HISTORY AND PHILOSOPHY OF LOGIC, **31** (May 2010), 177–184

Taylor & Francis
Taylor & Francis Group

Essay Review

DOMINIC HYDE, *Vagueness, Logic and Ontology*. Aldershot:Ashgate, 2008. xii + 226 pp. £55. ISBN 978-0-7546-1532-3.

NICHOLAS J. J. SMITH, *Vagueness and Degrees of Truth*. Oxford: Oxford University Press, 2008. vii + 341 pp. £40. ISBN 978-0-19-923300-7.

Reviewed by GRAHAM PRIEST. *Departments of Philosophy, University of Melbourne, University of St. Andrews, and the Graduate Center, City University of New York*

The problem of vagueness is a venerable one. The sorites paradox was discovered by the Megarians. Oddly, though, it lay dormant till recently. There is just a little discussion of it in Ancient logic, nothing at all in Medieval logic (to my knowledge); and, with a couple of odd exceptions, nothing in the contemporary period before about 1960. Then, after a couple of papers, it took off like a rocket. This is not because the topic is of some contemporary fad: indeed, the problems vagueness poses are pressing, deep, and hard. The question is why they were ignored for so long. I leave that for historians to worry about. It is, at any rate, now a very hot topic, and the books by Hyde and Smith are two recent and notable contributions to the debate. In what follows, I will first discuss Hyde's book, then Smith's.

Hyde's book develops a distinctive account of vagueness, in the process providing a careful review of alternative views. The discussions are always thoughtful, and show an impressive knowledge of the relevant literature. I learned much in the process of reading it. The pillars of Hyde's own view are two. First, he argues that vagueness is not merely semantic, but *de re*. In particular, there are vague objects in the world. Secondly, he argues for a truth-functional account of the semantics of vague language, endorsing, essentially, the Łukasiewicz logic $Ł_3$.

Chapter 1 is a first look at the phenomenon of vagueness, seeking a definition: it finds one in the existence of appropriate borderline cases which are not merely epistemic. The next chapter is a discussion of Russell's account of vagueness. Though many people have rejected Russell's view, Hyde defends it as an adequate account of vagueness *if* all vagueness is in representations – that is, if the world is precise. A natural thought at this point is that if the world itself is precise, we can, at least for theoretical purposes, reject vague language, and describe it in a precise language. Chapter 3 takes this view to task. Given, then, that we are stuck with language that is vague, what is its logic? Chapter 4 takes up the issue, arguing against accounts of this according to which super- or sub-valuation is employed on top of a three-valued logic. The duality between taking the middle value of the three-valued logic to be a truth value gap (and so deploying super-valuation) and taking it to be a truth value glut (and so deploying sub-valuation) is well brought out, showing that, as far as vagueness goes, there is little to choose between them. (Despite this, when Hyde goes

History and Philosophy of Logic ISSN 0144-5340 print/ISSN 1464-5149 online
http://www.informaworld.com DOI: 10.1080/01445341003594099

on to develop his own account, he settles for an understanding of the third value, 0.5, as a gap, on the ground of fighting one battle at a time. I found this a pity, since Hyde himself was one of the first people to put a paraconsistent account of vagueness on the table, an account which has recently been receiving growing attention.)

Having cleared the decks, the rest of the book develops Hyde's own account of vagueness. In chapter 5 it is argued that the world is not precise: there is *de re* vagueness. An object, a, is *de re* vague if there are objects (spatial, temporal, physical, and maybe of other kinds) such that there is no fact of the matter whether they are parts of a; in other words, there are statements of the form 'x is part of a' which have the value 0.5. Hyde argues (convincingly for me) that there are vague objects, and rejects some well-known objections. In particular, in 5.2 there is a careful analysis and rejection of the notorious 'Evans' argument'.

