Consonance and Dissonance
in Solutions to the Sorites

Nicholas J.J. Smith

to appear in Otavio Bueno and Ali Abasnezhad (eds)
On the Sorites Paradox (Springer)

1 Sorites

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[1] 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*:[2]

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:

---

[1] Except the last one, which has no object after it.

[2] $x_1, \ldots, x_n$ 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.

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]

3. $\therefore x_n$ is F.

A characteristic 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.[4] Where F is a vague predicate, on the other hand, the sorites argument will seem genuinely paradoxical: the premises 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.[5]

A requirement on any theory of vagueness is that it solve the sorites paradox. In the current literature on vagueness it is generally agreed that 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 isn't correct. The second task is to explain why the sorites reasoning is a paradox rather than a simple mistake. Thus, as well as locating the error in the argument, a 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.

In this paper I argue that there is a further constraint on responses to the second task: such responses should conform to the standard modus operandi in formal semantics, in which the semantic theory one develops is taken to be implicit in the ordinary usage of competent speakers. That is, it should not turn out that one's explanation of why ordinary speakers react to the sorites reasoning in the way they do depends on speakers *not* thinking that the semantics of vague language is governed by the theory one is advocating. As we shall see, out of the current main contenders for a theory of vagueness, only theories that posit degrees of truth can meet this further constraint of *consonance* between the theory of vagueness being advocated and the theory implicit in ordinary speakers' behaviour.

---

[3]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 premises each of which makes a particular claim. For simplicity, I shall generally write below of 'the second premise'. Relative to formulations involving a multitude of premises 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.

[4]We don't need to know which object it is: the point is just that there is such an object.

[5]This is not the only difference between vague and precise predicates. For more on the issue of defining vagueness see Smith [2008, ch.3].

The formal details behind the theories to be discussed are sketched in §2. In the remainder of the paper, solutions to the sorites following from the various theories are discussed and it is shown that only theories that posit degrees of truth yield solutions that meet the consonance constraint.

2   Theories

In this section I sketch some formal theories.[6] Throughout we consider a standard first order language $\mathscr{L}$ with individual constants $a, b, c, \ldots$ and predicates $P, Q, R, \ldots$ of each arity—although for the sake of an uncluttered presentation we shall generally mention only connectives and one-place predicates explicitly.

A *classical valuation* $\mathfrak{V}_c$ of $\mathscr{L}$ comprises:

- a set D (the *domain*)

- (where $\mathscr{I}$ is the set of individual constants of $\mathscr{L}$:)

  a function from $\mathscr{I}$ to D (assigning a *referent* to each individual constant)

- (where $\mathscr{P}$ is the set of one place predicates of $\mathscr{L}$:)

  a function from $\mathscr{P}$ to $\{0, 1\}^D$ (assigning an *extension* to each predicate: an extension is a total function from the domain to $\{0, 1\}$; objects sent to 1 are in the extension and objects sent to 0 are not in the extension).

A classical valuation can be extended to a classical *model* $\mathfrak{M}$ using the standard classical rules. In particular:

- the truth value of an atomic wff Pa is the value (1 or 0) to which the extension of P sends the referent of a

- the truth values of negations, conjunctions, disjunctions and so on are determined by the classical truth tables.

A classical model assigns a truth value (1 or 0) to each closed wff of $\mathscr{L}$.

In *epistemicism*, a language $\mathscr{L}$ is associated with a unique classical model, which we call the *intended model*.

In (classical) *plurivaluationism*, a language $\mathscr{L}$ is associated with a nonempty set of classical models. We call the members of this set the *acceptable models*.

---

[6]The main purpose of this section is to fix terminology—not to give an introduction of these views suitable for someone completely unfamiliar with them. For a full introduction and references to the relevant literature see Smith [2008, ch.2].

A *partial valuation* $\mathfrak{V}_p$ of $\mathscr{L}$ is just like a classical valuation except that the extension of a predicate can be a *partial* function from the domain to $\{0,1\}$. Objects sent to 1 are in the extension; objects sent to 0 are outside the extension; objects sent nowhere are neither in nor out.

