false (on the epistemicist view)—yet most competent speakers would neither assent to nor dissent from ‘ $x$  is bald’ (we hedge over such cases). So within the borderline cases there is a failure of match-up between meaning and use. We can picture the situation—say for the predicate ‘is tall’—as in Figure 1. Epistemicism entails a mismatch between use (on the left) and meaning (on the right).

Of course (MU) is just one proposal connecting meaning and use and denying it does not necessarily mean denying the general principle that use determines meaning. (MU) tells us that in Figure 1 the left hand side and the right hand side must *match*. The

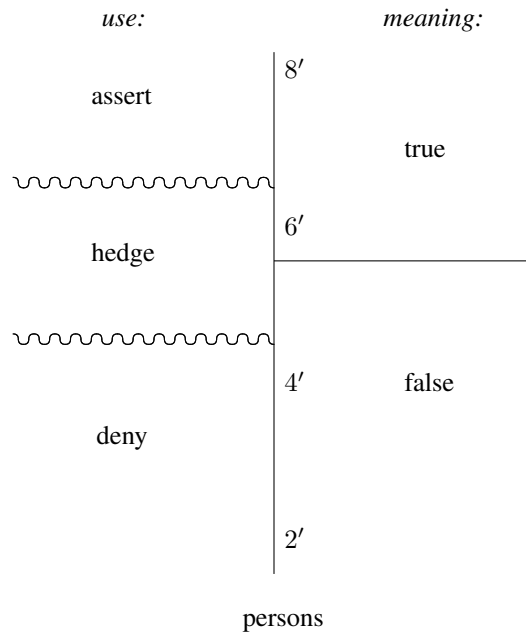


Figure 1. Use and meaning

claim that use determines meaning—that meaning *supervenes* on use, that there can be no difference in meaning without a difference in use—is the claim that the left hand side *determines* the right hand side: that if the right hand side of the picture were different, the left hand side would be different too. It is *consistent* to claim that the left hand side and the right hand side do not match *and* that the right hand side could not be different without the left hand side being different. This is the sort of line taken by Williamson [122, §4] and [123, §7.5]. The idea is that use does determine meaning—just not in any straightforward way. We have no idea of the mechanisms intervening between use and meaning: we do not know how (if at all) the right hand side of our picture would alter if we altered this or that bit of usage on the left hand side. Yet Williamson can still maintain that the right hand side could not be different unless the left hand side was too.

This is a consistent position but it does not provide a real response to the location problem—the problem of *how* precise extensions for vague predicates are determined: it leaves the connection between use and meaning mysterious. Hence it is by nature unsatisfying: we are told merely *that* meaning is determined by use when we want to know *how* this could possibly be the case (assuming—as the epistemicist thinks—that meaning is cleanly bipartite, and given that usage is fuzzily tripartite). In the absence of any explanation of how usage determines sharp boundaries for vague predicates, the most reasonable conclusion seems to be that the meaning-determining facts do *not* suffice to pick out meanings for vague predicates of the kind the epistemicist says they have.

### 2.3 Gaps and third values

In contrast to epistemicism, there are views according to which vague predicates have alethic borderline cases: objects  $x$  such that  $Px$  is neither true nor false (as opposed to being one or the other, we know not which). One way of spelling out this idea is to say that the extension of a predicate is a partial set: a *partial* function from objects to the classical truth values 1 and 0. Objects sent to neither value are the borderline cases. Where  $x$  is a borderline case of  $P$ ,  $Px$  will then lack a truth value. Another way is to posit a third truth value in addition to the two classical truth values and say that the extension of a predicate is a function from objects to the set of three truth values. Objects sent to the new third value are the borderline cases. Where  $x$  is a borderline case of  $P$ ,  $Px$  will then have the third truth value. For ease of exposition, I shall use the term ‘tripartite’ to cover both partial two-valued and three-valued approaches.

Tripartite approaches can accept (MU) and hence respond to the location problem in a straightforward way. The things sent to 1 by the extension of  $P$  are the things that competent speakers would confidently classify as  $P$ ; the things sent to 0 by the extension of  $P$  are the things that competent speakers would confidently classify as non- $P$ ; and the rest are the cases over which speakers would hedge. This is progress, admittedly—but of course a residual problem remains: usage is *fuzzily* tripartite (there is no precise boundary between assertion and hedging, nor between denial and hedging) while meaning (on tripartite views) is *sharply* tripartite.

Given tripartite extensions for predicates, the question now arises how to handle compound statements—conjunctions, negations and so on—where some components have a value other than 1 or 0.<sup>10</sup> One option is the recursive route, where we extend the two-valued truth tables of classical logic to cover the cases where one or both components lack a classical value. There are many options here; some prominent ones are shown in Figures 2(a), 2(b) and 2(c) (in which the  $*$  can be interpreted as a third truth value—on three-valued views—or no value—on partial two-valued views).<sup>11</sup>

Another option is to proceed via classical sharpenings.<sup>12</sup> Begin with a three-valued valuation  $V_3$  (i.e. a mapping from propositional constants to the set of three truth values). Now instead of extending this to a three-valued model (i.e. an assignment of truth values to all well formed formulas) using truth tables, say that a classical valuation  $V_2$  (i.e. a mapping from propositional constants to the set of two classical truth values) *extends*  $V_3$  iff whenever  $V_3$  assigns 1 (respectively, 0) to a proposition,  $V_2$  does too.<sup>13</sup> Consider all the classical valuations that extend  $V_3$ . Each one of them determines (via the classical truth tables) a classical model  $\mathfrak{M}_2$ ; call these models *extensions* of  $V_3$ . The rule for extending  $V_3$  to a three-valued model  $\mathfrak{M}_3$  (which we shall call an s-valuational model) is now as follows (where  $[\alpha]$  is the truth value of  $\alpha$ ):

- $[\alpha] = 1$  (respectively, 0) on  $\mathfrak{M}_3$  iff  $[\alpha] = 1$  (resp., 0) on all classical extensions.
- $[\alpha] = *$  on  $\mathfrak{M}_3$  otherwise (i.e. iff  $[\alpha] = 1$  on some classical extension and  $[\alpha] = 0$  on some classical extension).

<sup>10</sup> For simplicity I focus on connectives here; similar remarks apply to quantifiers.

<sup>11</sup> For further details on these logics see [108].

<sup>12</sup> For simplicity, I shall present the main idea using three truth values in the context of propositional logic, but the approach extends readily to first order logics and can also be formulated using two values and gaps.

<sup>13</sup> So  $V_2$  differs from  $V_3$  only where  $V_3$  assigns  $*$  to some basic proposition  $\alpha$ .



$\alpha$	$\beta$	$\neg\alpha$	$\alpha \wedge \beta$	$\alpha \vee \beta$	$\alpha \rightarrow \beta$	$\alpha \leftrightarrow \beta$	$\alpha$	$\beta$	$\neg\alpha$	$\alpha \wedge \beta$	$\alpha \vee \beta$	$\alpha \rightarrow \beta$	$\alpha \leftrightarrow \beta$
1	1	0	1	1	1	1	1	1	0	1	1	1	1
1	*		*	*	*	*	1	*		*	1	*	*
1	0		0	1	0	0	1	0		0	1	0	0
*	1	*	*	*	*	*	*	1	*	*	1	1	*
*	*		*	*	*	*	*	*		*	*	*	*
*	0		*	*	*	*	*	0		0	*	*	*
0	1	1	0	1	1	0	0	1	1	0	1	1	0
0	*		*	*	*	*	0	*		0	*	1	*
0	0		0	0	1	1	0	0		0	0	1	1

(a) Bočvar (aka Kleene weak) truth tables

(b) Kleene (strong) truth tables

$\alpha$	$\beta$	$\neg\alpha$	$\alpha \wedge \beta$	$\alpha \vee \beta$	$\alpha \rightarrow \beta$	$\alpha \leftrightarrow \beta$
1	1	0	1	1	1	1
1	*		*	1	*	*
1	0		0	1	0	0
*	1	*	*	1	1	*
*	*		*	*	1	1
*	0		0	*	*	*
0	1	1	0	1	1	0
0	*		0	*	1	*
0	0		0	0	1	1

(c) Łukasiewicz truth tables

Figure 2. Prominent tripartite truth tables

We can get a logic from this framework by picking a set of designated values: a wff is valid iff there is no *s*-valuational model on which it has a non-designated value and an argument is valid iff there is no *s*-valuational model on which the premisses all have designated values and the conclusion has a non-designated value. If we set 1 as the only designated value we get *supervaluationist* logic; if we set 1 and \* as designated we get *subvaluationist* logic.<sup>14</sup> Note that a formula gets a designated value on a model  $\mathfrak{M}_3$  (based on a valuation  $V_3$ ) in the supervaluationist framework if it is true on every classical model that extends  $V_3$ ; such a formula is sometimes called ‘supertrue’. In the subvaluationist framework, a formula gets a designated value on a model  $\mathfrak{M}_3$  (based on a valuation  $V_3$ ) if it is true on at least one classical model that extends  $V_3$ .

In the context of vagueness, it is standard to add a further refinement to this general framework. The base valuation—which is three-valued or partial—is taken to correspond to ordinary usage of vague language: clear cases/noncases of a predicate *P* are mapped to 1/0 by the extension of the predicate in the base model and the borderline cases are neither mapped to 1 nor to 0. A classical extension of the base model then represents a way of *sharpening* or *precisifying* all vague predicates: classifying their borderline cases as positive or negative cases. (Think for example of the way ‘adult’ might be sharpened in a certain legal context so that it applies to all and only the persons over 18.) Instead of considering *all* such extensions, we consider only the *admissible* ones: the ones that represent acceptable ways of sharpening all vague predicates at once.

<sup>14</sup> This unified presentation of supervaluationism and subvaluationism in terms of *s*-valuational models is inspired by [93].

So, for example, suppose that Bill and Ben are borderline tall, and Bill is slightly taller than Ben. Considering just 'Bill' and the predicate 'tall' it might be acceptable to put him in or out of the extension; likewise for the predicate 'short'. The same might hold for Ben. Yet it is not acceptable to put Bill in *both* the extensions of 'tall' and 'short' and it is not acceptable to put 'Bill' in the extension of 'short' and 'Ben' in the extension of 'tall'. In other words, borderline cases cannot always be classified independently when we are precisifying a predicate—and different predicates cannot always be precisified independently of one another.

The first part of the supervaluationist solution to the sorites paradox—saying what is wrong with the argument—is that the second premise is false no matter how we precisify and hence comes out as false (in the tripartite base model). The second part of the solution—saying why we are taken in—is as follows. Where  $x$  or  $x'$  is a borderline case of  $P$ , each statement of the form ' $x$  is  $P$  and  $x'$  is not  $P$ ' is true on one admissible extension and false on the others—and hence comes out neither true nor false (in the tripartite base model). So we cannot truly say of any object in the series that it is the last  $P$ . From here—the story goes—we (mistakenly) conclude that the second premise of the sorites argument is true.

Note that the second part of the solution involves the same sort of departure from the usual modus operandi in formal semantics that we saw in the case of epistemicism: the explanation of ordinary competent speakers' reactions to the sorites argument turns on their being fundamentally mistaken about the semantics of the predicates they are using [101]. If a speaker thought that her language works in the way the supervaluationist says it does then she would have no tendency to move from the non-truth of ' $x$  is the last  $P$ ' to the truth of 'there is no last  $P$ '. Consider an analogous case. We are rolling a die. We know that we can only say 'it is certainly the case that  $\Phi$ ' if  $\Phi$  will be true no matter how the die falls. So we cannot say 'it is certainly the case that we will roll 1'—or 2, 3, 4, 5 or 6. Yet we have no tendency to infer from this that 'we will roll one of the numbers 1 through 6' is false. On the contrary, it is certainly true: for it will hold however the die lands. According to the supervaluationist, however, this is precisely the kind of mistake ordinary speakers make in relation to the sorites paradox.

Moving from sorites susceptibility to borderline cases and blurry boundaries: unlike epistemicism, supervaluationism holds that vague predicates have genuinely alethic borderline cases—but it does not capture the idea that vague predicates draw blurry boundaries. On the contrary, the divisions between the clear cases and the borderline cases and between the clear noncases and the borderline cases are perfectly sharp. Supervaluationists have explored various approaches to this problem. One that we should mention here—because it bears a superficial similarity to fuzzy approaches—is the degree-theoretic form of supervaluationism. The machinery of a tripartite base model together with its admissible classical extensions is augmented with a normalised measure over the classical extensions. Formulas are then assigned not one of three truth values in the base model—but one of the infinitely many values in  $[0, 1]$ , according to the rule that the value of  $\alpha$  in the base model is the measure of the set of extensions in which  $\alpha$  is true. We can think of what is going on here as follows [110]. The values assigned to formulas in the base model are *probabilities* of full truth under complete precisification of the language. The original supervaluationist focusses on formulas that have 100% probability of truth

and the subvaluationist focusses on formulas that have non-zero probability of truth; the new degree-theoretic form of supervaluationism, on the other hand, is interested in all intermediate probabilities as well. Earlier we saw that, according to the epistemicist, the blurriness of the boundaries of vague predicates is of an epistemic sort: in themselves the boundaries are perfectly sharp but they are hidden behind a veil of ignorance. In degree supervaluationism, the blurriness of the boundaries of vague predicates is probabilistic: in themselves (as things stand in the base model, which represents language in its actual vague state) the boundaries are sharply tripartite (the extension of a predicate in the base model is a three-valued function or a partial two-valued function)—but the probability that a given object would end up in the extension *were* the language to be fully precisified can in general be anywhere between 0 and 100%.

#### 2.4 Plurivaluationism

In order to understand plurivaluationism—and how it differs from both epistemicism and supervaluationism—we need to distinguish pure model theory and model-theoretic semantics (MTS). MTS requires an additional notion that does not figure in pure model theory: its role is to distinguish one (or some) of the infinity of models of a given formal language countenanced in pure model theory as the one(s) relevant to questions of the (actual) meaning and truth (simpliciter) of utterances in some discourse. Such questions are of central interest in natural language semantics—but pure model theory cannot (fully) answer them: for it tells us only that a formula is true on this model and false on that one (etc.). If we want to know whether a given statement is true (simpliciter), then we need to single out a particular model (or a class of models): truth simpliciter will then be truth relative to this model. So we get a system of MTS by combining a system of pure model theory with some notion that plays the role of distinguishing some model(s) as the ones relevant to questions of (actual) meaning and truth (simpliciter) of utterances in some discourse. The simplest choice of model theory is classical model theory. The simplest choice of auxiliary notion is the idea that for each discourse, there is a *unique* relevant model: the ‘intended model’. Combining these two choices yields the ‘classical semantic picture’ [106, §1.2]. It is the version of MTS that underlies epistemicism.

Supervaluationism differs from epistemicism by retaining the idea of a unique intended model of vague discourse but rejecting the idea that it should be a classical model: instead it should be an s-valuational model. Plurivaluationism on the other hand countenances only classical models—but instead of holding that a vague discourse is associated with a single intended model it holds that it is associated with multiple equally acceptable models. These models are precisely the classical models that figure as admissible extensions of the (uniquely intended) tripartite base model in the supervaluationist picture. This can lead to plurivaluationism and supervaluationism being confused or conflated—but they are really quite different. On the supervaluationist picture, there is just one model that is relevant to questions of the (actual) meaning and truth (simpliciter) of utterances in a given vague discourse, and it is nonclassical. On the plurivaluationist picture, there are many models that are all equally relevant to questions of the (actual) meaning and truth (simpliciter) of utterances in a given vague discourse, and each of them is classical.

One natural motivation for plurivaluationism would be a fondness for classical models coupled with a conviction that epistemicism cannot adequately answer the location problem: the epistemicist cannot explain how usage singles out one model as the unique intended one (the model  $\mathfrak{M}$  such that the unique sharp extension of ‘tall’ that the epistemicist countenances is its extension on  $\mathfrak{M}$ )—and so we should not believe that there is a single intended model. Rather, there are many acceptable models: all the ones that are not ruled out as incorrect by the meaning-determining facts.