Evans' argument starts from the supposition that for certain objects, a and b, we have $\nabla a = b$ and $\neg\nabla a = a$, where ∇ is read as 'It is indeterminate that'. It then uses λ-abstraction to infer that $\lambda x(\nabla a = x)b$ and $\neg\lambda x(\nabla a = x)a$, and then $a \neq b$ on the basis of the principle of substitutivity of identicals, in the form $\forall P(a = b \rightarrow (Pa \leftrightarrow Pb))$. A central part of Hyde's analysis of the argument is to reject the λ-abstraction involved, on the ground that the condition $\nabla a = x$ does not express a property. This, I think, is not the optimal strategy. It raises the difficult question of what conditions do specify properties, an issue that is not resolved when Hyde specifies the formal semantics in Chapter 7, since he does not discuss second-order quantification. Nor does it address what to say about the validity of the plain vanilla version of the argument using just the two initial premises and substitutivity in schematic form, $a = b \rightarrow (A_x(a) \leftrightarrow A_x(b))$ (where $A_x(c)$ is the result of substituting c for free occurrences of x in A). The semantics of identity is also not explicitly addressed in Chapter 7; but I think that it is clear that the identity predicate, like all predicates, is to have an extension (comprising those things of which it is true), and a non-overlapping co-extension (comprising those things of which it is false). Now, provided that the extension of $=$ is $\{\langle x, x\rangle : x \in D\}$, where D is the domain of objects, we have $a = b \models A_x(a) \leftrightarrow A_x(b)$ but we do not have $\neg A_x(a), A_x(b) \models a \neq b$. Indeed, assuming that the value of ∇A is 1 if the value of A is 0.5, and 0 otherwise, Evans' argument is demonstrably invalid. Let a and b denote x and y, respectively; where $\langle x, y\rangle$ is in neither the extension nor the co-extension of $=$. Then $\neg\nabla a = a$ and $\nabla a = b$ both have value 1, even though $a \neq b$ has value 0.5.

Chapter 6 of the book contains a very interesting discussion of the relatively neglected but important issue of the counting of vague objects. Hyde's account is simple and natural: just cash out numerical statements in terms of first-order quantification in the standard way, and let the semantics take care of the rest.

Chapter 7 finally specifies Hyde's preferred semantics and logic. (I personally would have found it helpful to have this done before Chapter 5.) In 7.2 he defends the use of Ł$_3$ against objections to the effect that this does not respect 'penumbral connections'. Some of these concern conjunction and disjunction. Thus, for example, suppose that we have a sorites sequence of objects $a_0, \ldots, a_{100}$, where, say, a_0 is blue, and a_{100} is not blue. Let B be 'is blue', and let a_i be in the indeterminate area. Then, on the truth-functional account $Ba_i \wedge \neg Ba_i$ is itself indeterminate, not false as one might have expected. Hyde parlays this objection to a standstill. But I was less persuaded when the connections concerned the conditional. If a_{i+1} is also in the indeterminate area, (i) $\neg Ba_i \rightarrow \neg Ba_{i+1}$ and (ii) $\neg Ba_i \rightarrow Ba_{i+1}$ are both true. This seems wrong. Certainly, (i) seems true: if a_i is not blue then a_{i+1} is not, since it is less blue. But for the same reason (ii) is *not* true. Hyde suggests that the oddness of (ii) can be accounted for by the fact that, though true, it is unassertable, since its truth derives from the indeterminacy of Ba_i and Ba_{i+1}, which is stronger information, and so makes (ii) unassertable for reasons of conversational implicature. But if that were so, we would find a similar oddity with (i), which we do not.

The problem can be solved by moving to an account of the conditional that is at least modal (and many-valued).[1] Let $\rightarrow$ have the truth conditions of some sort of strict conditional, and let R be the accessibility relation. Suppose that wRw' iff the extension and co-extension of B at w are 'stretched out a bit further towards the middle' at w'. Then (i) will indeed come out to be true, but (ii) will come out to be false if there are accessible worlds where Ba_i is true and Ba_{i+1} is false.