A *three-valued valuation* is just like a classical valuation except that the extension of a predicate is a total function from the domain to $\{0, *, 1\}$, where $*$ is a third truth value, in addition to 1 and 0. Everything that we say below in terms of partial valuations can be reformulated in terms of three-valued valuations. I shall use the term 'tripartite' when I wish to speak generally of three-valued and partial two-valued setups.

A classical valuation $\mathfrak{V}_c$ *extends* a partial valuation $\mathfrak{V}_p$ if:

- $\mathfrak{V}_p$ and $\mathfrak{V}_c$ have the same domain

- $\mathfrak{V}_p$ and $\mathfrak{V}_c$ assign the same referents to names

- for all predicates P and for all x in the domain (and where P_p is the extension of P on $\mathfrak{V}_p$ and P_c is the extension of P on $\mathfrak{V}_c$):

    - if $P_p(x) = 1$ then $P_c(x) = 1$
    - if $P_p(x) = 0$ then $P_c(x) = 0$

    (i.e. the extension of P on $\mathfrak{V}_c$ just closes the gaps in the extension of P on $\mathfrak{V}_p$: it does not move anything from in (1) to out (0) or vice versa).

We call a classical model an *extension* of $\mathfrak{V}_p$ if it is determined (using the standard classical rules) by a classical valuation $\mathfrak{V}_c$ that extends $\mathfrak{V}_p$.

There are various ways of extending a partial valuation to a model, which assigns a truth value or a gap (i.e. no value) to each closed wff of $\mathscr{L}$. One way—call it the *recursive* way—mirrors the classical story as closely as possible:

- the truth value of an atomic wff *Pa* is the value (1, 0 or no value) to which the extension of P sends the referent of *a*

- the classical truth tables are extended to cover the cases where one or more component proposition has no value. Perhaps the best known option here is Kleene's strong tables—but there are many other options.[7]

Another way of extending a partial valuation to a model is the supervaluationist route. Given a partial valuation $\mathfrak{V}_p$, the associated *supervaluation* is a partial function v from closed wffs α to $\{0,1\}$ defined as follows:

$$v(\alpha) = \begin{cases} 1 & \text{iff } \alpha \text{ has the value 1 on every extension of } \mathfrak{V}_p \\ 0 & \text{iff } \alpha \text{ has the value 0 on every extension of } \mathfrak{V}_p \end{cases}$$

---

[7]See Smith [2012] for an introduction to some of the main ones.

Thus $v(\alpha)$ is undefined if α has the value 1 on some extension of $\mathfrak{V}_p$ and has the value 0 on some (other) extension of $\mathfrak{V}_p$.

Consider atomic wffs such as Pa for a moment. They can be assigned a value by the supervaluation in the way just specified—or alternatively (as in the recursive way of extending a partial valuation to a model) they can be assigned a value with reference only to the partial valuation (not its extensions): the truth value of an atomic wff Pa is the value (1, 0 or no value) to which the extension of P sends the referent of a. Both options yield the same values for atomic wffs.

One can refine the supervaluationist view by considering not *all* extensions of $\mathfrak{V}_p$ but only those satisfying certain conditions. These extensions are then called the *admissible extensions* and in the definition of the supervaluation v, 'extension' is replaced by 'admissible extension'. We shall use $E(\mathfrak{V}_p)$ to denote the set of all admissible extensions of $\mathfrak{V}_p$. The unrefined version of supervaluationism presented above is the special case of the refined version where every extension is admissible.

A variant of the supervaluationist framework is the degree-theoretic form of supervaluationism. Here we suppose there to be a normalised measure function μ on the powerset of $E(\mathfrak{V}_p)$. Where A is the set of admissible extensions on which α has the value 1, the supervaluation v is then defined thus:

$$v(\alpha) = \mu(A)$$

It may be possible to define μ only on some σ-field of subsets of $E(\mathfrak{V}_p)$, not on the full powerset. In that case v is defined only for α such that A is a measurable set. Decock and Douven have recently proposed a particular implementation of the degree-theoretic form of supervaluationism in the framework of conceptual spaces.[8] They show that the values assigned by the supervaluation v in their account behave in the same ways as the *verities* of Edgington [1997] and are formally probabilities.