As the plurivaluationist lacks the idea of a unique intended model of vague discourse, he cannot follow the route of identifying truth simpliciter for statements in that discourse with truth on the intended model. When it comes to truth, he can only say, of each acceptable model, whether a statement is true or false on that model. As there is no further semantic machinery beyond the individual classical models, there can be no further overarching semantic fact, such as that the statement is ‘true simpliciter’. However, the plurivaluationist can say something about assertability that is analogous to what the supervaluationist says about truth. For the supervaluationist, truth simpliciter is truth on the unique intended tripartite base model—and (in the s-valuationist semantics) this means truth on all admissible extensions of the base model. The plurivaluationist can say that when a statement is true on all acceptable models, we can simply assert it; when a statement is false on all acceptable models, we can simply deny it; and when a statement is true relative to some acceptable models and false relative to others, we can neither simply assert it nor simply deny it.

Here’s a useful way of thinking about the plurivaluationist view. When I utter ‘Bob is tall’ (say), I say many things at once: one claim for each acceptable model. Thus we have semantic indeterminacy—or equally, semantic plurality. However, if all the claims I make are true (or false) then we can talk as if I make only one claim, which is true (or false). Figuratively, think of a shotgun fired (once) at a target: many pellets are expelled, not just one bullet; but if all the pellets go through the bullseye, then we can harmlessly *talk as if* there was just one bullet, which went through.

This approach will lead to a solution to the sorites that is analogous to the supervaluationist solution: the second premise is false on every acceptable model (recall that each such model is classical) but speakers fail to see this because where  $x$  or  $x'$  is a borderline case of  $P$ , each statement of the form ‘ $x$  is  $P$  and  $x'$  is not  $P$ ’ is true on one acceptable model and false on the others—and hence cannot be asserted. Again, this involves a departure from the usual modus operandi in formal semantics: the explanation of ordinary competent speakers’ reactions to the sorites turns on their not realising that the semantics of their discourse is plurivaluationist [101]. Consider an analogous case. We are canvassing opinions about football and (suppose for the sake of argument) there is a convention that one can say ‘the man in the street believes  $\Phi$ ’ only when *everyone* canvassed believes  $\Phi$ . Now suppose that each person canvassed has a favourite team but there is no single team that is everyone’s favourite. So we cannot say ‘the man on the street’s favourite team is  $X$ ’ (or  $Y$  or  $Z$  etc., through all the teams). Yet we have no tendency to infer from this that we cannot say ‘the man on the street has a favourite team’. On the contrary, this is clearly something that we can assert (in the imagined circumstances) and it may well convey important information (e.g. it excludes the possibilities that some people just don’t care about football or have several equally favoured teams).

Turning to borderline cases and blurry boundaries: like epistemicism and unlike supervaluationism, plurivaluationism gives a non-alethic account of borderline cases. On each acceptable model, ' $x$  is  $P$ ' is true or false. There is no further machinery beyond the classical models—and so there is no level at which ' $x$  is  $P$ ' lacks one of the two classical truth values. It is just that when it has different truth values on different acceptable models, we can neither simply assert nor simply deny that  $x$  is  $P$ . As for blurry boundaries: like supervaluationism, plurivaluationism makes no room for them; the divisions between the clear cases and the borderline cases and between the clear noncases and the borderline cases are perfectly sharp.

## 2.5 Contextualism

Combining a recursive tripartite approach with some additional machinery yields contextualist theories of vagueness, which have played a significant role in the recent literature. Let's suppose that at any particular time, a vague discourse has a unique intended model, which (in the most prominent versions of contextualism) is tripartite and recursive. In supervaluationism, the idea of a complete precisification of the language plays a central role. In contextualism, a more local and less idealised form of precisification plays a central role: the sort of precisification where we partially precisify a predicate, by classifying one of its borderline cases as a positive or negative case. For example, we might partially precisify 'tall' by deeming that Bob—a borderline case—is to count as tall. We can understand what is going on here as a change in the intended model: the intended model  $\mathfrak{M}$  of the discourse at some time  $t$  is one in which Bob is sent neither to 1 nor to 0 by the extension of 'tall'; the intended model  $\mathfrak{M}'$  of the discourse at some slightly later time  $t'$ —where this extended discourse now includes the act of stipulating that Bob is tall—is one in which Bob is sent to 1 by the extension of 'tall'. According to contextualists, this will typically not be the only difference between  $\mathfrak{M}$  and  $\mathfrak{M}'$ : other persons who are similar in height to Bob and who were also borderline cases of 'tall' in  $\mathfrak{M}$  will be positive cases in  $\mathfrak{M}'$ . Exactly what prompts the change from  $\mathfrak{M}$  to  $\mathfrak{M}'$  and exactly how  $\mathfrak{M}'$  differs from  $\mathfrak{M}$  are matters over which different contextualists differ—but one prominent way to change the intended model is to stipulate that a (hitherto) borderline case is  $P$  or that it is not  $P$ . Note that even after such an act of precisification the intended model will still be tripartite: it will generally not be a classical model of the sort that supervaluationists take as extensions of their tripartite base models. Note also that because truth simpliciter is truth on the intended model, 'Bob is tall' will be neither true nor false as uttered at  $t$  and true as uttered at  $t'$ : at  $t'$ , this sentence is still neither true nor false on model  $\mathfrak{M}$  (the model itself does not change) but model  $\mathfrak{M}$  is no longer the intended model of the discourse;  $\mathfrak{M}'$  is now the intended model—the one that is relevant to questions of the (actual) meaning and truth (simpliciter) of utterances in the discourse.

Contextualists think that this combination of a tripartite approach within each model with a dynamic story about how acts of stipulation change which model is intended yields a satisfying theory of vagueness. In particular, the solution to the sorites will be along the following lines. The second premise fails to be true on the intended model (at any stage of the discourse). We nevertheless find it plausible because first, it is not false, and second, if we suppose that some object  $x$  in the sorites series is  $P$ , we thereby affect which model is intended in such a way that 'it is not the case that ( $x$  is  $P$  and  $x'$  is not  $P$ )'

and ‘if  $x$  is  $P$  then  $x'$  is  $P$ ’ become true (even if they were not so beforehand). Again, however, note that this solution to the sorites turns on speakers not believing that the contextualist story is the correct account of vague language [101]. A speaker who believed that classifying a borderline case  $x$  as  $P$  or as not  $P$  changes the intended model in such a way as to render these classifications—and analogous statements about objects similar to  $x$ —true would still have no reason to think that the second premise of the sorites argument was true (relative to any model that might be the intended one at any point in time).

Contextualism gives an alethic account of borderline cases. As for blurry boundaries, the contextualist will say that the kind of blurriness involved is a matter of rapid shifting. At any point in time, the extension of a vague predicate is a sharply tripartite three-valued or partial two-valued set—but *which* such set is the correct extension tends to shift as a conversation proceeds. In particular, attempts to locate the boundary between the  $P$ ’s and the borderline cases—by saying that  $x$  is  $P$  and  $x'$  is not—cause the boundary to move elsewhere: the act of deeming  $x$  to be  $P$  triggers a shift to a model in which  $x$  and things sufficiently similar to it—which includes  $x'$ —are in the extension of  $P$ .

### 3 Fuzzy theories of vagueness

We noted earlier that a system of MTS has two components: the models themselves (of some particular sort) and some notion like that of the ‘intended model’ that plays the role of distinguishing one (or some) of the infinity of models as the one(s) relevant to questions of the (actual) meaning and truth (simpliciter) of utterances in some discourse. We’ll touch on the second component in Section 3.3 and then come back to it in more detail later. For now, let’s look at the models. In fuzzy theories of vagueness, the kind of models involved will be fuzzy models. Here—as in any kind of model—we can distinguish two components: a part that deals with the nonlogical vocabulary and a part that treats the logical vocabulary. We shall discuss these in turn—under the headings ‘degrees of truth’ (Section 3.1) and ‘truth rules’ (Section 3.2) respectively. For each component, there are many choices available: so there are many kinds of fuzzy model.

#### 3.1 Degrees of truth

It will be helpful to start by considering classical models—say of a standard first order language, which has the usual logical vocabulary (connectives, quantifiers, variables) and the following nonlogical vocabulary:

- Individual constants:  $a, b, c, \dots$
- $n$ -place predicates (for each  $n > 0$ ):  $P^n, Q^n, R^n, \dots$

A classical model has some machinery dealing with the nonlogical vocabulary. Each primitive nonlogical symbol is assigned a value: for an individual constant  $a$  its value  $[a]$  is an object in the domain of the model; for an  $n$ -place predicate  $P$  its value  $[P]$  is a function from the set of  $n$ -tuples of members of the domain to the set  $\{0, 1\}$  of classical truth values (such a function can be thought of as a *set* of  $n$ -tuples: the  $n$ -tuples sent to 1 are in the set and the  $n$ -tuples sent to 0 are not in the set). This machinery suffices to

determine a truth value for each closed wff that involves no logical vocabulary, via the following principle:

$$[P^n a_1 \dots a_n] = [P^n]([a_1], \dots, [a_n]).$$

That is, the truth value of  $P^n a_1 \dots a_n$  is the value (0 or 1) to which the value of the predicate sends the  $n$ -tuple containing the values of the names (in the order in which they appear after the predicate).

There is more to classical models: there is the part that determines the truth value on a model of a closed formula that contains logical vocabulary. But we shall return to this below. Already, at the stage of treating the nonlogical vocabulary, fuzzy models will differ from classical models. Specifically, the guiding idea of fuzzy models is that representing the extension of a predicate as a crisp set—a function from objects to  $\{0, 1\}$  (where objects sent to 1 are in the set and objects sent to 0 are out of the set)—is inadequate for vague predicates: because vague predicates have borderline cases and blurry boundaries. The thought is to replace the set  $\{0, 1\}$  of classical truth values with a set  $D$  of *degrees of truth* so that the extension of a vague predicate can be represented as a function from objects to  $D$  in a way that does justice to its vagueness: to its blurriness and to its possession of borderline cases. In particular, this will mean that  $D$  contains an element corresponding to the classical value 1—representing complete truth—and an element corresponding to the classical value 0—representing complete falsity—and many other elements in between. The gradual fading off of possession of the property  $P$  as we move along the sorites series is then modelled by the mapping of successive elements in the series to values in  $D$  that get gradually further from 1 and closer to 0. The elements of  $D$  (like the classical values 0 and 1) thus play two roles: they serve as degrees of property possession and degrees of truth. The value to which the extension of  $P$  sends the object  $x$  represents the degree to which  $x$  possesses the property picked out by  $P$ . This will then also be the degree of truth of the sentence  $Pa$ , when  $x$  is the referent of  $a$ . So the degree of truth of  $Pa$  is the degree to which the referent of  $a$  possesses the property picked out by  $P$ . This is exactly the kind of relationship we saw in the classical case: the truth value of  $Pa$  is the membership value of the referent of  $a$  in the extension of  $P$ . The only difference is that now these values do not have to be full-on (1) or full-off (0):  $D$  contains a host of values in between 1 and 0; this models the idea that objects can possess properties to intermediate degrees and (correspondingly) sentences can be true to intermediate degrees.

So far we have spoken in a general way about the set  $D$  of degrees of truth. When it comes time to be more specific about the exact content and structure of  $D$ , there are many options. Facts about the structure of  $D$  will show up in relationships amongst sentences. For example, if  $D$  is equipped only with an ordering and not with anything like a metric structure or a notion of distance between values, then we will be able to say things like ‘ $\alpha$  is truer than  $\beta$ ’ but we will not be able to say things like ‘ $\alpha$  is just a little bit truer than  $\beta$ , while  $\beta$  is a lot more true than  $\gamma$ ’ or ‘ $\alpha$  and  $\beta$  are very close in truth value’. For another example, if  $D$  is equipped with a linear ordering then for any two sentences  $\alpha$  and  $\beta$ , either one will be truer than the other or they will have exactly the same degree of truth—while if  $D$  is only partially ordered, then it might be that  $\alpha$  and  $\beta$  have distinct degrees of truth *without* either of them being truer than the other. So,

views about the appropriate structure of the set  $D$  of degrees of truth will be influenced by intuitions about such matters. Different options have been explored in the literature. Here it is useful to distinguish between a more concrete kind of approach and a more abstract kind of approach.

On the concrete sort of approach we pick some particular antecedently known structure as  $D$ . The most common option here is to let  $D$  be  $[0, 1]$ , comprising all the real numbers between 0 and 1 inclusive (together with all the usual structure: e.g. arithmetical operations such as multiplication, notions of distance between elements, and so on). On this approach, each primitive nonlogical symbol is assigned a value as follows: for an individual constant  $a$  its value  $[a]$  is an object in the domain of the model (as in the classical view); for an  $n$ -place predicate  $P$  its value  $[P]$  is a function from the set of  $n$ -tuples of members of the domain to the set  $[0, 1]$  of degrees of truth (as in the classical view—except that the set  $\{0, 1\}$  of classical truth values has been replaced by the set  $[0, 1]$  of degrees of truth). As before, such a function can be thought of as a *set* of  $n$ -tuples: but this time it is not a crisp set (where every object is in—sent to 1—or out—sent to 0) but a *fuzzy set*: the number  $x \in [0, 1]$  to which an  $n$ -tuple is sent represents the *degree* of membership of that  $n$ -tuple in the set. As before, this machinery suffices to determine a degree of truth for each closed wff that involves no logical vocabulary, via the following principle:

$$[P^n a_1 \dots a_n] = [P^n]([a_1], \dots, [a_n]).$$

That is, the degree of truth of  $P^n a_1 \dots a_n$  is the value to which the value of the predicate sends the  $n$ -tuple containing the values of the names (in the order in which they appear after the predicate). The statement of the principle is exactly the same as in the classical case—but this time, the truth value of  $P^n a_1 \dots a_n$  may be any real number between 0 and 1 inclusive.

On the abstract sort of approach we do not take a particular set of objects (together with a known structure thereon) as  $D$ —we just impose some structural properties and suppose that  $\mathbf{D}$  is some structure satisfying these properties. For example, we might specify that  $\mathbf{D} = \langle D, \vee, \wedge, \&, \rightarrow, 0, 1 \rangle$  is to be a *residuated lattice*,<sup>15</sup> that is:

1.  $\langle D, \vee, \wedge, 0, 1 \rangle$  is a lattice with least element 0 and greatest element 1.
2.  $\langle D, \&, 1 \rangle$  is a commutative monoid (i.e.  $\&$  is associative and commutative and for all  $x \in D$ ,  $x \& 1 = x$ ).
3.  $\rightarrow$  is the residuum of  $\&$  (i.e. for all  $x, y, z \in D$ ,  $x \& y \leq z$  iff  $x \leq y \rightarrow z$ ).

Or we might specify that  $\mathbf{D}$  is to be an *MTL-algebra* (a residuated lattice satisfying also the condition of *prelinearity*: for all  $x, y \in D$ ,  $(x \rightarrow y) \vee (y \rightarrow x) = 1$ ), or a *BL-algebra* (an MTL-algebra satisfying also the condition of *divisibility*: for all  $x, y \in D$ ,  $x \& (x \rightarrow y) = x \wedge y$ ) or a *BL-chain* (a BL-algebra where the underlying lattice is linearly ordered). There are also other options besides these ones that are studied in the literature (e.g. MV-algebras, product algebras, Gödel algebras...; see e.g. [10] and elsewhere in this

<sup>15</sup> Properly, a *bounded integral commutative* residuated lattice or  $\text{FL}_{\text{ew}}$ -algebra; see [34]. This simplification of terminology is common in the fuzzy logic literature; see [10, p. 22].