Chapter 7 finishes with a discussion of the vexed issue of 'higher-order vagueness'. I found this chapter the most unsatisfying part of the book. The thing that makes a two-valued account of vagueness seem wrong is that in a standard sorites progression one can locate no exact point at which values go from *true* to *false*. In the three-valued approach, this fact is accommodated, since there is, indeed, no such point. But exactly the same problem arises again. One can locate no exact point at which truth values go from *true* to *neither true nor false* or *neither true nor false* to *false*. This is one way of stating the problem of higher-order vagueness as it applies to Hyde's solution.

The way out, says Hyde, is simply to iterate the construction: it should be possible for claims about whether or not something is indeterminate in truth value to themselves be indeterminate in truth value. This is a plausible enough thought. The problem is how to execute it. The devil is in the details. In fact, Hyde suggests that (a) the metalanguage itself should be three-valued and (b) $Ł_3$ should be generalised to an infinite-valued logic to model the discriminations required; he does not spell out the details of this, but I take it to be $Ł_{\aleph}$. As far as I can tell, Hyde seems to think that these are part of the same picture. They are not. A three-valued logic is still a three-valued logic, even if the underlying logic of the theory in which it is couched is itself three-valued. Using $Ł_{\aleph}$ will deliver theorems of the form $\exists v(v(p) = 0.75)$. This is not a theorem of the three-valued semantics couched in classical logic; even less couched in $Ł_3$. However, they are both plausible ways of trying to cash out the basic idea. I do not think that (b) will work, for familiar reasons. In a standard sorites, the truth values of the relevant sentences start at 1, then fall monotonically, till they reach 0. There is, then, a last point at which the value is 1. Yet the existence of such a point is just as counter-intuitive as the existence of any other precise cut-off point. (To the extent that Hyde has an answer to this, it would seem to be that we just have to accept it (p. 208). This does not seem very helpful.) The situation with suggestion (a) is less clear, and this is largely because how to proceed is itself less than clear. Hyde does not help us out in the matter. What is necessary is an axiomatic specification of the metatheory to which we can apply the deductive machinery of $Ł_3$. Since a full theory of validity is couched in set theory, we need a set theory based on $Ł_3$. How much of the standard picture about what is and is not valid in $Ł_3$ will be forthcoming in the new context is one of the first issues that needs to be addressed. But perhaps something more modest is sufficient for addressing the worry. We do not need a theory of truth-in-an-evaluation; simply a theory of truth. We would expect axioms

[1] See ch. 11a of G. Priest, *Introduction to Non-Classical Logic: from If to Is*, Cambridge: Cambridge University Press, 2008 [hereafter, INCL].

such as $\neg\exists x(Tx \wedge Fx)$, where T is the truth predicate and F the falsity predicate, and if Ix is $\neg Tx \wedge \neg Fx$, we would also expect to have $\exists x Ix$. Assuming some such theories to have been set up, it is not at all clear that higher-order worries are avoided. Take our blue sorites series above. How is the state of affairs concerning this to be described in the metalanguage? A natural suggestion is that the theory will specify $T\langle Ba_i\rangle$ for some of the a_i in the sequence, but not for others. But in that case, there will be a last i such that $T\langle Ba_i\rangle$ holds in the description of the situation, and we have our old problem back. Maybe there are things that can be said to address the problem. Maybe not. I think that the best one can say about this strategy is that a lot more work needs to be done on it before its viability can be assessed properly.

Smith's book, to which I now turn, provides a rather different account of vagueness. The book is clear, focussed, technically deft, and has impressive vision. Like Hyde, Smith argues that vagueness is in the world, but he offers a fuzzy-logic account of vagueness, based, essentially, on $Ł_{\aleph}$.

Part 1 of the book lays the foundations for what is to come. Chapter 1 describes a general semantical framework into which all standard theories fit. A semantics is provided by a domain of objects and a space of truth values with suitable operators. An interpretation of the language maps (amongst other things) predicates to functions from (n-tuples of) objects in the domain to truth values – call these *predicate values*. Smith identifies vagueness in the world with a non-classical set of truth values of a certain kind, and vagueness in language with the existence of multiple acceptable interpretation functions. I think that it might have been preferable to take interpretations to comprise, as well as a set of objects, a set of potential predicate values. (So that interpretations map names to objects and predicates to one of these.) This makes it clear that there are vague properties *in re*. (Even one who thinks that the world is precise, can, after all, admit the existence of non-standard truth values: they just don't get used!) In a sense, the extra machinery is redundant, but conceptually things are clearer.