A *fuzzy valuation* is just like a classical valuation except that the extension of a predicate is a total function from the domain to $[0, 1]$, the set of all real numbers between 0 and 1 inclusive—known in this context as *degrees of truth*.[9] The *recursive* way of extending a fuzzy valuation to a fuzzy model, which assigns a degree of truth to each closed wff of $\mathscr{L}$, follows the same general plan as the classical way of extending a valuation to a model. First, the truth value of an atomic wff Pa is the value to which the extension of P sends the referent of a. Second, we want to fix truth values for formulas containing connectives. Earlier, when describing

---

[8]See Decock and Douven [2012], Douven and Decock [2014] and Decock and Douven [2014]. This works builds on Douven et al. [2013].

[9]$[0, 1]$ is a standard choice of degrees of truth in fuzzy logics but not the only possible choice. For some other options see Smith [2015].

the classical case, we spoke of truth tables—but we could also spell out the classical story in a slightly different way. Let $[\alpha]$ be the truth value of α (on some model). For each n-place connective $\triangledown$ we specify a corresponding operation $\blacktriangledown$ (of arity n) on the classical truth values and then stipulate that

$$[\triangledown(\alpha_1, \ldots, \alpha_n)] = \blacktriangledown([\alpha_1], \ldots, [\alpha_n])$$

For example, corresponding to the negation connective $\neg$ we have the unary operation $1 - x$ 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 $[\alpha \vee \beta] = \max\{[\alpha], [\beta]\}$. This alternative way of specifying the move from classical valuations to models ends up, of course, with exactly the same truth conditions. Now in the fuzzy case, we can proceed in an analogous way: we specify an operation on the set $[0, 1]$ of degrees of truth corresponding to each connective. Here there are many live options, of which we shall mention a few of the most prominent.

First consider Zadeh logic (Figure 1). Negation, conjunction and disjunction

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

Figure 1: Zadeh logic

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.

Second, consider what we may call Philosophers' Fuzzy Logic or PFL (Figure 2).[10] 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.

---

[10]The reason for the name is that many philosophers write as if 'fuzzy logic' just *is* PFL.

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

Figure 2: Philosophers' Fuzzy Logic

Third, consider t-norm fuzzy logics. A t-norm is a binary function $\wedge$ on $[0, 1]$ satisfying the conditions shown in Figure 3. A t-norm logic is specified by picking

commutativity:	$x \wedge y$	$= y \wedge x$
associativity:	$(x \wedge y) \wedge z$	$= x \wedge (y \wedge z)$
non-decreasing in 1st argument:	$x_1 \leq x_2$	$\Rightarrow x_1 \wedge y \leq x_2 \wedge y$
non-decreasing in 2nd argument:	$y_1 \leq y_2$	$\Rightarrow x \wedge y_1 \leq x \wedge y_2$
unit:	$1 \wedge x$	$= x$

Figure 3: Conditions on t-norms

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:[11]

$$
x \to y = \max\{z : x \wedge z \leq y\}
$$

and the negation the precomplement of the conditional:

$$
\neg x = x \to 0
$$

Figures 4, 5 and 6 show 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.[12]

---

[11]The residuum exists iff the t-norm is left-continuous.