Handbook). Note that on this abstract kind of approach we don't care what particular objects are in  $D$ : we just care about the way that they are structured. Nevertheless, it would be odd to call the elements of  $D$  *degrees of truth* if there were (say) only two of them: 0 and 1. Yet the Boolean algebra of classical truth values can be taken as an example of many of these abstract structures. Hence when considering vagueness one tends to have in mind particular examples of these structures—for example,  $[0, 1]$ .

### 3.2 Truth rules

Let's return to classical models for our first order language. We now have enough machinery to determine a truth value for each closed wff that involves no logical vocabulary—but in order to determine a truth value for each closed wff that involves logical vocabulary we need further machinery: one piece per logical operator (connective or quantifier). For each  $n$ -place connective  $\nabla$  we specify a corresponding operation  $\blacktriangledown$  (of arity  $n$ ) on the classical truth values and then stipulate that

$$[\nabla(\alpha_1, \dots, \alpha_n)] = \blacktriangledown([\alpha_1], \dots, [\alpha_n]).$$

For example, corresponding to the negation connective  $\neg$  we have the unary operation  $x \rightarrow 1$  on  $\{0, 1\}$  and corresponding to the conjunction and disjunction connectives  $\wedge$  and  $\vee$  we have the binary operations  $\min\{x, y\}$  and  $\max\{x, y\}$  on  $\{0, 1\}$ ; thus  $[\neg\alpha] = 1 - [\alpha]$ ,  $[\alpha \wedge \beta] = \min\{[\alpha], [\beta]\}$  and  $[\alpha \vee \beta] = \max\{[\alpha], [\beta]\}$ . For the universal (existential) quantifier, the truth value of the closed formula  $\forall x\alpha$  ( $\exists x\alpha$ ) is determined by considering all the truth values that one gets by taking the free variable  $x$  in  $\alpha$  to denote some object in the domain—one truth value per object in the domain—and then applying an infinitary analogue of the conjunction (disjunction) operation to these truth values.

In the fuzzy case, one option is to proceed in an analogous way: to specify an operation on the set  $D$  of degrees of truth corresponding to each logical operator. How specifically one does this will depend on the structure imposed on  $D$  at the first stage of specifying fuzzy models (i.e. the stage of specifying the degrees of truth, examined in Section 3.1). But the general idea is the same in all cases: one uses the structure in the set of truth degrees to define logical operators.

In the case where one has taken  $[0, 1]$  as the set of truth degrees—together with (as we mentioned) all its usual structure (e.g. arithmetical operations such as multiplication, notions of distance between elements, and so on)—there is plenty to work with. For universal and existential quantification, it is standard to define these using the infimum and supremum operations respectively (which are indeed infinitary analogues of one kind of conjunction and disjunction: the min conjunction and max disjunction, which we shall encounter shortly). For the connectives of propositional logic, on the other hand, there are multiple live options; let me just mention a few which have played a prominent role in theories of vagueness. First consider Zadeh logic (Figure 3). Negation, conjunction and disjunction are defined precisely as above in classical logic—although this time the operations on truth values take all reals in  $[0, 1]$  as inputs and outputs, not just 1 and 0. The conditional is defined in terms of negation and conjunction—or equivalently negation and disjunction—in precisely the way that is familiar from classical logic. Likewise, the biconditional is defined in terms of conditional and conjunction in the familiar classical way.

$$\begin{aligned}
[\neg\alpha] &= 1 - [\alpha] \\
[\alpha \wedge \beta] &= \min\{[\alpha], [\beta]\} \\
[\alpha \vee \beta] &= \max\{[\alpha], [\beta]\} \\
[\alpha \rightarrow \beta] &= [\neg\alpha \vee \beta] = [\neg(\alpha \wedge \neg\beta)] \\
[\alpha \leftrightarrow \beta] &= [(\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)]
\end{aligned}$$

Figure 3. Zadeh logic

$$\begin{aligned}
[\neg\alpha] &= 1 - [\alpha] \\
[\alpha \wedge \beta] &= \min\{[\alpha], [\beta]\} \\
[\alpha \vee \beta] &= \max\{[\alpha], [\beta]\} \\
[\alpha \rightarrow \beta] &= \begin{cases} 1 & \text{if } [\alpha] \leq [\beta] \\ 1 - [\alpha] + [\beta] & \text{if } [\alpha] > [\beta] \end{cases} \\
[\alpha \leftrightarrow \beta] &= [(\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)]
\end{aligned}$$

Figure 4. Philosophers' Fuzzy Logic

Second, consider what we may call Philosophers' Fuzzy Logic or PFL (Figure 4).<sup>16</sup> Negation, conjunction and disjunction are the same as in Zadeh logic. The definition of the biconditional looks the same as in Zadeh logic but note that the result is different, because the conditional featuring in the definition is different. As for the conditional, the idea is this: if the antecedent isn't truer than the consequent, then the conditional is true to degree 1; but if the antecedent is truer than the consequent, then whatever the difference between them, the conditional falls precisely *that* far short of complete truth.

Third, consider t-norm fuzzy logics. A t-norm is a binary function  $\wedge$  on  $[0, 1]$  satisfying the conditions:

$$\begin{aligned}
x \wedge y &= y \wedge x \\
(x \wedge y) \wedge z &= x \wedge (y \wedge z) \\
x_1 \leq x_2 &\Rightarrow x_1 \wedge y \leq x_2 \wedge y \\
y_1 \leq y_2 &\Rightarrow x \wedge y_1 \leq x \wedge y_2 \\
1 \wedge x &= x \\
0 \wedge x &= 0.
\end{aligned}$$

A t-norm logic is specified by picking a t-norm and taking it to be the conjunction operation, and then defining the other operations (conditional, negation and so on) in certain specific ways. Notably, the conditional is taken to be the residuum of the t-norm (the residuum exists iff the t-norm is left-continuous) and the negation the precomplement of the conditional:

$$x \rightarrow y = \max\{z : x \wedge z \leq y\} \quad \neg x = x \rightarrow 0.$$

<sup>16</sup>The reason for the name is that many philosophers write as if 'fuzzy logic' just *is* PFL.

	Łukasiewicz logic	Gödel logic	Product logic
$x \wedge y$	$= \max(0, x + y - 1)$	$= \min(x, y)$	$= x \cdot y$
$x \rightarrow y$	$= \begin{cases} 1 & \text{if } x \leq y \\ 1 - x + y & \text{if } x > y \end{cases}$	$= \begin{cases} 1 & \text{if } x \leq y \\ y & \text{if } x > y \end{cases}$	$= \begin{cases} 1 & \text{if } x \leq y \\ y/x & \text{if } x > y \end{cases}$
$\neg x$	$= 1 - x$	$= \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{otherwise} \end{cases}$	$= \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{otherwise} \end{cases}$

Figure 5. Three prominent t-norm logics

Figure 5 shows the conjunctions, conditionals and negations in three prominent t-norm logics.

It is common in these logics to define a second, ‘weak’ (or ‘lattice’) conjunction (with the t-norm conjunction then termed ‘strong’). In all these logics, the weak conjunction is the same as the min operation used to define conjunction in Zadeh logic.<sup>17</sup>

Instead of taking an antecedently known structure such as  $[0, 1]$  as degrees of truth, one might have specified a structure of degrees of truth in an abstract way—for example, as  $\langle D, \vee, \wedge, \&, \rightarrow, 0, 1 \rangle$  with certain constraints imposed on  $\vee, \wedge, \&$  and so on. In this case the obvious way to proceed will be to associate a connective in the logical language with an operator of the same arity in the algebra of truth degrees: either one directly given in the definition of the algebra, or one defined in terms of those. So, in the examples discussed above, one can again get *two* conjunctions: one defined as  $\wedge$  and one as  $\&$ .

I said near the beginning of this section that when it comes to specifying truth values for closed wffs that involve logical vocabulary, *one* approach is to proceed in an analogous way to the classical approach: to specify an operation on the set of degrees of truth corresponding to each logical operator. We have just explored some particular options within this approach. Another kind of approach is also possible: one where we assign truth values in a non-degree-functional way. For example, given a fuzzy valuation  $V$ —an assignment of a referent to each name and a fuzzy set of  $n$ -tuples to each  $n$ -place predicate—we can define the notion of a classical extension of this valuation in an obvious way: it is a classical model  $\mathfrak{M}$  that assigns the same referents to names as  $V$ , and assigns extensions to predicates in such a way that whenever the value of  $P$  on  $V$  sends an  $n$ -tuple to 1 (or to 0), its value on  $\mathfrak{M}$  also sends that  $n$ -tuple to 1 (or, respectively, to 0). We can then assign closed formulas degrees of truth in—for example—the supervaluationist way, the subvaluationist way or the degree supervaluationist way.

<sup>17</sup> So in Gödel logic, there is no difference between the strong and weak conjunction. Note that the min conjunction of Figure 4 is definable using the operations for Łukasiewicz logic of Figure 5 as  $\alpha \wedge (\alpha \rightarrow \beta)$  and the Łukasiewicz t-norm conjunction of Figure 5 is definable using the negation and conditional of Figure 4 (which are the same as the negation and conditional for Łukasiewicz logic of Figure 5) as  $\neg(\alpha \rightarrow \neg\beta)$ . So ‘PFL’ and ‘Łukasiewicz logic’ pick out different perspectives (rather than different logics): from the PFL perspective, the Łukasiewicz t-norm conjunction is ignored (it is not even defined, let alone put to any use: only min conjunction is considered); from the Łukasiewicz logic perspective, the Łukasiewicz t-norm conjunction is of central importance (although not necessarily to the *exclusion* of min conjunction, which may *also* be considered).

### 3.3 The intended model

I have distinguished two parts to fuzzy MTS: the part that tells us about fuzzy models and the part that tells us which models are relevant to questions of the (actual) meaning and truth (simpliciter) of utterances in a given vague discourse. Within the first part, I have again distinguished two aspects: the ‘degrees of truth’ part—discussed in Section 3.1—which provides the machinery for modelling degrees of property possession and for assigning degrees of truth to closed wffs involving no logical vocabulary; and the ‘truth rules’ part—discussed in Section 3.2—which provides the machinery for assigning degrees of truth to closed wffs involving logical vocabulary.

What about the second part of fuzzy MTS—the part that introduces some notion like that of the ‘intended model’ that plays the role of distinguishing one (or some) of the infinity of models as the one(s) relevant to questions of the (actual) meaning and truth (simpliciter) of utterances in some discourse? For the moment we shall work with the simplest possible assumption: each vague discourse is associated with a unique intended fuzzy model. Later we shall return to this issue and find reason to explore a plurivaluationist alternative.

## 4 Arguments for

In this section we look at some of the major considerations in favour of adopting a theory of vagueness based on fuzzy logic.<sup>18</sup>

### 4.1 The nature of vagueness

All the theories of vagueness examined in Section 2 revolve in one way or another around the idea of sharpening the boundaries of vague predicates: either the models involved are already classical (in the case of epistemicism and plurivaluationism) or else the theory gives a central place to the idea of removing borderline cases (either completely, in the case of supervaluationism, or partially, in the case of contextualism). Now of course *one* thing we can do with vague predicates is sharpen them. But we also employ them in their natural, unsharpened state—blurry boundaries and all. We do so very successfully all the time. One motivation for fuzzy theories is that they offer an account which does justice to vagueness *as it is*—without one eye (or both) on how it could be *removed* (or reduced). Of course at this stage this does not amount to an argument for fuzzy theories. It does, however, capture part of the motivation of many who are attracted to fuzzy theories—and furthermore something like this line of thought can be turned into a real argument to the conclusion that only fuzzy theories can do justice to the phenomenon of vagueness. We examine this argument now.

In Section 1.1 above we introduced vague predicates via three characteristics: blurry boundaries, borderline cases and sorites susceptibility. This can be compared to explaining what water is by saying it’s a clear potable liquid that falls as rain and boils at 100°C. This will help someone who doesn’t know what water is to identify samples of it—but it still leaves open the question of the underlying nature or essence of water: of what water fundamentally *is*, that explains why it has these characteristics. The same goes for the

<sup>18</sup> More detailed versions of the arguments in Section 4.1 and Section 4.2.1 can be found in [106].

case of vagueness. It would be desirable to know the underlying essence of vagueness: to understand its fundamental nature and explain why it has these surface characteristics (blurry boundaries, borderline cases and sorites susceptibility).

Let's briefly consider some proposals for a fundamental definition of vagueness that do not work—before getting to a proposal that does work. For a start, one might think that perhaps one of the three characteristics of vague predicates is not merely a surface feature but is in fact the underlying essence of vagueness. However this thought does not pan out. Having borderline cases cannot be fundamental to vagueness, because a predicate could have *sharply delineated* borderline cases—borderline cases without blurry boundaries—and then it would not be vague. Generating sorites paradoxes cannot itself be what vagueness fundamentally consists in: for surely there is some feature that vague predicates possess and precise ones lack that *explains why* the former generate sorites paradoxes—and it is then this feature (rather than sorites-susceptibility, which this feature explains) that is fundamental to vagueness. Nor will it do to say that having blurry boundaries is the essence of vagueness—in the absence of a further explanation of what we mean by 'blurry boundaries'. Do we mean the purely epistemic sort of blur that the epistemicist countenances, or the probabilistic kind of blur that the degree supervaluationist countenances—or something else?

A different proposal that has been mentioned in the literature is that vagueness is semantic indeterminacy of the sort involved in plurivaluationism. However, such indeterminacy cannot be fundamental to vague predicates, because we can easily imagine predicates that exhibit such indeterminacy but are not vague—for example 'gavagai' or 'mass'. If Quine [88, ch. 2] and Field [29] are right, then these predicates exhibit semantic indeterminacy—but they do not generate sorites paradoxes, nor do their extensions have blurred boundaries: hence they are not vague.

A proposal that has received considerable discussion in the literature is that vague predicates are *tolerant*. There are various ways of formulating tolerance but for our purposes the following will be most useful. Recall that in a sorites series for the predicate  $P$ , each object is extremely similar to the object after it in all respects relevant to the application of  $P$  (for short, in ' $P$ -relevant respects'). A predicate  $P$  is tolerant iff it satisfies the following principle (for any objects  $a$  and  $b$ ):

**Tolerance** If  $a$  and  $b$  are very similar in  $P$ -relevant respects then  $Pa$  and  $Pb$  are identical in respect of truth.

(Two sentences are identical in respect of truth if they have the same truth value, or both lack a truth value, or in general have exactly the same truth status—where the possible truth statuses of a sentence will depend on the system of MTS in question.) The problem with Tolerance as a definition of vagueness is that it leads to contradiction. Suppose that we have a sorites series for  $P$ —so ' $x$  is  $P$ ' is true when  $x$  is the first object in the series and false when  $x$  is the last object in the series, and adjacent pairs of objects in the series are very similar in  $P$ -relevant respects. Then, given Tolerance, for every object  $x$  in the series, the claim ' $x$  is  $P$ ' is both true and false.

This brings us to a proposal which, I shall argue, does work: that vagueness is fundamentally Closeness. That is, a predicate  $P$  is vague iff it satisfies the following principle (for any objects  $a$  and  $b$ ):

**Closeness** If  $a$  and  $b$  are very similar in  $P$ -relevant respects then  $Pa$  and  $Pb$  are very similar in respect of truth.

This yields what we wanted in a fundamental definition of vagueness: an explanation of why vague predicates have the three characteristics mentioned at the outset. We'll return to sorites susceptibility in Section 4.2.1; here let's consider the other two characteristics.