Against this background, Chapter 2 locates the standard approaches to vagueness. An important distinction is drawn between two views that are normally run together: supervaluationism, and what Smith calls *plurivaluationism*. In both, there is a non-classical set of truth values (*not* necessarily three-membered), and a bunch of admissible interpretation functions. In supervaluationism, only one of these is the intended interpretation. Super-truth conditions of formulas are given non-recursively in terms of the admissible classical precisifications of the intended interpretation. The notion of validity that goes naturally with this is the preservation of super-truth. By contrast, in plurivaluationism, all the admissible interpretations are intended interpretations, and truth conditions are given recursively for each interpretation. The notion of validity that goes naturally with this is preservation of truth – or more generally, designation – in all such interpretations.

Part 2 of the book argues for a fuzzy semantics (i.e., a semantics where the truth values are the real numbers between 0 and 1, inclusive) in a novel way. First, in Chapter 3, a definition of vagueness is mounted and defended. Essentially, a predicate, F, is vague[2] iff the *Closeness Condition* is satisfied: whenever a and b are very close in F-relevant aspects then 'Fa' and 'Fb' are close (NB, *not* necessarily identical) in respect of truth. The analysis is, it seems to me, very plausible.

[2] On pp. 158–159, Smith argues that all vagueness reduces to predicate vagueness. I doubt this: I do not see how the vagueness of quantifiers, like 'most', can be so reduced.

Continuity (though not necessarily in a technical sense) does seem to go hand-in-glove with vagueness. Chapter 4 then argues that the only account of vague semantics which fits this definition of vagueness is a fuzzy one. (Smith does point out that a very large finite number of truth values might do in some cases. But having an infinite number of values makes things uniform.) Obviously, the argument puts a lot of weight on the notions of *F-relevance* and *closeness*, and much time is spent articulating these notions. Though what is said is very plausible, I think that the defenders of other views may well find some wiggle-room here. Consider someone, for example, who takes there to be three truth values, *true* (*only*), *false* (*only*), and *both*; where, in a sorties progression, the statements start off as *true*, become *both* in the middle range, and then *false*. The Closeness Condition might well be taken to fail at the boundaries. But let *a* and *b* be points close to each other, and on either side of the boundary between *true* and *both*. Then *Fa* and *Fb* are, in fact, close with respect to truth: they are identical! Of course, they are not close with respect to falsity. It will be pointed out that at the other boundary between the *boths* and the *falses*, the relevant statements will be identical with respect to falsity but *not* truth. Fair enough. But we may take 'close in respect of truth' to mean identical in truth value *or* identical in falsity value. Smith, no doubt, would insist on a conjunction, rather than a disjunction. But I am not sure that there are non-question-begging grounds for this. Smith would say, I am sure, that in the three-valued cases there is a jolt (to use his nice phrase) at each boundary; but equally, there is an intuitive jolt, in the continuum-valued case, between having a value $= 1$ and having a value < 1; in both cases it seems to be impossible to locate anywhere where the semantic transition can be said to occur in any but an arbitrary fashion. Smith says (p. 192) that the location problem, and the 'jolt' problem are really distinct problems; but I doubt this. If we could locate a relevant transition point, then we would accept a jolt there quite happily. Conversely, if we cannot locate a relevant transition point, then any point where such is supposed to be located will feel like a jolt: a punctual change where none is indicated.

In Part 3 of the book, Smith defends his account against objections. Higher-order vagueness is left for Chapter 6; Chapter 5 deals with the others. Here we find discussions of validity, penumbral connections, assertability, degrees of belief, and many other interesting topics.