[12]So in Gödel logic, there is no difference between the strong and weak conjunction. Note that the min conjunction of Figure 2 is definable using the operations of Figure 4 as $\alpha \wedge (\alpha \to \beta)$ *and* the Łukasiewicz t-norm conjunction of Figure 4 is definable using the negation and conditional of Figure 2 (which are the same as the negation and conditional of Figure 4) as $\neg(\alpha \to \neg\beta)$. So 'PFL'

$$x \wedge y \quad = \quad \max(0, x + y - 1)$$
$$x \rightarrow y \quad = \quad \begin{cases} 1 & \text{if } x \leq y \\ 1 - x + y & \text{if } x > y \end{cases}$$
$$\neg x \quad = \quad 1 - x$$

Figure 4: Łukasiewicz logic

$$x \wedge y \quad = \quad \min(x, y)$$
$$x \rightarrow y \quad = \quad \begin{cases} 1 & \text{if } x \leq y \\ y & \text{if } x > y \end{cases}$$
$$\neg x \quad = \quad \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{otherwise} \end{cases}$$

Figure 5: Gödel logic

$$x \wedge y \quad = \quad x \cdot y$$
$$x \rightarrow y \quad = \quad \begin{cases} 1 & \text{if } x \leq y \\ y/x & \text{if } x > y \end{cases}$$
$$\neg x \quad = \quad \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{otherwise} \end{cases}$$

Figure 6: Product logic

Once we settle on a choice of operations on $[0, 1]$ corresponding to the connectives, we fix a particular kind of fuzzy model. From there, we still have two choices for a theory of vagueness—corresponding to the choice between epistemicism and (classical) plurivaluationism, both of which are built on the (same) notion of a classical model. The *basic fuzzy theory of vagueness* holds that a language $\mathscr{L}$ is associated with a unique fuzzy model, which we call the *intended model*. This view stands to fuzzy models (of a certain sort) as the epistemicist view stands to classical models. In *fuzzy plurivaluationism*, a language $\mathscr{L}$ is associated with a nonempty set of fuzzy models; we call the members of this set the *acceptable models*. This view stands to fuzzy models (of a certain sort) as (classical) plurivaluationism stands to classical models.

3 Dissonance

According to the epistemicist, vague predicates—like precise ones—have crisp sets as their extensions. This is not to say that there is no difference 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.

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 find the argument compelling—is more subtle. 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. 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* had sharp but unknowable boundaries. Then she would not think for a moment that the second premise of a sorites argument for *P*

---

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).

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. 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. Yet 'bearfast' 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 for a second that the second premise of the sorites argument is *true*.

Thus the epistemicist solution to the sorites exhibits dissonance: if ordinary speakers accepted that vague predicates work the way epistemicists say such predicates work, they would not behave as they do around vague predicates. In particular, they would regard sorites arguments for vague predicates as obvious mistakes rather than genuine paradoxes.

Let's turn now to supervaluationism. 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 is assigned the value 0 by the supervaluation. The second part of the solution—saying why we find the argument compelling—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. 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. 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' (for each x) 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 Φ' if Φ 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

10

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.

Now let's consider (classical) plurivaluationism. When it comes to truth, the plurivaluationist can only say, of each acceptable model, whether a statement is true or false on that model. (Unlike in supervaluationism, where we also have a supervaluation and a tripartite model, in plurivaluationism there is no further semantic machinery beyond the individual classical models.) However the plurivaluationist can say something about assertibility that is analogous to what the supervaluationist says about truth. For the supervaluationist, a statement that is true on all admissible extensions of the tripartite valuation is assigned the value 1 by the supervaluation. 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 argument turns on their not realising that the semantics of their discourse is plurivaluationist. For consider an analogous case. We are canvassing opinions about (say) football and (suppose for the sake of argument) there is an accepted convention that one can say 'the man in the street believes Φ' only when *everyone* canvassed believes Φ. 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).

Now let's consider a view we haven't yet mentioned: *contextualism*.[13] 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: each admissible extension corresponds to a legitimate way of completely precisifying the language. 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 time 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 contextualists differ—but typically the basic 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 valuations. 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.

Contextualists think that this combination of a three-valued approach within each model with a dynamic story about how acts of stipulation change which tripartite 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

---

[13]Again, for a more detailed introduction and references to the relevant literature see Smith [2008, ch.2]. Note that there are various contextualist positions in the literature which differ in sometimes subtle but significant ways.

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).

Once again, however, this solution to the sorites turns on speakers not believing that the contextualist story is the correct account of vague language. 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.