**Blurry boundaries:** Suppose that we have a range of (possible) objects, some of which are  $F$  and some of which are not, and where we can get from any object to any other by passing between objects that are very similar in  $F$ -relevant respects. (For example, consider a sorites series for  $F$ .) If  $F$  conforms to Closeness, then its extension cannot consist in a sharp line between the  $F$ 's and the non- $F$ 's: for then we would have two objects  $a$  and  $b$  which are very similar in  $F$ -relevant respects, with  $a$  on one side of the line and  $b$  on the other, so that  $Fa$  is true and  $Fb$  false—and this would violate Closeness. Rather,  $F$ -ness must gradually fade away as one travels further from the definite  $F$  objects. To take a concrete example, consider the term 'red', and suppose that it conforms to Closeness. This term does not cut a sharp band out of the rainbow: as one moves across the points of the rainbow, small steps in red-relevant respects—which in this case correspond to small steps in space—can never, given Closeness, make for big changes in the truth of the claim that the point one is considering is red. By small steps one can move from full-fledged red points to full-fledged non-red points: but there is no sharp boundary between them that can be crossed in one small step. Thus Closeness yields an explanation of the blurred boundaries phenomenon.

**Borderline cases:** If a predicate satisfies Closeness then it admits of borderline cases. Consider a predicate  $F$  that conforms to Closeness, and a sorites series  $x_1, \dots, x_n$  for  $F$ .  $Fx_1$  is true and  $Fx_n$  is false; but given Closeness, it cannot be that there is an  $i$  such that  $Fx_i$  is true and  $Fx_{i+1}$  is false. There must then be sentences  $Fx_i$  which are neither true nor false—and the corresponding objects  $x_i$  are borderline cases for  $F$ .

Thus, if we define vagueness in terms of Closeness, we can explain why vague predicates have blurry boundaries and admit of borderline cases—and also why they generate sorites paradoxes (although we shall not discuss that until Section 4.2.1). Other advantages of the Closeness definition are that it accommodates the intuitions that have been taken to support the idea that vague predicates satisfy Tolerance, without generating contradictions; and it accommodates, within the definition of vagueness itself, the intuitions that lead those who wish to define vagueness in terms of possession of borderline cases to posit an additional phenomenon—'higher-order vagueness'—over and above mere vagueness.<sup>19</sup> So there are good reasons to think that conforming to Closeness is the essence of vagueness. But if that is so, then it leads to an argument in favour of fuzzy theories of vagueness. In brief, the argument is that no other kind of theory can allow for the existence of vague predicates.

First consider the epistemicist view. It will be helpful to consider a concrete example. Suppose we have a strip of paper that is red at the left end and orange at the right end, and in between gradually changes colour from red to orange. Now consider the sentence 'Point  $x$  is red', for each point  $x$  on the strip. According to the epistemicist,

<sup>19</sup> For details see [105] and [106, ch. 3].

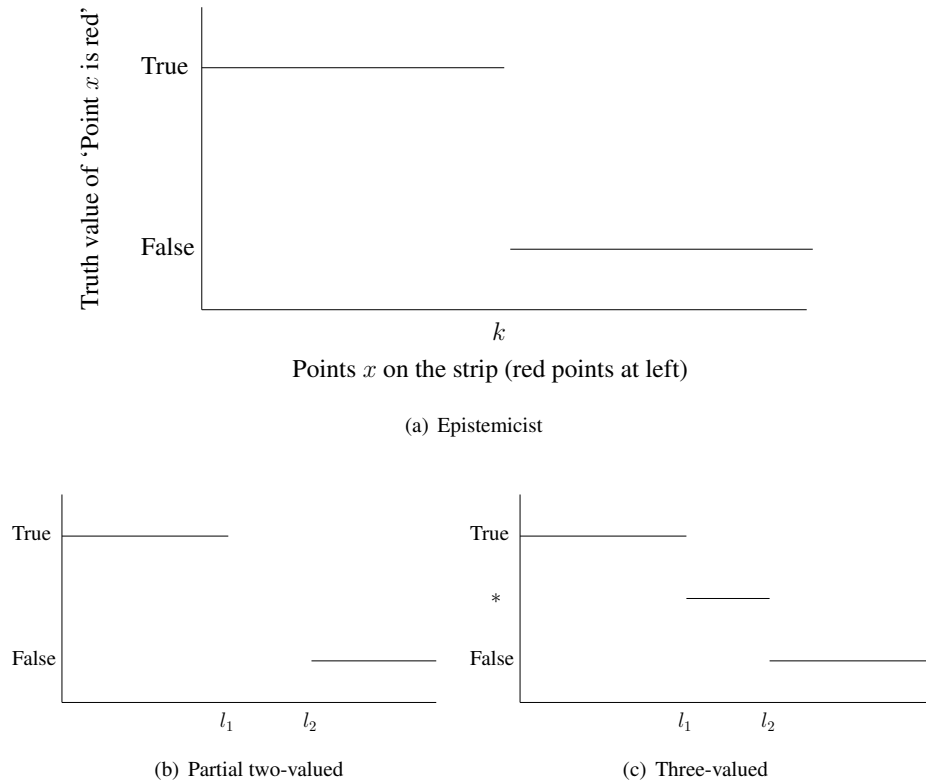


Figure 6. The strip according to epistemicist, partial two-valued and three-valued views

every one of these sentences is either true or false. Figure 6(a) represents the situation according to the epistemicist. This flouts Closeness: there are points  $k$  (shown on the diagram) and  $k + \Delta$  (a point to the right of, and arbitrarily close to,  $k$ ) whose colours are *very* close together, while the sentence ‘Point  $k$  is red’ is True and the sentence ‘Point  $k + \Delta$  is red’ is False. This sort of flouting of Closeness is unavoidable for the epistemic theory. The epistemicist, then, cannot allow the existence of predicates  $F$  that both conform to Closeness and have associated sorites series. Thus—given that vagueness is Closeness—he cannot allow for the existence of vague predicates (that have associated sorites series).

Going plurivaluationist does not solve the problem. Any acceptable model must make  $F$  true of the things at the beginning of the sorites series for  $F$  and false of the things at the end. But (on the plurivaluationist view) every acceptable model is classical. Thus, on *every* acceptable model,  $F$  will violate Closeness (for the reasons just seen). But if Closeness is violated on every acceptable model, then—on the plurivaluationist picture—it is violated everywhere. There are only the acceptable models, and so there simply is nowhere else for Closeness to be accommodated.

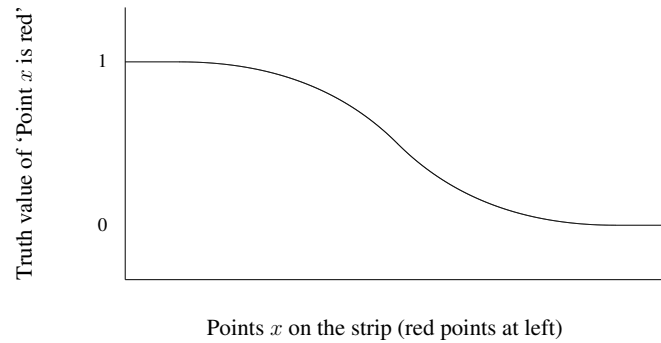


Figure 7. The strip according to the fuzzy view

Tripartite views fare no better. Consider again our strip of paper. Figure 6(b) represents the situation according to the truth gap view and Figure 6(c) represents the situation according to the third value view. According to these views, there are points  $l_1$  and  $l_1 + \Delta$  whose colours are very close together, while the sentence ‘Point  $l_1$  is red’ is True and the sentence ‘Point  $l_1 + \Delta$  is red’ has the third status (i.e. a third value, or a gap); and there are points  $l_2$  and  $l_2 + \Delta$  whose colours are very close together, while the sentence ‘Point  $l_2$  is red’ has the third status and the sentence ‘Point  $l_2 + \Delta$  is red’ is False. But if one sentence is True and another False, then they are very far apart in respect of truth—and so no third truth status can be *very close* to *both* of them. Thus, to the extent that a sentence is very close to True, it is *not* very close to False, and vice versa. Hence Closeness must be violated by at least one of the pairs of points  $l_1/l_1 + \Delta$  or  $l_2/l_2 + \Delta$ —and given the natural assumption that the third status is symmetric with respect to truth and falsity, Closeness will be violated at both places.

This problem arises for all tripartite views: going supervaluationist (rather than recursive) makes no difference. Even if ‘Point  $l_1$  is red’ is classically true/false on almost exactly the same admissible extensions as ‘Point  $l_1 + \Delta$  is red’, at best this makes the two sentences very similar in the respects that determine truth. If one sentence is True and the other has the third status, then they are not very similar in respect of truth.<sup>20</sup>

If Closeness gives the essence of vagueness, then to allow that there are vague predicates (with associated sorites series) we need to be able to accommodate the idea of a predicate  $F$  where  $Fx$  is true when  $x$  is the first object in the series and false when  $x$  is the last object in the series—and  $Fx$  and  $Fx'$  are always very similar in respect of truth, when  $x$  and  $x'$  are adjacent objects in the series. This requires *degrees of truth*. For example, consider a view on which the set of truth values is  $[0, 1]$ . In this case we can easily accommodate Closeness; for example, we could handle the strip of paper as in Figure 7. Here there are no points  $x$  and  $x'$  that are very close together on the strip but where  $Fx$  and  $Fx'$  are not very similar in respect of truth.

<sup>20</sup> Or more precisely: the third status cannot be very similar in respect of truth *both* to True and to False.



Taking  $[0, 1]$  as the set of truth degrees would, then, be one way to accommodate Closeness. It is not the only way—and furthermore, not *any* set of degrees of truth will suffice. For example, if the only structure on the degrees of truth was an ordering—if there was no meaningful way to say that two truth values were *very close* or that the *difference* between them was *very small* or something else along these lines—then (even if there were infinitely many of these degrees of truth) we would not be able to accommodate Closeness.

A natural response on the part of non-fuzzy theorists to the foregoing argument for degrees of truth from the Closeness definition of vagueness is to say that the definition is subtly wrong. Suppose that we defined vagueness not in terms of Closeness—which features the notion of similarity in respect of truth—but in terms of a variant principle in which truth is replaced by, say, warranted assertability. In that case we could not conclude that degrees of truth are required if vague predicates are to be allowed. We shall return to this line of thought in Section 4.2.1. In the meantime, to sum up the present section: the first big advantage of fuzzy theories of vagueness (at least fuzzy theories that take  $[0, 1]$  as the set of degrees of truth—or some other structure that has the right properties to accommodate predicates that satisfy Closeness and have associated sorites series) is that they are motivated by the best account we have of what vagueness fundamentally consists in—viz., conformity to Closeness.

## 4.2 Solving the sorites

The second big advantage of fuzzy theories is that they alone are in a position to solve the sorites paradox without departing from the basic modus operandi of formal semantics, which assumes that speakers' linguistic behaviour flows from their semantic competence (modulated by performance factors) [101]. It is not assumed that speakers have a full understanding of the semantics of their language. However, it is assumed that speakers have some sort of implicit grasp of this semantics—and that this grasp manifests itself in their behaviour. That is why speakers' linguistic behaviour is evidence for formal semantic theories. It goes entirely against the grain of this kind of approach to posit a semantic theory  $T$  for a language which is such that if speakers thought that  $T$  was the correct semantics for their language, then they would use their language quite differently from the way they actually use it. This is, as we saw, the kind of position that non-fuzzy theories of vagueness find themselves in when it comes to the sorites paradox. Ordinary speakers respond to sorites paradoxes in the following sort of way: they find the first premise undeniable; they are strongly inclined to accept the second premise; they find no fault in the reasoning leading from the premisses to the conclusion; and yet they find the conclusion unacceptable. However, a speaker who thought that epistemicism (for example) gives the correct semantics of the predicate  $P$  would have no inclination at all to think that the second premise was true (and—as we have seen—similar remarks apply to other non-fuzzy theories of vagueness). Fuzzy theories, on the other hand, can offer a semantic theory  $T$  for vague language which does have the following desired feature: if speakers thought that theory  $T$  gave the correct semantics for a vague predicate  $P$ , they would respond to a sorites argument for  $P$  in just the way they *do* ordinarily respond to sorites paradoxes.

There are two different kinds of fuzzy approach to the sorites in the literature. One solves the paradox using only resources already available at the ‘degrees of truth’ stage of formulating a fuzzy theory; the other posits specific resources at the ‘truth rules’ stage and uses those to solve the paradox. Let’s examine the two approaches in turn.<sup>21</sup>

#### 4.2.1 Sorites via degrees of truth

We said in Section 4.1 that if we define vagueness as Closeness then we can explain why vague predicates generate sorites paradoxes. Let’s now see why this is so. Distinguish two readings of the second premise in a sorites argument (i.e. two different claims that one might understand the second premise to be making): the Closeness reading and the Tolerance reading. On the Closeness reading, the second premise expresses (a particular consequence of) the claim that  $F$  conforms to Closeness:  $x$  and  $x'$  are very similar in  $F$ -relevant respects, so if  $x$  is  $F$  then for all practical purposes we can just say that  $x'$  is  $F$  too. On the Tolerance reading, the second premise expresses (a particular consequence of) the claim that  $F$  conforms to Tolerance:  $x$  and  $x'$  are very similar in  $F$ -relevant respects, so if  $x$  is  $F$  then—not just for all practical purposes but without qualification— $x'$  is  $F$  too (the two claims  $Fx$  and  $Fx'$  are *exactly the same* in respect of truth). Now note the following:

1. On the Closeness reading, the conclusion does not follow from the premises.
2. If one thinks that  $F$  conforms to Closeness (and not to Tolerance), one will nevertheless tend to go along with claims that  $F$  is Tolerant in normal situations.
3. On the Tolerance reading, the conclusion does follow from the premises.

Claim 3 is obvious. The basic idea behind claim 1 is this: each successive statement  $Fx$  must be *very similar* in respect of truth to the one before, but need not be *exactly the same* in respect of truth—and so by the end, the final statement  $Fx_n$  may be simply false. In fact one can formalise the sorites argument (under a Closeness reading) and show that it is invalid.<sup>22</sup> This leaves Claim 2. Suppose we accept that a predicate  $F$  satisfies Closeness, but not Tolerance. The very fact that we accept Closeness will mean that in many ordinary circumstances, we act as if we believe Tolerance. For Closeness tells us that a negligible or insignificant difference between  $a$  and  $b$  in  $F$ -relevant respects makes for at most a negligible or insignificant difference between  $Fa$  and  $Fb$  in respect of truth—and a negligible or insignificant difference is one that we are entitled to ignore for practical purposes. So for practical purposes, when there is a negligible or insignificant difference between  $a$  and  $b$  in  $F$ -relevant respects, we will simply ignore any difference between  $Fa$  and  $Fb$  in respect of truth, and so treat them as being identical in respect of truth. This practice is licensed by our acceptance of Closeness. However the practice has its limits: sometimes we find ourselves in a situation in which Tolerance cannot be accepted as a useful approximation of Closeness—and a sorites series is precisely such a context. Consider an analogy. We do not believe that dust particles are weightless: we believe that the weight of a dust particle is negligible. (This is the analogue for dust particles of Closeness without Tolerance.) But this very

<sup>21</sup> Historically, the ‘truth rules’ kind of approach came first.

<sup>22</sup> This is done in one way in [106, p. 271] and in several different ways in [42].

belief licenses us to act as if we believe that dust particles are weightless: given that the weight of a dust particle is negligible, we do well to ignore it! We do not demand that the delicatessen assistant remove all specks of dust from the scale arms before weighing our smallgoods, and we do not wash and dry our hair (to remove all dust particles) before weighing ourselves. So our belief that the weight of a dust particle is negligible licenses us to accept the claim that dust particles are weightless as a useful approximation of our real belief, in ordinary circumstances. (This is the analogue for dust particles of Tolerance.) Nevertheless, the claim that dust particles are weightless is revealed as a mere approximation to our real belief—something that we act as if we believe, in ordinary circumstances, but not something we actually believe—in certain situations, for example when we are arranging to empty the bag from the dust extraction system at our carpentry shop, which weighs 85kg when full (of nothing but dust particles). The claim that a dust particle weighs nothing is a useful approximation to our true belief, except when we come across many dust particles together, at which time we see clearly that the claim is just an *approximation* to what we really believe, which is that the weight of a dust particle is extremely small. Similarly, the claim that the predicate ‘heap’ (say) applies equally to two things that differ negligibly in heap-relevant respects is a useful approximation of the real belief of someone who accepts Closeness (but not Tolerance), except when she encounters many pairs of things that differ negligibly in heap-relevant respects put together—a sorites series—at which time it becomes clear that the claim is just an *approximation* of what she really believes, which is that the difference in truth value between ‘ $x$  is a heap’ and ‘ $y$  is a heap’ is at most very small, when  $x$  and  $y$  are very similar in heap-relevant respects.