Smith defines an inference to be valid if, in all interpretations, when all the premises have truth value *greater than* 0.5, the conclusion has value *greater than or equal* to 0.5. The inferences that are valid according to this definition (provided that we do not use the Łukasiewicz conditional; we will come to this in a second) are, interestingly, exactly the classically valid inferences. This may placate the conservative. However, it is hard to see why conservatism should be a virtue here. As Smith shrewdly observes (p. 3), classical logic was developed in the context of mathematical reasoning. Vague language was entirely ignored. There is absolutely no *a priori* reason why, then, validity should be the same if we are dealing with a wider class of situations including vague ones. And I think there are reasons to prefer a different definition of validity. As Smith notes, in effect, in 5.4, how true something has to be to be assertable is not absolute, but depends on the context. 'This is a new motor bike' must have a higher degree of truth to be assertable by a bike salesman than at a bike rally. Now we standardly look to validity to preserve assertability. Smith's definition does not necessarily. (Just take the assertable things to be those with value greater than 0.5.) I therefore think it better to define an inference to be valid if it never decreases the degree of truth; that is, it works in all contexts (see INCL, 11.4).

Smith takes the conditional to be, not the Łukasiewicz $\rightarrow$, but the material $\supset$ ($\neg A \vee B$). For, he notes, the sorites argument where the conditional premise is expressed by $\supset$ seems just as urgent, and a uniform solution is to be expected. This insistence is, in fact, somewhat undercut by what he says about identity. There are sorites arguments where the major premises are identities, and the rule of inference employed is the substitutivity of identicals (see INCL, 25.5). If a uniform solution is to be provided for all forms of the sorites, identity needs to be a vague predicate. But, Smith asserts in a footnote (p. 271), identity is to be understood as a crisp predicate.[3] I think it is better just to admit that there are different versions of the paradox, the machinery of a solution saying different things about the different versions.

On Smith's account, the major premise of a sorites may fall below 0.5. (Let the value of Ba_i be 0.6, and that of Ba_{i+1} be 0.4; then the value of $Ba_i \supset Ba_{i+1}$ is 0.4.) He has, therefore, to explain why we so naturally take all the conditionals to be true. He says (p. 270f.) that we do so because we take them to be an expression of the fact that Ba_i and Ba_{i+1} have identical truth values. I find this implausible, if only because conditionals are not symmetric, whilst identity is. (A person using a sentence may not mean exactly what the sentence means, but there had better be some kind of explanation of how an utterer manages to get a hearer to understand it in the way intended.) Moreover, identifying the English (indicative) conditional with the material conditional is fraught with well-known objections (see INCL, 1.9). Better to let the Łukasiewicz conditional stand (where the sorites conditionals are all very close to having the value 1). If one does this, one has to deal with the problems of penumbral connection which the conditional poses (and which Smith does not address explicitly, since they are subsumed, for him, by the ones concerning the extensional connectives). As with Hyde, I think that Smith would be better off taking the conditional to be a modalised Łukasiewicz conditional, which avoids these problems. Indeed, it is straightforward enough to build a modalised (in fact, relevant) account of the conditional on top of a continuum of truth values (see INCL, 11.7).

The most novel and interesting part of this chapter, it seems to me, is section 5.3, where Smith discusses the connection between degrees of truth and degrees of belief, and how these notions fit into the Lewis/Stalnaker account of 'keeping score' in language games. This is an important but neglected issue, and Smith is much to be commended for seeing its importance. We are to suppose that at the current stage of matters, the set of worlds in play, as possible candidates for the actual, is W. Any sentence, S, has a degree of truth $|S|_w$, at each world, w. We also assume that there is a probability function, μ, such that for all $w \in W$, $\mu(w)$ is the (subjective) probability that that is the actual world.[4] Smith's novel idea is to take the appropriate degree of belief for S to be its 'expectation'. In the finite case, this is $\Sigma_{w \in W}\ |S|_w.\mu(w)$. (The infinite case is discussed, but not defined explicitly. What is intended, I take it, is that the expectation, E, is (the Lebesgue) $\int_W\ |S|_w.dP$, where P is the probability distribution of $|S|_w$. (That is, $E = \int_W\ |S|_w.f(w)dw$, where f is the density function corresponding to P, if it has one.)