4 Consonance

A crucial advantage of fuzzy theories of vagueness is that they 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. It is not assumed that speakers have a full understanding of the semantics of their language, but it is assumed that they have some sort of implicit grasp of it, which manifests itself in their behaviour. That is why such 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 in the previous section—the kind of position that other 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 premises to the conclusion; and yet they find the conclusion unacceptable. However, a speaker who thought that epistemicism, say, gives the correct semantics of the predicate P would have no inclination at all to think that the second premise was true. 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, then they would respond to a sorites argument for P in just the way that they *do* ordinarily respond to sorites paradoxes.[14]

---

[14] As we shall see, the key to the solution is being able to say that certain sentences are true to a degree very close to 1 (complete truth). So doesn't the degree-theoretic form of supervaluationism share this advantage with fuzzy theories? No: because (as I argue in a separate paper [Smith,

13

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.

$\vdots$

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

$\therefore$ 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.[15]

Note that the explanation of why the sorites is compelling is sometimes put in terms of ordinary speakers being fooled—of their *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 previous paragraph, we do not suppose that speakers mistakenly think 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

---

2014]) the degrees posited by the degree supervaluationist cannot be degrees of *truth*—and indeed quite generally, the values assigned by a supervaluation cannot be *truth values*.

[15]Cf. [Forbes, 1983, pp.243–4], [Forbes, 1985, pp.171–2], [Williamson, 1994, pp.123–4], [Paoli, 2003, p.365] and Paoli [2014].

the statement—at least in normal contexts and until trouble was seen to arise. In everyday contexts we *round* values up or down. This is not sloppy: it is useful. Occasionally it leads to trouble—but in those cases we just refrain from rounding!

Consider an example. We do not believe that dust particles are weightless: we believe that the weight of a dust particle is negligible. But this very belief licenses us to accept the claim that dust particles are weightless as a useful approximation of our real belief, in ordinary circumstances. 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. Nevertheless, the claim that dust particles are weightless is revealed as merely a useful 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 very very small.

Similarly, the view being proposed about the sorites is not that speakers really believe (mistakenly, given the fuzzy semantics) that the sorites conditionals are fully true. Rather, they believe that they are very nearly true, and so in accordance with a *generally useful* practice of rounding up or down—of ignoring very small differences—they go along with them. Of course when they later see where the argument leads, they baulk. This behaviour is exactly what we would predict of a speaker who implicitly adopts the fuzzy semantic theory currently under consideration. Hence the fuzzy solution to the sorites exhibits consonance, not dissonance.

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 premises. 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.[16] Or suppose we read 'if. . . then. . . ' in the argument as (say) the Zadeh—or Gödel—conditional: in that case some premises would have degrees of truth of around 0.5—or even less—and so again we lose the explanation given above of why the argument is compelling. (There is no generally useful practice of ignoring *large* differences: of rounding *that* far.)

---

[16]See [Machina, 1976, pp.69–75], [Williamson, 1994, pp.123–4] and [Priest, 1998, p.332].

So far we have considered just one formulation of the sorites argument: one involving multiple conditional premisses. 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. Wright [1987, pp.251–2] took this to be a problem for fuzzy views because he thought that they could not handle conjunctive formulations of the paradox—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).
>
> $\vdots$
>
> It is not the case that (pile 2 is a heap and pile 1 is not a heap).
>
> $\therefore$ Pile 1 is a heap.

However, Wright was assuming that fuzzy logic is PFL, in which only min conjunction is considered. Using min conjunction, 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 if we move from the more limited PFL perspective to the wider perspective of 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)' [Paoli, 2014]. 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 resources. Wright's point does not count against fuzzy approaches in general: it counts only in favour of some fuzzy approaches over others.[17]

5 Conclusion

It is widely recognised that in order to solve the sorites paradox we need to do more than simply develop a logic or semantics for vagueness on which the argument is invalid or unsound. In addition to showing what is wrong with the argument (e.g. where the reasoning fails to be truth-preserving, or which premise fails to be true) we must explain why ordinary speakers find the argument compelling: we must explain why the sorites is a paradox and not a simple mistake.