Putting together the ingredients just laid out, we can now see why a predicate  $F$  that conforms to Closeness generates a sorites paradox. The sorites argument is compelling because when the second premise is taken to express Tolerance, (a) the argument is valid and (b) the second premise is (initially) acceptable (to one who accepts Closeness but not Tolerance) because accepting Closeness licenses one to accept Tolerance as a useful approximation of what one believes, in ordinary situations. Nevertheless, the argument is ultimately flawed because it leads one to see that in situations involving sorites series one cannot happily use Tolerance as an approximation of Closeness, and must work with Closeness itself—and when the major premise is taken to express Closeness (and not Tolerance), the argument is invalid. Thus, someone who believes that  $F$  conforms to Closeness (and not Tolerance) can be expected to react to a sorites argument in precisely the way ordinary speakers do react: he will initially find the premisses acceptable and will agree that the conclusion follows from them—but he will not ultimately go on to accept that the conclusion is actually true.

We have now explained why predicates that satisfy Closeness will generate sorites paradoxes. This adds a piece that was earlier left missing to the argument that Closeness is the correct fundamental definition of vagueness. In light of this, let’s turn now to our current task, which is to show that fuzzy theories can solve the sorites paradox in a way that does not require ordinary speakers to lack an understanding of the semantics of their own vague predicates. We saw in Section 4.1 that accommodating vagueness (defined as Closeness) requires positing degrees of truth (with a certain sort of structure). Now any fuzzy theory that employs degrees of truth with the right kind

of structure—for example  $[0, 1]$ —has the resources to distinguish predicates that satisfy Closeness (and not Tolerance) from predicates that satisfy Tolerance. But once we can meaningfully distinguish Closeness and Tolerance, a solution to the sorites paradox is ready to hand. A sorites argument (as presented in English) can be understood in two ways: the second premise might express Closeness or Tolerance. On the Tolerance reading the argument is unsound (vague predicates are not Tolerant: they merely conform to Closeness) and on the Closeness reading it is invalid. This is why the argument is mistaken. However the argument is also compelling (to one who thinks that the predicate involved in the sorites argument conforms to Closeness but not Tolerance)—for the reason seen above. Believing that  $F$  conforms to Closeness, one is thereby inclined to accept the second premise on the Tolerance reading (as a useful approximation); and as the argument is valid on that reading, one is led to the conclusion. One does not then simply accept the conclusion, however: the argument is compelling, but it is not ultimately convincing. At that point one realises that one is in a special situation in which Tolerance is not a useful approximation of Closeness—and when one goes back and reads the second premise as expressing Closeness, the conclusion no longer follows.

Before turning to the second approach to solving the sorites within fuzzy theories, we need to tie up one more loose end. As indicated earlier, proponents of non-fuzzy theories of vagueness might try to block the argument for fuzzy theories presented in Section 4.1 by claiming that vagueness should be defined in terms not of Closeness but of some variant principle that differs from Closeness by substituting something such as the notion of *warranted assertability* in place of the notion of *truth*—for example:

**A-Closeness** If  $a$  and  $b$  are very similar in  $P$ -relevant respects then  $Pa$  and  $Pb$  are very similar in respect of warranted assertability.

Note that there is no argument from A-Closeness to degrees of truth: epistemicists, for example, should have no trouble accommodating A-Closeness. But we are now in a position to see why A-Closeness (unlike Closeness) does not provide an adequate definition of vagueness: there is no connection between A-Closeness and the sorites. Suppose that we have a predicate  $F$  that conforms to A-Closeness and suppose that we have a sorites series for  $F$ . Will the corresponding sorites argument be compelling? We have no reason to think so. Given A-Closeness, and the fact that adjacent objects in the sorites series are very close in  $F$ -relevant respects, we know that for any object in the series, to whatever extent it is justifiable, reasonable or warranted to assert, believe or judge that it is  $F$ , it is to a very similar extent justifiable, reasonable or warranted to assert, believe or judge that the next object is  $F$ . But we have no reason at all to conclude from this that the second premise of the sorites argument is true. Whether or not the second premise is compelling depends crucially on *why* we think A-Closeness holds for  $F$ . If it holds because  $F$  satisfies Closeness, then the second premise does become compelling, as we have seen. However, if  $F$  satisfies A-Closeness but not Closeness—say because  $F$  works in the way the epistemicist thinks vague predicates work (it draws sharp but unknowable boundaries)—then we have no reason at all to accept the second premise, and so the sorites argument will not be compelling in the slightest. This shows that A-Closeness—by itself—yields no account of sorites-susceptibility. If we define vague

predicates in terms of A-Closeness, we are left with no understanding of why they generate sorites paradoxes. Hence A-Closeness—unlike Closeness—does not provide an adequate definition of vagueness.

#### 4.2.2 Sorites via truth rules

In the previous section we saw that simply positing a structure of degrees of truth (such as  $[0, 1]$ ) that allows for predicates that satisfy Closeness without satisfying Tolerance opens the way to a solution to the sorites paradox. Now we turn to a different strategy for solving the sorites within fuzzy theories: one that turns on truth rules—in particular, on positing particular truth conditions for the connectives used in formulating the sorites argument.

Consider a version of the paradox that concerns a series of piles of sand 1 through 10,000, where pile  $i$  has  $i$  grains of sand and each pile is of a very similar shape to its neighbour(s). Consider the following sorites argument:

Pile 10,000 is a heap.

If pile 10,000 is a heap then pile 9,999 is a heap.

If pile 9,999 is a heap then pile 9,998 is a heap.

⋮

If pile 2 is a heap then pile 1 is a heap.

∴ Pile 1 is a heap.

Let's suppose that 'if...then...' here is read as the Łukasiewicz conditional and that we define validity as follows: on every model on which every premise is true to degree 1, the conclusion is true to degree 1. Then we get the following solution to the sorites. The problem with the argument is that, although it is valid, it is unsound (i.e. it is not the case that every premise is true to degree 1). The first premise is true to degree 1. As for the conditionals, at first both antecedent and consequent are true to degree 1, and so are the conditionals. As we move along the series, we get to a point at which the antecedents are ever so slightly more true than the consequents. In this region, the conditionals are true to a degree ever so slightly less than 1. This continues for a while until both antecedent and consequent are true to degree 0, and hence the conditionals are true to degree 1 again. So why is the argument compelling? Because all the premises are *very nearly* true to degree 1. In normal contexts, we are naturally inclined simply to accept something as true when it is very nearly true—this is a useful approximation. Of course, once we see where the argument leads, we may well reconsider.<sup>23</sup>

<sup>23</sup> See [31, pp. 243–4], [32, pp. 171–2], [123, pp. 123–4], [82, p. 365] and [81]. Note that the explanation of why the sorites is compelling is sometimes put in terms of ordinary speakers being fooled—mistaking near truth for full truth. But this just gives away the advantage of fuzzy theories, which is that they can explain speakers' reactions to the sorites paradox without resorting to the view that speakers are mistaken about the semantic facts. In the explanation given in the text above, we do not suppose that speakers mistakenly think

Note that the solution depends on a particular choice of truth conditions for the conditional and a particular definition of validity. Suppose we instead define validity as follows: on every model, the truth value of the conclusion is greater than or equal to the infimum of the truth values of the premisses. Then modus ponens (for the Łukasiewicz conditional) and the above sorites argument are invalid—and so we lose the explanation given above of why the argument is compelling.<sup>24</sup> Or suppose we read ‘if...then...’ in the argument as (say) the Zadeh or Gödel conditional: in that case some premisses would have degrees of truth of around 0.5, and so again we lose the explanation given above of why the argument is compelling.

So far we have considered just one formulation of the sorites argument: one involving multiple conditional premisses. But what about other formulations? If we can solve the sorites only when it is formulated in one particular way, then we have not really solved the underlying problem. Crispin Wright took this to be a problem for fuzzy views because he thought that they could not handle conjunctive formulations of the paradox [124, pp. 251–2]—for example:

Pile 10,000 is a heap.

It is not the case that (pile 10,000 is a heap and pile 9,999 is not a heap).

It is not the case that (pile 9,999 is a heap and pile 9,998 is not a heap).

⋮

It is not the case that (pile 2 is a heap and pile 1 is not a heap).

∴ Pile 1 is a heap.

However, Wright was assuming that fuzzy logic is PFL. In this logic, some of the premisses of the argument just given have degrees of truth of around 0.5, and so indeed we cannot run the kind of explanation given above of why the argument is compelling. However, the problem dissolves in Łukasiewicz logic: if we take the ‘and’ to be strong conjunction, then ‘If pile 2 is a heap then pile 1 is a heap’ is equivalent to ‘It is not the case that (pile 2 is a heap and pile 1 is not a heap)’ [81]. Thus the only moral here is the one we already saw: this kind of solution to the sorites paradox is not universally applicable but requires a careful choice of fuzzy logic.

## 5 Arguments against

In this section we look at—and respond to—the two major objections to fuzzy theories of vagueness.<sup>25</sup>

that the premisses are fully true: rather we exploit the fact that someone who takes a statement to be extremely close to fully true would naturally just go along with the statement, at least in normal contexts and until trouble was seen to arise.

<sup>24</sup> See [72, pp. 69–75], [123, pp. 123–4] and [85, p. 332].

<sup>25</sup> More detailed versions of the arguments in Section 5.1 can be found in [106] and [107]. More detailed versions of the arguments in Section 5.2 can be found in [106] and [102]. Further objections are considered and responded to in [106].

### 5.1 Artificial precision

The first objection to fuzzy theories is that they involve *artificial precision*:<sup>26</sup>

[Fuzzy logic] imposes artificial precision. . . [T]hough one is not obliged to require that a predicate either definitely applies or definitely does not apply, one *is* obliged to require that a predicate definitely applies to such-and-such, rather than to such-and-such other, degree (e.g. that a man 5 ft 10 in tall belongs to *tall* to degree 0.6 rather than 0.5) — Haack [38, p. 443]

One immediate objection which presents itself to [the fuzzy] line of approach is the extremely artificial nature of the attaching of precise numerical values to sentences like ‘73 is a large number’ or ‘Picasso’s *Guernica* is beautiful’. In fact, it seems plausible to say that the nature of vague predicates precludes attaching precise numerical values just as much as it precludes attaching precise classical truth values.— Urquhart [119, p. 108]

[T]he degree theorist’s assignments impose precision in a form that is just as unacceptable as a classical true/false assignment. In so far as a degree theory avoids determinacy over whether *a* is *F*, the objection here is that it does so by enforcing determinacy over the *degree* to which *a* is *F*. All predications of “is red” will receive a unique, exact value, but it seems inappropriate to associate our vague predicate “red” with any particular exact function from objects to degrees of truth. For a start, what could determine which is the correct function, settling that my coat is red to degree 0.322 rather than 0.321? — Keefe [57, p. 571]

In a nutshell, the problem is that it is artificial/implausible/inappropriate to associate each vague *predicate* in natural language with a particular function that assigns one particular real number between 0 and 1 to each object (the object’s degree of possession of the property picked out by that predicate); likewise, it is artificial/implausible/inappropriate to associate each vague *sentence* in natural language with a particular real number between 0 and 1 (the sentence’s degree of truth).

The first thing to note is that the problem only arises for fuzzy theories with the following two features: they take  $[0, 1]$  as the set of truth degrees; and they posit a unique intended model of each vague discourse. For it is only theories of this kind that *do* associate each vague predicate with a particular function from objects to reals between 0 and 1 (i.e. the extension of the predicate on the unique intended model) and each sentence with a particular real between 0 and 1 (i.e. the sentence’s degree of truth on the unique intended model). For ease of exposition, let’s call such theories ‘basic fuzzy theories’. (Note that there are many of them, not just one: they differ amongst themselves over the truth conditions of the logical operators.) We’ll also refer to  $[0, 1]$  as the ‘standard fuzzy truth values’ or sftv’s for short.

The next thing to do is determine the true nature and source of the problem. Haack offers no diagnosis; Urquhart maintains that the *nature of vague predicates* precludes

<sup>26</sup> For further statements of the problem see [19, pp. 521–2], [36, p. 332], [37, p. 54], [64, p. 462, p. 481], [72, p. 187], [74], [94, pp. 223–4], [98, p. 46], [118, p. 11], [123, pp. 127–8] and [58, pp. 113–4].

attaching precise numerical values; Keefe asks what could *determine* which is the correct function, settling that her coat is red to degree 0.322 rather than 0.321. Assuming that vagueness is correctly defined in terms of Closeness, we can see that Keefe is on the right track and Urquhart is not: the problem with the fuzzy view turns not on considerations having to do with the nature of vagueness, but rather on considerations having to do with the way in which the meanings of our terms are fixed. In order to see this, it will help to consider epistemicism again. As we saw in Section 4.1, epistemicism conflicts with the nature of vagueness: it cannot allow for the existence of predicates that satisfy Closeness (and have associated sorites series). As we saw in Section 2.2, epistemicism also runs into problems concerning the determination of meaning: it seems that the meaning-determining facts itemized in Section 2.2 do not suffice to pick out a particular height dividing the tall from the non-tall, and so on. The epistemicist therefore faces two distinct problems:

1. The *existence* of a sharp drop-off from true to false in a sorites series: this conflicts with the nature of vagueness.
2. The *particular location* of the drop-off: this conflicts with our best views about how meaning is determined.

I refer to these two problems as the *jolt problem* and the *location problem* respectively.

Now let's return to basic fuzzy theories. As we saw in Section 4.1, they do not face a version of the jolt problem: they can allow for the existence of predicates that satisfy Closeness (and have associated sorites series) and so they do not conflict with the nature of vagueness. However they *do* face a version of the location problem. It seems that the meaning-determining facts itemized in Section 2.2 do *not* suffice to pick out a particular function from objects to  $[0, 1]$  representing the extension of 'is tall' (and similarly for other vague predicates). So the fundamental problem underlying the artificial precision worry is that basic fuzzy theories conflict with our best available views about how the meanings of our expressions are determined.

In light of this diagnosis, let's consider some responses to the problem that have been proposed in the literature. We'll see several that don't work before getting to one that does. Recall that the problem arises for basic fuzzy theories, which have two key ingredients: they take sftv's as truth degrees; and they posit a unique intended model of each vague discourse. Responses to the problem can be categorised according to whether they abandon the first ingredient or the second.

The first response we shall consider—*fuzzy epistemicism*—actually tries to avoid the problem without abandoning either ingredient. This view is modelled after epistemicism, which attempts to retain the classical semantic picture (i.e. classical model theory plus a unique intended model) in the face of vagueness. The fuzzy epistemicist holds that each vague sentence does indeed have a unique sftv as its truth value—but we do not (cannot) know which value it is.<sup>27</sup> However, once we are clear that the artificial precision problem concerns the determination of meaning, this approach evidently fails

<sup>27</sup> Fuzzy epistemicism is mentioned by Copeland [19, p. 522] and developed in more detail by MacFarlane [71]. Machina [72, p. 187, n. 8] could also be interpreted as hinting at such a view when he writes of “difficulties about how to assign degrees of truth to propositions”; Keefe in [57, p. 571] and [58, p. 115] interprets him in this way and criticises his view on this basis.