[3] Which also raises the question of what, on his view, one should say about the Evans' argument.

[4] Smith actually works with a probability function defined on sets of worlds. Given countable additivity, the two approaches are interdefinable in the finite case.

And so we come at last to Chapter 6 and higher-order vagueness. Smith starts by noting that the fuzzy solution thus far advocated is 'too precise'. It is arbitrary to assign a sentence, S, a truth value of 0.996 instead of 0.997. His solution is to take the intended interpretation of our language to be under-determined by the linguistic practice of a set of speakers. What the practice determines is a whole class of acceptable interpretations. Some will assign S a value of 0.976; some will assign it a value of 0.997, etc. But obviously the problem recurs. How does practice determine a unique such set? Thus, let V be the set of admissible evaluations. Consider the evaluation, f, of the monadic predicate, P, such that $f(P)(x) = Sup\{\upsilon(P)(x) : \upsilon \in V\} + \varepsilon(x)$, where $\varepsilon(x)$ is some appropriate function, whose values are in the order of $10^{-10^{10}}$. This is not in V; yet it is hard to see how practice could be so discriminating as to keep it out.

The problem of over-precision also arises in another predictable location. In a sorites progression (say from blue to not blue), any intended interpretation will deliver a last object of which it is true to say that it is blue to degree 1. This, as we have already observed, seems wrong. Locating any such point would seem to be arbitrary. One can defuse this problem by appealing to the fact that the point will vary from acceptable interpretation to acceptable interpretation. But as ever, the problem has not left us. As Smith notes (p. 305) there will still be a last point which is [blue on every interpretation]. This seems just as bad.

Smith's response to these problems is to insist that the requisite precision is, indeed, there. The (precise) set of admissible interpretations is determined by a variety of considerations including, especially, what the group of speakers of the language say, or are disposed to say, about the various states of affairs at issue. How, exactly, the determination works is not spelled out in detail (pp. 286–287). However, I find it hard to see how such considerations will determine a unique and precise set – for this reason if no other: the speakers concerned must, presumably, be competent speakers. And competence comes by degrees. (When does a child become a competent speaker?) Given that the set of speakers (and so of actual and dispositional behaviour in question) is itself vague, a precise determination seems implausible.

An obvious thought at this point is to take the set of admissible interpretations to be vague, by taking Smith's own account to be given in a vague metalanguage. Unlike Hyde, however, he rejects this possibility (6.2.1). As we have already seen, such projects are indeed problematic. But for all this, they may be possible; and I certainly was not persuaded by the reason Smith gives for writing them off. His reason is that we must have a well-grounded prior understanding of the metatheoretic machinery if our account is to be of value. This, I think, is wrong. Metatheoretic projects always involve a certain amount of boot-strapping. This is just as true if we are giving a classical metatheory for a classical language, an intuitionist metatheory for an intuitionist language, or a fuzzy metatheory for a fuzzy language. Machinery and understanding develop in tandem. This may be very difficult, but not impossible.

To conclude: for me, both Hyde's and Smith's solutions to the problem of vagueness fail to clear the hurdle of higher-order vagueness. Optimists may say that we have, at least, solutions to the problem of vagueness; we are just left with some details to mop up. Pessimists (of whom I am one in this regard), will say that higher-order vagueness is not some problem over and above vagueness; it is just the problem of vagueness as it arises with respect to a putative solution – just as the so-called 'extended paradoxes' of self-reference are not really paradoxes additional to the standard ones, but are just the way that the phenomenon in question forces itself upon us, given a certain theoretical setting. Be that as it may, both Hyde's and Smith's books are to be much welcomed and praised. Even though many books and papers have been written on the topic in the last 40 years, these two books have new insights to offer, and advance our understanding of the terrain in notable ways. Both are a must-read for anyone interested in vagueness.[5]

[5] Many thanks to Hyde and Smith for helpful comments on earlier drafts of this review, which greatly improved it.