---

[17]There is a second way of solving the sorites within a fuzzy framework that does not depend on any particular choice of fuzzy truth functions—and which also exhibits consonance. This solution is presented in Smith [2008]; for the argument that it exhibits consonance see Smith [2015].

I have argued that in giving this explanation we should furthermore conform to the standard modus operandi of formal semantics, where we assume that speakers have some sort of implicit grasp of the semantics of their language, which manifests itself in their behaviour. Of course we do not assume they have a full grasp of the semantics: but we do assume that they do not have utterly mistaken views. In particular, we should not 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. In that case their linguistic behaviour would provide no evidence at all for semantic theory T—whereas the standard modus operandi takes the behaviour of ordinary speakers as data for the formulation of semantic theories. Thus, a desired feature for a semantic theory T for vague language is that it satisfy the following schema:

> If speakers thought that theory T gave the correct semantics for a vague predicate P, then they would respond to a sorites argument for P in just the way that they *do* ordinarily respond to sorites paradoxes.

I have argued that amongst existing solutions to the sorites, only those based on degrees of truth are in a position to meet this requirement. Other theories of vagueness involve the idea that ordinary speakers *make a mistake* when confronted with sorites reasoning. Thus there is dissonance between the views of the semantic theorist and the views attributed to ordinary speakers. The fuzzy theory, on the other hand, can solve the sorites in a way that exhibits consonance between the semantics proposed by the theorist and the semantics of which an implicit, rough, inchoate grasp is assumed to guide the behaviour of ordinary speakers.[18]

References

Lieven Decock and Igor Douven. What is graded membership? *Noûs*, 2012. doi: 10.1111/nous.12003.

Lieven Decock and Igor Douven. Two recent degree-theoretic approaches to the sorites paradox. 2014.

Igor Douven and Lieven Decock. What verities may be. 2014.

Igor Douven, Lieven Decock, Richard Dietz, and Paul Égré. Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic*, 42(1):137–160, 2013.

Dorothy Edgington. Vagueness by degrees, 1997. In Keefe and Smith [1997], pages 294–316.

---

[18]Many thanks to Francesco Paoli for very helpful comments on a draft version.

Graeme Forbes. Thisness and vagueness. *Synthese*, 54:235–59, 1983.

Graeme Forbes. *The Metaphysics of Modality*. Clarendon Press, Oxford, 1985.

Rosanna Keefe and Peter Smith, editors. *Vagueness: A Reader*. MIT Press, Cambridge MA, 1997.

Kenton F. Machina. Truth, belief, and vagueness. *Journal of Philosophical Logic*, 5:47–78, 1976. Reprinted in Keefe and Smith [1997], pages 174–203.

Francesco Paoli. A really fuzzy approach to the sorites paradox. *Synthese*, 134: 363–387, 2003.

Francesco Paoli. Truth degrees, closeness, and the sorites. 2014.

Graham Priest. Fuzzy identity and local validity. *Monist*, 81:331–42, 1998.

Nicholas J.J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, Oxford, 2008. Paperback 2013.

Nicholas J.J. Smith. Many-valued logics. In Gillian Russell and Delia Graff Fara, editors, *The Routledge Companion to Philosophy of Language*, pages 636–51. Routledge, London, 2012.

Nicholas J.J. Smith. Truthier than thou: Truth, supertruth and probability of truth. 2014.

Nicholas J.J. Smith. Fuzzy logics in theories of vagueness. In Petr Cintula, Christian Fermüller, and Carles Noguera, editors, *Handbook of Mathematical Fuzzy Logic - Volume 3*, Studies in Logic, Mathematical Logic and Foundations. College Publications, London, 2015.

Timothy Williamson. *Vagueness*. Routledge, London, 1994.

Crispin Wright. Further reflections on the sorites paradox. *Philosophical Topics*, 15:227–90, 1987.