Figure 8. Bob's degree of tallness

to solve the problem. The problem concerns how there could *be* a unique function that is the extension of 'is tall' (say), given that our usage (etc.) does not suffice to pick out a unique such function. Saying that we do not or cannot *know* which function it is misses the point of the problem.

Let's turn then to a second proposal, due to [103]—the *blurry sets* view, which tries to solve the artificial precision problem by moving from  $[0, 1]$  to a different set of truth values: *degree functions*, which are functions from  $[0, 1]^*$  to  $[0, 1]$ .<sup>28</sup> Suppose  $f: [0, 1]^* \rightarrow [0, 1]$  is the truth value of 'Bob is tall' (B). The idea is that the value that  $f$  assigns to the empty sequence—say, 0.5—is a *first approximation* to Bob's degree of tallness/the degree of truth of (B). The values assigned by  $f$  to sequences of length 1 then play two roles. First, they rate possible first approximations. The higher the value assigned to  $\langle x \rangle$ , the better  $x$  is as a first approximation to Bob's degree of tallness/the degree of truth of (B). If  $f(\langle 0.3 \rangle) = 0.5$ , then we say that it is 0.5 true that Bob is tall to degree 0.3; if  $f(\langle 0.5 \rangle) = 0.7$ , then we say that it is 0.7 true that Bob is tall to degree 0.5; and so on. Second, the assignments to sequences of length 1 jointly constitute a second level of approximation to Bob's degree of tallness/the degree of truth of (B). Together, these assignments determine a function  $f_{\langle \rangle}: [0, 1] \rightarrow [0, 1]$ . We regard this as encoding a density function over  $[0, 1]$ , and we require that its centre of mass is at  $f(\langle \rangle)$  (Figure 8(a)).<sup>29</sup> The same thing happens again when we move to the values assigned to sequences of length 2: these values play two roles. First, they rate possible ratings of first approximations. The higher the value assigned to  $\langle x, y \rangle$ , the better  $y$  is as a rating of  $x$  as a first approximation to Bob's degree of tallness/the degree of truth of (B). If  $f(\langle 0.5, 0.7 \rangle) = 0.8$ , then we say that it is 0.8 true that it is 0.7 true that Bob is tall to degree 0.5; if  $f(\langle 0.4, 0.5 \rangle) = 0.3$ , then we say that it is 0.3 true that it is 0.5 true that Bob is tall to degree 0.4; and so on. Second, the assignments to sequences of length 2

<sup>28</sup> On notation: where  $S$  is a set,  $S^*$  is the set of all finite sequences of elements of  $S$  (including the empty or null sequence).

<sup>29</sup> I am here giving a rough picture of the view: in order to make it precise, various definitions and assumptions are required—see [103] for the details.

jointly constitute a third level of approximation to Bob's degree of tallness/the degree of truth of (B). Together, the assignments made by  $f$  to sequences  $\langle a, x \rangle$  of length 2 whose first member is  $a$  determine a function  $f_{\langle a \rangle} : [0, 1] \rightarrow [0, 1]$ . This can be seen as encoding a density function, and we require that its centre of mass is at  $f(\langle a \rangle)$  (Figure 8(b)). And so the story goes, ad infinitum. Figuratively, we can picture a degree (of truth or property-possession) as a region of varying shades of grey spread between 0 and 1 on the real line. If you focus on any point in this region, you see that what appeared to be a point of a particular shade of grey is in fact just the centre of a further such grey region. The same thing happens if you focus on a point in this further region, and so on. The region is blurry all the way down: no matter how much you increase the magnification, it will not come into sharp focus. On this view, statements of the form 'The degree of truth of "Bob is tall" is 0.4' need not be simply true or false: they may themselves have intermediate degrees of truth. So rather than exactly one sentence of the form 'The degree of truth of "Bob is tall" is  $x$ ' being true and the others false, many of them might be true to various degrees. Thus there is a sense in which sentences in natural language which predicate vague properties of objects are *not* each assigned just one particular sftv. Be that as it may, however, the artificial precision problem has not been solved. In just the way that they fail to determine a unique *classical* set (function from objects to classical truth values) or a unique *fuzzy* set (function from objects to fuzzy truth values) as the extension of 'is tall', the meaning-determining facts do not suffice to pick out a unique *blurry set* (function from objects to degree functions) as the extension of 'is tall'.

The third proposal that we shall consider—the *ordinal approach*—is that when we assign sftv's to sentences, the only thing about the assignments that is meaningful is the relative *ordering* of the values assigned. Views of this sort have been advocated by—amongst others—Goguen, Machina and Hyde:<sup>30</sup>

We certainly do not want to claim there is some *absolute* [fuzzy] set representing 'short'... It appears that many arguments about fuzzy sets do not depend on particular values of functions... This raises the problem of *measuring* fuzzy sets... Probably we should not expect particular numerical values of shortness to be meaningful (except 0 and 1), but rather their *ordering*... degree of membership may be measured by an *ordinal scale*. [36, pp. 331–2]

the assignment of exact values usually doesn't matter much... what is of importance instead is the ordering relation between the values of various propositions. [72, p. 188]

The foregoing account... requires only a totally-ordered dense set of values. The choice of a specific value from among the infinitely many possible... is arbitrary except in so far as it preserves ordering requirements... No significance attaches to the choice of value apart from these ordering requirements. [46, p. 207]

There are actually two quite different ways of spelling out this idea in detail.<sup>31</sup>

<sup>30</sup> See also [96, p. 29], [37, p. 59], [40, pp. 162–3] and [121].

<sup>31</sup> For a more detailed discussion see [109].

On the first way of looking at things, the truth degrees are not sftv's. The real interval  $[0, 1]$  comprises some entities, together with some structure—an order structure, a metric structure—and some operations—addition, subtraction, and so on. Now suppose we retain the entities and the order structure, but discard the metric structure and any operations definable in terms of it or vice versa (e.g. subtraction). This gives us a new structure—and its elements are the truth values of the new sort of model theory now under consideration. Figuratively, one can think of the new structure as a rubbery unit interval, fixed at each end: its end-points have fixed positions, but between them, none of the other elements has a fixed position. They can be squeezed or stretched left or right at will—but they can never leapfrog one another: their *order* is fixed. Let us refer to the elements of this structure—the rubbery unit interval—as rtv's, and to models in which the truth degrees are rtv's as rubbery models (i.e. models that assign rubbery sets to predicates—where a rubbery set is a function from the domain to the rtv's—and rtv's to sentences). On the first way of spelling out the ordinal approach, then, each vague discourse is associated with a unique intended model—but it is a rubbery model, not a standard fuzzy model (i.e. a model involving sftv's).

On the second way of spelling out the ordinal approach, there are no rtv's, only sftv's. We represent the facts about the truth of statements by assigning sftv's to them, but there are many acceptable ways of doing so: for any acceptable way of mapping sentences to sftv's, any mapping obtained by composing it with an order-preserving and endpoint-fixing transformation of  $[0, 1]$  is equally acceptable. Thus there is no unique intended model of a discourse: there is a whole class of acceptable models, differing from one another by such transformations of the sftv's assigned—but they are all standard fuzzy models.

On either way of spelling out the ordinal approach to the artificial precision problem, a serious difficulty looms. In trying to resolve the conflict between basic fuzzy theories and the facts about how meanings are determined, we have generated a new conflict: with facts about the nature of vagueness (which was an area where basic fuzzy theories had no problem). On the view that truth is measured on an ordinal scale, it *makes no sense* to say that two sentences  $P$  and  $Q$  are *very close* in respect of truth: it makes sense to say only that one sentence is *true* than another. But the idea of two sentences being *very close* in respect of truth is at the heart of the Closeness definition—and so a view that makes no room for this notion lacks the resources to accommodate predicates that satisfy Closeness.<sup>32</sup>

The fourth and final response to the artificial precision problem that we shall consider here is *fuzzy plurivaluationism*, due to [106].<sup>33</sup> Fuzzy plurivaluationism is just like classical plurivaluationism except that its models are fuzzy, not classical. It stands to basic fuzzy theories of vagueness (on which a vague discourse is associated with a unique intended standard fuzzy model) in just the way that classical plurivaluationism stands to the classical semantic picture (on which a discourse is associated with a unique intended classical model). That is, everything about the original view is retained (so in the classical case, only standard classical models are countenanced, and in the

<sup>32</sup> Recall that it was already noted in Section 4.1 that if the only structure on the truth degrees was an ordering then we would not be able to accommodate Closeness.

<sup>33</sup> For a more detailed presentation and motivation of this view, see [106, §2.5, ch. 6]. See also [9].

fuzzy case, only standard fuzzy models are countenanced), *except* the idea that each discourse is associated with a unique intended model. The latter idea is replaced with the thought that each discourse is associated with multiple acceptable models. Which models? The fuzzy plurivaluationist answer to this question is different from the answer given by the second way of spelling out the ordinal approach discussed above. On the ordinal approach, the acceptable models are closed under order-preserving and endpoint-fixing transformations of the truth degrees. On the plurivaluationist approach, by contrast, the acceptable models are those models not ruled out as incorrect by the meaning-determining facts. (For example, suppose that speakers in a discourse are all disposed to count certain paradigm individuals as tall and certain others as definitely not tall; in that case, acceptable models of the discourse must assign ‘tall’ an extension that maps the former individuals to 1 and the latter to 0.) As we have seen, the meaning-determining facts do not suffice to single out a unique intended standard fuzzy model of vague discourse: this is precisely what generates the artificial precision problem for basic fuzzy theories. *A fortiori*, there are multiple acceptable models (where ‘acceptable’ means ‘not ruled out as incorrect by the meaning-determining facts’). Note that fuzzy plurivaluationism avoids the problem faced by the ordinal view of not being able to say meaningfully that two sentences are very close in respect of truth. One plausible constraint on acceptable models is that vague predicates come out as satisfying Closeness relative to them. Hence if  $x$  and  $y$  are very similar in  $F$ -relevant respects,  $Fx$  and  $Fy$  must be very similar in respect of truth on every acceptable model (if  $F$  is vague). In that case we will indeed be able to talk as if these two sentences are simply very close in respect of truth.

Fuzzy plurivaluationism successfully solves the artificial precision problem. The problem is that the meaning-determining facts do not suffice to pick out a unique standard fuzzy model of vague discourse as the intended model. The fuzzy plurivaluationist solution is to abandon the notion of a unique intended model in favour of the idea of multiple acceptable models—where an acceptable model is one that is not ruled out as incorrect by the meaning-determining facts. As the problem is precisely that there is not a unique acceptable fuzzy model of vague discourse—because too many models are compatible with the constraints imposed by the meaning-fixing facts—it follows *a fortiori* that fuzzy plurivaluationism—the view that there are multiple equally correct models—is correct. The upshot of fuzzy plurivaluationism is that ‘Bob is tall’ (say) does *not* have a uniquely correct sftv: it is assigned multiple different sftv’s—one on each acceptable model—and none of these is more correct than any of the others. This was the desired result: that it was not the case on basic fuzzy theories was precisely the artificial precision problem. Furthermore, we have avoided assigning a unique sftv to each vague sentence *for the right reason*: because doing so is incompatible with the meaning-determining facts.

## 5.2 Truth functionality

The second objection to fuzzy theories is that they are incompatible with ordinary usage of compound propositions in the presence of borderline cases. This is a very common objection and there are many different versions of it to be found in the literature. Many of the objectors falsely assume that fuzzy logic is simply PFL—but some versions

of the objection are more general: they are directed against any truth(-degree)-functional account. Here is a sample of the objections:

1. *Fine I and Osherson and Smith I.* Suppose that a certain blob is on the border of pink and red and let  $P$  be the sentence ‘the blob is pink’ and  $R$  the sentence ‘the blob is red’—so  $P$  and  $R$  are neither clearly true nor clearly false. Fine thinks that  $P \vee R$  is clearly true and that  $P \wedge R$  is clearly false. This is not predicted by a fuzzy account based on PFL.<sup>34</sup> On a related note, Osherson and Smith think that where  $Ax$  means that  $x$  is an apple,  $Aa \wedge \neg Aa$  should be true to degree 0 and  $Aa \vee \neg Aa$  should be true to degree 1, whatever  $a$  is. This conflicts with PFL.<sup>35</sup>
2. *Kamp.* Kamp thinks that  $\alpha \wedge \neg\alpha$  is clearly true to degree 0, even when  $[\alpha] = 0.5$ ; and that  $\alpha \wedge \alpha$  is clearly true to a degree strictly greater than 0, when  $[\alpha] = 0.5$ . Assuming  $[\neg\alpha] = 1 - [\alpha]$ , no truth function for  $\wedge$  can predict this. So this is an argument not just against PFL but against any degree-functional account that agrees with PFL about negation.<sup>36</sup>
3. *Fine II.* With  $P$  and  $R$  as in 1, Fine claims that  $P \wedge P$  is equivalent to  $P$  and hence is neither clearly true nor clearly false, while (as already discussed)  $P \wedge R$  is clearly false. Given that  $P$  and  $R$  have the same degree of truth, this is an argument against any degree-functional account of conjunction.<sup>37</sup>
4. *Osherson and Smith II.* Consider an apple (a), illustrated thus:



(a)

Osherson and Smith claim that (a) is psychologically less prototypical of the concept ‘apple’ than of the concept ‘striped apple’. If we equate prototypicality with degree of membership/truth and take ‘striped apple’ to be formed from the two components ‘striped’ and ‘apple’ by intersection/conjunction, then this is an objection to the minimum rule for conjunction—and more generally to any account according to which  $[\alpha \wedge \beta]$  can never be strictly greater than  $[\alpha]$  (or  $[\beta]$ ).<sup>38</sup>

<sup>34</sup> See [30, pp. 269–70].

<sup>35</sup> See [79, pp. 45–6]. Osherson and Smith present their argument in terms of degrees of membership of objects in sets rather than degrees of truth of statements.

<sup>36</sup> See [53, p. 546].

<sup>37</sup> See [30, p. 269].

<sup>38</sup> See [79, pp. 43–5]. Essentially the same objection (using the example of *pussy willow* and *willow*) was made earlier by Kay [56, p. 153]. Osherson and Smith also present a dual objection concerning alleged disjunctions that are apparently less true than either disjunct [79, pp. 46–8].

The objections are often based solely on the intuitions of their authors—for example:

we would have  $[\phi \wedge \neg\phi] = \frac{1}{2}$ , which seems absurd. For how could a logical contradiction be true to *any* degree? [53, p. 546]<sup>39</sup>

a given object  $x$  may be a triangle (say) to degree 0.9;  $f_{\Delta}(x) = 0.9$ . If the complement of  $f_{\Delta}$  represents ‘is not a triangle’ and union disjunction, then  $\max(f_{\Delta}, 1 - f_{\Delta})$  should represent ‘is a triangle or isn’t a triangle’ *and should be the constant 1 function*; but it isn’t. [119, p. 108; my emphasis]<sup>40</sup>

the same value must be given to (e) ‘if Tek is tall then Tek is not tall’ as to (f) ‘if Tek is tall then Tek is tall’ since the respective values of the antecedent and consequent of these two conditionals are the same. But (f) is intuitively true and (e) is not, so again no choice of value will capture our intuitions about both of these cases [58, p. 97].

So one perfectly legitimate response is to express alternative intuitions—for example:<sup>41</sup>

[Fine] claims that ‘red’ and ‘pink’, even though vague and admitting of borderline cases of applicability, are nevertheless logically connected so that to say of some color shade that it is both red and pink is obviously to say something false. I must confess being completely insensitive to that intuition of a penumbral connection. [72, p. 77, n. 2]

An aspect of the theory of fuzzy sets which Osherson and Smith find objectionable is that, in the theory, the union of  $A$  and its complement,  $A'$ , is not, in general, the whole universe of discourse. This relates, of course, to the long-standing controversy regarding the validity of the *principle of the excluded middle*. . . The principle of the excluded middle is not accepted as a valid axiom in the theory of fuzzy sets *because it does not apply to situations in which one deals with classes which do not have sharply defined boundaries*. [126, p. 292; second emphasis mine]

the failure of contradiction and excluded-middle laws is typical of fuzzy logic as emphasized by many authors. This is natural with gradual properties like ‘tall’. [21, p. 152]

Sometimes the objections are based on empirical data, but there are also considerable empirical data in support of various fuzzy theories. For example, studies in [5], [6], [92] and [97] show a significant willingness of subjects to agree with statements such as

<sup>39</sup> I have omitted a superscript and a subscript from Kamp’s notation because they add complexity that is irrelevant in the present context. Cf. [30, p. 270] “Surely  $P \& \neg P$  is false even though  $P$  is indefinite”; [50, p. 199] “This consequence is absurd, because a self-contradiction surely merits a truth value of zero”; [123, p. 136] “How can an explicit contradiction be true to any degree other than 0?”; [55, p. 134] “According to the most familiar versions of fuzzy logic the degree to which  $a$  satisfies the conjunctive concept *apple which is not an apple* is . . . greater than 0. Clearly this is not the right result”; and [43, pp. 366–7] “. . . an obviously counterintuitive conclusion”.

<sup>40</sup> Cf. [55, pp. 146–7], [16, pp. 389–90] and [43, pp. 366–7].

<sup>41</sup> Cf. also [51, pp. 287–8], [65, p. 141], [33, pp. 323–4], [58, p. 164] (NB Keefe, unlike fellow supervaluationists Fine and Kamp & Partee, does not take  $Fa \vee \neg Fa$  to be assertable when  $a$  is a borderline case of  $F$ ), [12, p. 578], [13, p. 31], [106, p. 86] and [93, p. 341].

‘ $X$  is tall and not tall’ and ‘The circle is near the square and it is not near the square’.<sup>42</sup> Furthermore, we must also be very careful to ensure that any data that are established are actually relevant to the assessment of fuzzy theories. For example, as we saw above, Osherson and Smith [79, pp. 43–5] take it as a datum that “There can be no doubt that [(a)] is psychologically less prototypical of an apple. . . than of an apple-with-stripes” and then take this to mean that (a)’s degree of membership in the set of striped apples is greater than its degree of membership in the set of apples. But one could accept the former datum about *prototypicality* and yet maintain that when it comes to degrees of *membership* (and *truth*), (a) is a degree 1 member of ‘striped apple’ and of ‘apple’ (and ‘(a) is a striped apple’ and ‘(a) is an apple’ are both quite simply true, to degree 1).<sup>43</sup>

At this point, then, it has not been established that there are consistent patterns of ordinary usage that pose even a *prima facie* threat to fuzzy theories of vagueness. But suppose for the sake of argument that we take the intuitions of the objectors as genuine data. There are in fact many ways in which fuzzy theories could accommodate such data. We can divide them broadly into semantic and pragmatic approaches; and we can further divide the semantic approaches into those that employ degree-functional truth rules and those that employ non-degree-functional truth rules. Let’s consider these approaches in turn.

**Degree-functional semantics:** Some of the intuitions of the objectors can be accommodated straightforwardly using fuzzy resources that have already been mentioned in this chapter. For example:

1. *Fine I and Osherson and Smith I.* In Łukasiewicz t-norm logic, when  $[P] = [R] = 0.5$ ,  $[P \vee R] = 1$  and  $[P \wedge R] = 0.44$ . This meets Fine’s desiderata. Likewise, in Łukasiewicz logic, *whatever* the degree of truth of  $\alpha$ ,  $[\alpha \wedge \neg\alpha] = 0$  and  $[\alpha \vee \neg\alpha] = 1$  [13, p. 31], [11, p. 138] and [81].<sup>45</sup> This meets Osherson and Smith’s desiderata.
2. *Kamp.* We do not have to define  $[\neg\alpha] = 1 - [\alpha]$ . In Gödel t-norm logic,  $\alpha \wedge \neg\alpha$  is 0 true and  $\alpha \wedge \alpha$  is 0.5 true. In product t-norm logic,  $\alpha \wedge \neg\alpha$  is 0 true and  $\alpha \wedge \alpha$  is 0.25 true. This meets Kamp’s desiderata.
3. *Fine II.* In Łukasiewicz t-norm logic, when  $[P] = [R] = 0.5$ , ‘ $P$  and  $P$ ’ is true to degree 0.5 where ‘and’ is read as weak conjunction and ‘ $P$  and  $R$ ’ is true to degree 0 where ‘and’ is read as strong conjunction [81]. This meets Fine’s desiderata.

Furthermore, there are oceans of functions on  $[0, 1]$  available to fuzzy theorists that we have not yet mentioned. For example, recall Osherson and Smith’s idea that (a)’s degree of membership in ‘striped apple’ should be greater than its degree of membership in ‘apple’. This sort of situation could be modelled by taking ‘striped apple’ to be formed from the sets ‘striped’ and ‘apple’ not by a conjunction/intersection operation

<sup>42</sup> See [102] for further discussion and references.

<sup>43</sup> The point that degrees of membership and truth on the one hand and degrees of typicality on the other hand need to be carefully distinguished has been made by numerous authors including [126, p. 293], [100, pp. 51–2], [52], [55, p. 131, p. 133], [80, p. 191], [12, p. 578] and [11, pp. 132–3].

<sup>44</sup> Where  $\wedge$  is the Łukasiewicz t-norm and  $\vee$  is its dual:  $x \vee y = 1 - ((1 - x) \wedge (1 - y))$ .

<sup>45</sup> Where  $\wedge$  is the Łukasiewicz t-norm and  $\vee$  is its dual.

but by an averaging operation [12, p. 578]. For another example, one can introduce a ‘determinately’ operator  $\Delta$  [114] and [7]:

$$[\Delta\alpha] = \begin{cases} 1 & \text{if } [\alpha] = 1 \\ 0 & \text{if } [\alpha] \neq 1 \end{cases}$$

When  $[\alpha] = 0.5$ , then even using the conjunction and negation of PFL:

$$[\Delta\alpha \wedge \Delta\neg\alpha] = [\Delta\alpha \wedge \neg\Delta\alpha] = 0$$

**Pragmatics:** A moment ago we saw one way in which apparently non-degree-functional data could be accommodated within a degree-functional semantics: Fine’s idea that ‘ $P$  and  $R$ ’ should be clearly false while ‘ $P$  and  $P$ ’ is neither clearly true nor clearly false—which seems to rule out a degree-functional treatment of ‘and’—could be accommodated by reading ‘and’ as weak conjunction in ‘ $P$  and  $P$ ’ and as strong conjunction in ‘ $P$  and  $R$ ’. It is not easy, however, to make this sort of approach—which posits an ambiguity in ‘and’—work in detail,<sup>46</sup> so it is important to note that there is a second way in which apparently non-degree-functional data could be accommodated within a theory that employed a degree-functional semantics: via pragmatics. Suppose that a consistent pattern of usage of the sort Fine finds intuitive was in fact discovered. There is no reason why a fuzzy theorist should not accommodate this by saying that while ‘and’ is unambiguously interpreted as some particular truth function, nevertheless the *assertability* of conjunctions sometimes goes by a different rule, viz.: assert  $\Phi$  iff  $\Phi$  would certainly be true no matter how the language were precisified. (Compare the following well-known view [66], [47], [48] and [69]: the truth conditions of the indicative conditional ‘If  $A$  then  $B$ ’ are the same as those of the material conditional  $A \supset B$ ; however *assertability* for indicative conditionals goes not by truth but by conditional probability (of  $B$  given  $A$ ).) In other words, if the data appear to support a supervaluationist picture, this does not in fact mean that we need a supervaluationist semantics: a supervaluationist approach could instead be located in the pragmatics.

**Non-degree-functional semantics:** Finally, even if it were for some reason important to treat some connective in a non-degree-functional way (i.e. in the semantics), this does not rule out fuzzy theories! As mentioned in Section 3.2, there is no reason why fuzzy theories cannot use non-degree-functional approaches at the ‘truth rules’ stage of developing fuzzy models.

## 6 Historical remarks and further reading

### Section 1

The characterisation of vagueness in terms of borderline cases can be traced at least as far as [83], while the blurred boundaries characterisation can be traced at least as far as Frege’s statement that if we represent concepts in extension by areas on a plane, then vague concepts do not have sharp boundaries, but rather fade off into the background (*Grundgesetze* vol. II, §56) [8, p. 259]. A few authors still introduce vagueness in terms of just one characteristic (e.g. possession of borderline cases) but since at least [60] the three-feature characterisation has become standard.

<sup>46</sup> See [77, p. 13] for one suggestion in this area and [28, pp. 200–1] for discussion.



## Section 2

Advocates of epistemic theories of vagueness include [18], [17], [112, ch. 6], [113], [122], [123, ch. 7–8] and [44]. For recursive tripartite approaches to vagueness see [116], [117] and [93]. For supervaluationist approaches see [75], [22], [30], [53], [84], [14] and [58, ch. 7–8]. For a subvaluationist approach to vagueness see [45]. For the degree-theoretic form of supervaluationism see [53], [54, pp. 234–5], [67, pp. 228–9] and [68, pp. 69–70]. Plurivaluationism was explicitly distinguished from supervaluationism in [106]; prior to that, the two views were conflated in the literature. [86] is a possible early example of a plurivaluationist approach. Contextualist treatments of vagueness include [54], [115], [89], [90], [111, ch. 7], [26] and [99]. Note that there are various contextualist theories in the literature which differ in more or less subtle ways (see e.g. [91, p. 245], [27, p. 329], [2], [3] and [1]) and that the necessarily brief Section 2.5 presents just one prominent kind of contextualist position; nevertheless a version of the objection that the contextualist solution to the sorites turns on speakers not believing that the contextualist story is the correct account of vague language applies also to other versions of contextualism (e.g. note that the contextualist solutions to the sorites considered in [3] turn on attributing *confusion* to ordinary speakers; see also [1, §3.1] and [2, §4]; and cf. also [59, §§4–5]).

## Section 3

The study of infinite-valued logics begins with Łukasiewicz; see e.g. [70]. Fuzzy set theory was born in [125]; a related earlier idea is Post's notion of an  $n$ -valued set [73, p. 47, p. 98]. For overviews of the history of fuzzy logic and set theory see [39, ch. 10], [78, ch. 8], [62] and [41]. For technical introductions to fuzzy logic and set theory, see e.g. [63], [39], [78], [76] and of course this Handbook. Note that 'Zadeh logic' has other names including 'Kleene–Zadeh logic' and sometimes simply 'fuzzy logic'. For discussion of non-truth-functional fuzzy logics see e.g. [35] and [41, p. 298]; cf. also [23–25]. Key early sources for the fuzzy view of vagueness are [36], [64] and [72]; [15] is a relevant earlier work. [106] is a recent comprehensive elaboration and defence of a fuzzy theory of vagueness.

## Acknowledgements

I am very grateful to Stewart Shapiro and Chris Fermüller for detailed comments on a draft version.

## BIBLIOGRAPHY

- [1] Jonas Åkerman. Contextualist theories of vagueness. *Philosophy Compass*, 7(7):470–480, 2012.
- [2] Jonas Åkerman and Patrick Greenough. Vagueness and non-indexical contextualism. In Sarah Sawyer, editor, *New Waves in Philosophy of Language*, pages 8–23. Palgrave Macmillan, 2009.
- [3] Jonas Åkerman and Patrick Greenough. Hold the context fixed—vagueness still remains. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds: Vagueness, its Nature, and its Logic*, pages 275–288. Oxford University Press, Oxford, 2010.
- [4] Ken Akiba and Ali Abasnezhad, editors. *Vague Objects and Vague Identity: New Essays on Ontic Vagueness*. Springer, 2014.
- [5] Sam Alxatib and Jeff Pelletier. On the psychology of truth-gaps. In Rick Nouwen, Robert Van Rooij, Uli Sauerland, and Hans-Christian Schmitz, editors, *Vagueness in Communication*, pages 13–36. Springer, 2011.

- [6] Sam Alxatib and Jeff Pelletier. The psychology of vagueness: Borderline cases and contradictions. *Mind and Language*, 26:287–326, 2011.
- [7] Matthias Baaz. Infinite-valued Gödel logic with 0-1-projections and relativisations. In Petr Hájek, editor, *Gödel'96: Logical Foundations of Mathematics, Computer Science, and Physics*, volume 6 of *Lecture Notes in Logic*, pages 23–33. Springer, Brno, 1996.
- [8] Michael Beaney, editor. *The Frege Reader*. Blackwell, Oxford, 1997.
- [9] Libor Běhounek. In which sense is fuzzy logic a logic for vagueness? In Thomas Lukasiewicz, Rafael Peñaloza, and Anni-Yasmin Turhan, editors, *Proceedings of the First Workshop on Logics for Reasoning about Preferences, Uncertainty, and Vagueness (PRUV 2014)*, pages 26–39, Vienna, 2014.
- [10] Libor Běhounek, Petr Cintula, and Petr Hájek. Introduction to mathematical fuzzy logic. In Petr Cintula, Petr Hájek, and Carles Noguera, editors, *Handbook of Mathematical Fuzzy Logic - Volume 1*, volume 37 of *Studies in Logic, Mathematical Logic and Foundations*, pages 1–101. College Publications, London, 2011.
- [11] Radim Bělohlávek and George J. Klir. Fallacious perceptions of fuzzy logic in the psychology of concepts. In Radim Bělohlávek and George J. Klir, editors, *Concepts and Fuzzy Logic*, pages 121–148. MIT Press, Cambridge, MA, 2011.
- [12] Radim Bělohlávek, George J. Klir, Harold W. Lewis III, and Eileen C. Way. On the capability of fuzzy set theory to represent concepts. *International Journal of General Systems*, 31:569–585, 2002.
- [13] Radim Bělohlávek, George J. Klir, Harold W. Lewis III, and Eileen C. Way. Concepts and fuzzy sets: Misunderstandings, misconceptions, and oversights. *International Journal of Approximate Reasoning*, 51:23–34, 2009.
- [14] Brandon Bennett. Modal semantics for knowledge bases dealing with vague concepts. In Anthony G. Cohn, Leonard Schubert, and Stuart C. Shapiro, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 6th International Conference (KR-98)*, pages 234–244. Morgan Kaufmann, 1998.
- [15] Max Black. Vagueness: An exercise in logical analysis. *Philosophy of Science*, 4(4):427–455, 1937.
- [16] Nicolao Bonini, Daniel Osherson, Riccardo Viale, and Timothy Williamson. On the psychology of vague predicates. *Mind and Language*, 14:377–393, 1999.
- [17] Richmond Campbell. The sorites paradox. *Philosophical Studies*, 26:175–191, 1974.
- [18] James Cargile. The sorites paradox. *British Journal for the Philosophy of Science*, 20:193–202, 1969. Reprinted in [61] pp. 89–98.
- [19] B. Jack Copeland. Vague identity and fuzzy logic. *Journal of Philosophy*, 94:514–534, 1997.
- [20] Steven Davis and Brendan S. Gillon, editors. *Semantics: A Reader*. Oxford University Press, Oxford, 2004.
- [21] Didier Dubois and Henri Prade. Can we enforce full compositionality in uncertainty calculi? In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 149–154. AAAI Press, 1994.
- [22] Michael Dummett. Wang's paradox. *Synthese*, 30:301–324, 1975. Reprinted in [61] pp. 99–118.
- [23] Dorothy Edgington. Validity, uncertainty and vagueness. *Analysis*, 52:193–204, 1992.
- [24] Dorothy Edgington. Vagueness by degrees. In Rosanna Keefe and Peter Smith, editors, *Vagueness: A Reader*, pages 294–316. MIT Press, Cambridge, MA, 1997.
- [25] Dorothy Edgington. The philosophical problem of vagueness. *Legal Theory*, 7(4):371–378, 2001.
- [26] Delia Graff Fara. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28: 45–81, 2000. Published under 'Delia Graff'.
- [27] Delia Graff Fara. Profiling interest relativity. *Analysis*, 68(4):326–335, 2008.
- [28] Christian G. Fermüller. Comments on 'Vagueness in language: The case against fuzzy logic revisited' by Uli Sauerland. In Petr Cintula, Christian G. Fermüller, Lluís Godo, and Petr Hájek, editors, *Understanding Vagueness: Logical, Philosophical, and Linguistic Perspectives*, volume 36 of *Studies in Logic*, pages 199–202. College Publications, London, 2011.
- [29] Hartry Field. Theory change and the indeterminacy of reference. *The Journal of Philosophy*, 70: 462–481, 1973.
- [30] Kit Fine. Vagueness, truth and logic. *Synthese*, 30:265–300, 1975. Reprinted with corrections in [61] pp. 119–150.
- [31] Graeme Forbes. Thisness and vagueness. *Synthese*, 54:235–259, 1983.
- [32] Graeme Forbes. *The Metaphysics of Modality*. Clarendon Press, Oxford, 1985.
- [33] Gy. Fuhrmann. "Prototypes" and "Fuzziness" in the logic of concepts. *Synthese*, 75:317–347, 1988.
- [34] Nikolaos Galatos, Peter Jipsen, Tomasz Kowalski, and Hiroakira Ono. *Residuated Lattices: An Algebraic Glimpse at Substructural Logics*, volume 151 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, Amsterdam, 2007.

- [35] Giangiacomo Gerla. *Fuzzy Logic—Mathematical Tool for Approximate Reasoning*, volume 11 of *Trends in Logic*. Kluwer and Plenum Press, New York, 2001.
- [36] Joseph Amadee Goguen. The logic of inexact concepts. *Synthese*, 19(3–4):325–373, 1969.
- [37] Joseph Amadee Goguen. Fuzzy sets and the social nature of truth. In Madan M. Gupta, Rammohan K. Ragade, and Ronald R. Yager, editors, *Advances in Fuzzy Set Theory and Applications*, pages 49–67. North-Holland, Amsterdam, 1979.
- [38] Susan Haack. Do we need “fuzzy logic”? *International Journal of Man–Machine Studies*, 11:437–445, 1979.
- [39] Petr Hájek. *Metamathematics of Fuzzy Logic*, volume 4 of *Trends in Logic*. Kluwer, Dordrecht, 1998.
- [40] Petr Hájek. Ten questions and one problem on fuzzy logic. *Annals of Pure and Applied Logic*, 96: 157–165, 1999.
- [41] Petr Hájek. What is mathematical fuzzy logic. *Fuzzy Sets and Systems*, 157(5):597–603, 2006.
- [42] Petr Hájek and Vilém Novák. The sorites paradox and fuzzy logic. *International Journal of General Systems*, 32(4):373–383, 2003.
- [43] James A. Hampton. Typicality, graded membership, and vagueness. *Cognitive Science*, 31:355–384, 2007.
- [44] Paul Horwich. *Truth*. Clarendon Press, Oxford, second revised edition, 1998.
- [45] Dominic Hyde. From heaps and gaps to heaps of gluts. *Mind*, 108:641–660, 1997.
- [46] Dominic Hyde. *Vagueness, Logic and Ontology*. Ashgate, Aldershot, 2008.
- [47] Frank Jackson. On assertion and indicative conditionals. In Frank Jackson, editor, *Conditionals*, pages 111–135. Oxford University Press, Oxford, 1979.
- [48] Frank Jackson. *Conditionals*. Basil Blackwell, Oxford, 1987.
- [49] Frank Jackson and Graham Priest, editors. *Lewisian Themes: The Philosophy of David K. Lewis*. Oxford University Press, Oxford, 2004.
- [50] Philip N. Johnson-Laird. *Mental Models*. Harvard University Press, Cambridge, MA, 1983.
- [51] Gregory V. Jones. Stacks not fuzzy sets: An ordinal basis for prototype theory of concepts. *Cognition*, 12:281–290, 1982.
- [52] Charles W. Kalish. Essentialism and graded membership in animal and artifact categories. *Memory & Cognition*, 23:335–353, 1995.
- [53] Hans Kamp. Two theories about adjectives. In Edward L. Keenan, editor, *Formal Semantics of Natural Language*, pages 123–155. Cambridge University Press, Cambridge, 1975. Reprinted in [20] pp. 541–562.
- [54] Hans Kamp. The paradox of the heap. In Uwe Mönnich, editor, *Aspects of Philosophical Logic*, pages 225–277. D. Reidel, Dordrecht, 1981.
- [55] Hans Kamp and Barbara Partee. Prototype theory and compositionality. *Cognition*, 57:129–191, 1995.
- [56] Paul Kay. A model-theoretic approach to folk taxonomy. *Social Science Information*, 14:151–166, 1975.
- [57] Rosanna Keefe. Vagueness by numbers. *Mind*, 107:565–579, 1998.
- [58] Rosanna Keefe. *Theories of Vagueness*. Cambridge University Press, 2000.
- [59] Rosanna Keefe. Vagueness without context change. *Mind*, 116(462):275–292, 2007.
- [60] Rosanna Keefe and Peter Smith. Introduction: theories of vagueness. In Rosanna Keefe and Peter Smith, editors, *Vagueness: A Reader*, pages 1–57. MIT Press, Cambridge, MA, 1997.
- [61] Rosanna Keefe and Peter Smith, editors. *Vagueness: A Reader*. MIT Press, Cambridge, MA, 1997.
- [62] George J. Klir. Foundations of fuzzy set theory and fuzzy logic: A historical overview. *International Journal of General Systems*, 30(2):91–132, 2001.
- [63] George J. Klir and Bo Yuan. *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Prentice Hall, 1995.
- [64] George Lakoff. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2(4):458–508, 1973.
- [65] George Lakoff. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press, Chicago, 1987.
- [66] David Lewis. Probabilities of conditionals and conditional probabilities. In Frank Jackson, editor, *Conditionals*, pages 76–101. Oxford University Press, Oxford, 1976.
- [67] David Lewis. General semantics. In *Philosophical Papers*, volume I, pages 189–232. Oxford University Press, New York, 1983.
- [68] David Lewis. Survival and identity. In *Philosophical Papers*, volume I, pages 55–77. Oxford University Press, New York, 1983.
- [69] David Lewis. Probabilities of conditionals and conditional probabilities II. In Frank Jackson, editor, *Conditionals*, pages 102–110. Oxford University Press, Oxford, 1986.

- [70] Jan Łukasiewicz and Alfred Tarski. Untersuchungen über den Aussagenkalkül. *Comptes Rendus des Séances de la Société des Sciences et des Lettres de Varsovie, cl. III*, 23(iii):30–50, 1930.
- [71] John MacFarlane. Fuzzy epistemicism. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds: Vagueness, its Nature, and its Logic*, pages 438–463. Oxford University Press, Oxford, 2010.
- [72] Kenton F. Machina. Truth, belief, and vagueness. *Journal of Philosophical Logic*, 5(1):47–78, 1976. Reprinted in [61] pp. 174–203.
- [73] Grzegorz Malinowski. *Many-Valued Logics*. Clarendon Press, Oxford, 1993.
- [74] Vincenzo Marra. The problem of artificial precision in theories of vagueness: A note on the rôle of maximal consistency. *Erkenntnis*, 79(5):1015–1026, 2013.
- [75] Henryk Mehlberg. *The Reach of Science*. University of Toronto Press, Toronto, 1958. Extract from §29, pp. 256–259 reprinted under the title ‘Truth and Vagueness’ in [61] pp. 85–88.
- [76] Hung T. Nguyen and Elbert A. Walker. *A First Course in Fuzzy Logic*. Chapman and Hall/CRC, Boca Raton, FL, second edition, 2000.
- [77] Vilém Novák and Antonín Dvořák. Fuzzy logic: A powerful tool for modeling of vagueness. Technical Report 158, Institute for Research and Applications of Fuzzy Modeling, University of Ostrava, Ostrava, 2011.
- [78] Vilém Novák, Irina Perfilieva, and Jiří Močkoř. *Mathematical Principles of Fuzzy Logic*. Kluwer, Dordrecht, 2000.
- [79] Daniel N. Osherson and Edward E. Smith. On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9:35–58, 1981.
- [80] Daniel N. Osherson and Edward E. Smith. On typicality and vagueness. *Cognition*, 64:189–206, 1997.
- [81] Francesco Paoli. Truth degrees, closeness, and the sorites. To appear in Otavio Bueno and Ali Abasnezhad, editors, *On the Sorites Paradox*. Springer.
- [82] Francesco Paoli. A really fuzzy approach to the sorites paradox. *Synthese*, 134:363–387, 2003.
- [83] Charles Sanders Peirce. Vague. In James Mark Baldwin, editor, *Dictionary of Philosophy and Psychology*, page 748. Macmillan, New York, 1902.
- [84] Manfred Pinkal. *Logic and Lexicon: The Semantics of the Indefinite*, volume 56 of *Studies in Linguistics and Philosophy*. Kluwer, Dordrecht, 1995. Translated by Geoffrey Simmons.
- [85] Graham Priest. Fuzzy identity and local validity. *Monist*, 81:331–342, 1998.
- [86] Marian Przełęcki. Fuzziness as multiplicity. *Erkenntnis*, 10:371–380, 1976.
- [87] Hilary Putnam. The meaning of ‘meaning’. In *Mind, Language and Reality*, volume 2 of *Philosophical Papers*, pages 215–271. Cambridge University Press, Cambridge, 1975.
- [88] Willard Van Orman Quine. *Word and Object*. MIT Press, Cambridge, MA, 1960.
- [89] Diana Raffman. Vagueness without paradox. *Philosophical Review*, 103:41–74, 1994.
- [90] Diana Raffman. Vagueness and context relativity. *Philosophical Studies*, 81:175–192, 1996.
- [91] Diana Raffman. How to understand contextualism about vagueness: Reply to Stanley. *Analysis*, 65: 244–248, 2005.
- [92] David Ripley. Contradictions at the borders. In Rick Nouwen, Robert Van Rooij, Uli Sauerland, and Hans-Christian Schmitz, editors, *Vagueness in Communication*, pages 169–188. Springer, 2011.
- [93] David Ripley. Sorting out the sorites. In Koji Tanaka, Francesco Berto, Edwin Mares, and Francesco Paoli, editors, *Paraconsistency: Logic and Applications*, pages 329–348. Springer, 2013.
- [94] Bertil Rolf. Sorites. *Synthese*, 58:219–250, 1984.
- [95] Gideon Rosen and Nicholas J.J. Smith. Worldly indeterminacy: A rough guide. *Australasian Journal of Philosophy*, 82(1):185–198, 2004. Reprinted in [49] pp. 196–209.
- [96] David H. Sanford. Borderline logic. *American Philosophical Quarterly*, 12:29–39, 1975.
- [97] Uli Sauerland. Vagueness in language: The case against fuzzy logic revisited. In Petr Cintula, Christian G. Fermüller, Lluís Godó, and Petr Hájek, editors, *Understanding Vagueness: Logical, Philosophical and Linguistic Perspectives*, volume 36 of *Studies in Logic*, pages 185–198. College Publications, London, 2011.
- [98] Stephen P. Schwartz. Intuitionism versus degrees of truth. *Analysis*, 50:43–47, 1990.
- [99] Stewart Shapiro. *Vagueness in Context*. Oxford University Press, Oxford, 2006.
- [100] Edward E. Smith and Daniel N. Osherson. Compositionality and typicality. In Stephen Schiffer and Susan Steele, editors, *Cognition and Representation*, pages 37–52. Westview Press, Boulder, CO, 1988.
- [101] Nicholas J.J. Smith. Consonance and dissonance in solutions to the sorites. To appear in Otavio Bueno and Ali Abasnezhad, editors, *On the Sorites Paradox*. Springer.
- [102] Nicholas J.J. Smith. Undead argument: The truth-functionality objection to fuzzy theories of vagueness. To appear in *Synthese*, doi: 10.1007/s11229-014-0651-7.
- [103] Nicholas J.J. Smith. Vagueness and blurry sets. *Journal of Philosophical Logic*, 33:165–235, 2004.

- [104] Nicholas J.J. Smith. A plea for things that are not quite all there: Or, Is there a problem about vague composition and vague existence? *Journal of Philosophy*, 102:381–421, 2005.
- [105] Nicholas J.J. Smith. Vagueness as closeness. *Australasian Journal of Philosophy*, 83:157–183, 2005.
- [106] Nicholas J.J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, Oxford, 2008. Paperback 2013.
- [107] Nicholas J.J. Smith. Fuzzy logic and higher-order vagueness. In Petr Cintula, Christian G. Fermüller, Lluís Godo, and Petr Hájek, editors, *Understanding Vagueness: Logical, Philosophical, and Linguistic Perspectives*, volume 36 of *Studies in Logic*, pages 1–19. College Publications, London, 2011.
- [108] Nicholas J.J. Smith. Many-valued logics. In Gillian Russell and Delia Graff Fara, editors, *The Routledge Companion to Philosophy of Language*, pages 636–651. Routledge, London, 2012.
- [109] Nicholas J.J. Smith. Measuring and modeling truth. *American Philosophical Quarterly*, 49:345–356, 2012.
- [110] Nicholas J.J. Smith. Truthier than thou: Truth, supertruth and probability of truth. To appear in *Noûs*, doi: 10.1111/nous.12108.
- [111] Scott Soames. *Understanding Truth*. Oxford University Press, New York, 1999.
- [112] Roy Sorensen. *Blindspots*. Clarendon Press, Oxford, 1988.
- [113] Roy Sorensen. *Vagueness and Contradiction*. Oxford University Press, Oxford, 2001.
- [114] Gaisi Takeuti and Satoko Titani. Globalization of intuitionistic set theory. *Annals of Pure and Applied Logic*, 33:195–211, 1987.
- [115] Jamie Tappenden. The liar and sorites paradoxes: Toward a unified treatment. *Journal of Philosophy*, 90:551–577, 1993.
- [116] Michael Tye. Vague objects. *Mind*, 99:535–557, 1990.
- [117] Michael Tye. Sorites paradoxes and the semantics of vagueness. *Philosophical Perspectives*, 8: 189–206, 1994. Reprinted in [61] pp. 281–293.
- [118] Michael Tye. Vagueness: Welcome to the quicksand. *Southern Journal of Philosophy*, 33(Suppl.):1–22, 1995.
- [119] Alasdair Urquhart. Many-valued logic. In Dov M. Gabbay and Franz Guenther, editors, *Handbook of Philosophical Logic*, volume III, pages 71–116. D. Reidel, Dordrecht, 1986.
- [120] Robert Van Rooij. Vagueness and linguistics. In Giuseppina Ronzitti, editor, *Vagueness: A Guide*, pages 123–170. Springer, 2011.
- [121] Brian Weatherston. True, truer, truest. *Philosophical Studies*, 123:47–70, 2005.
- [122] Timothy Williamson. Vagueness and ignorance. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 66:145–162, 1992. Reprinted in [61] pp. 265–280.
- [123] Timothy Williamson. *Vagueness*. Routledge, London, 1994.
- [124] Crispin Wright. Further reflections on the sorites paradox. *Philosophical Topics*, 15:227–290, 1987.
- [125] Lotfi A. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965.
- [126] Lotfi A. Zadeh. A note on prototype theory and fuzzy sets. *Cognition*, 12:291–297, 1982.

NICHOLAS J.J. SMITH  
 Department of Philosophy  
 Main Quadrangle A14  
 University of Sydney  
 NSW 2006 Australia  
 Email: njjsmith@sydney.edu.au