Fuzzy Identity and Local Validity

1. Two Sorts of Sorites

Standard sorites paradoxes can always be put into a simple canonical form, employing the sole inference *modus ponens*. For example, consider the following paradox. Take a continuum of colours going from red to blue, and let $a_1, \ldots, a_m$ be a sequence of segments of this continuum (starting at the red end) such that each segment is phenomenologically indistinguishable in colour from its immediate neighbours. Let Fx be the predicate 'x is red'. Then the untrue conclusion Fa_m can be inferred from the premises Fa_0 and $Fa_n \rightarrow Fa_{n+1}$ (with $0 \leq n < m$).

There is, however, another, and less familiar kind of kind of sorites that uses, not *modus ponens*, but the substitutivity of identicals, more specifically, the transitivity of identity. For example, given the same sequence as before, let r refer to the hue of red of a_0, and consider the function symbol, f, whose intuitive interpretation is 'the (phenomenological) colour of'. Then the conclusion $fa_m = r$ can be inferred from the premises $fa_0 = r$ and $fa_n = fa_{n+1}$ (with $0 \leq n < m$); premises that have exactly the same plausibility as the original sorites.

The parallels between the two sorts of paradox can be made closer still. For note that the converses of the conditionals which are the premises of the first kind of sorites are not at all contentious (if a_{n+1} is red, certainly a_n is red). Hence, we could easily take as premises the *bi*conditionals (for which, of course, *modus ponens* is equally valid). But, intuitively, biconditionality is to sentences what identity is to terms, namely an equivalence relation satisfying substitutivity in extensional contexts).[1]

2. One Kind of Solution

There are, of course, many suggestions as to how the standard sorites paradoxes should be solved. All face well known objections. In particular, all standard accounts face problems concerning "higher order" vagueness:

"Fuzzy Identity and Local Validity" by Graham Priest,
The Monist, vol. 81, no. 2, pp. 331–342.

they appear to impute more precision to the semantic situation than is actually there.[2] It is not my intention to address any of this here. Solutions can be divided into exactly two kinds: those that deny the truth of at least one of the conditional premises, and those that deny the validity of the inference employed.[3] In this paper, I will just focus upon one aspect of solutions of the second kind.

Such solutions must find a way of showing how an application of *modus ponens* can fail. For example, a natural suggestion is to suppose that truth comes by degrees, which can be measured, conventionally, by numbers in the closed interval [0, 1].[4] For reasons that will become clearer later, I will reverse the usual conventions, and let 0 mark the true end of the continuum. Sentences receive a truth value in this interval, and truth conditions can be given for connectives in a natural way, as, for example, in Łukasiewicz's continuum-valued logic. In this, the value of $\alpha \rightarrow \beta$ is the value of $\beta \dot{-}$ the value of α. ($\dot{-}$, here, is truncated subtraction, i.e., $x \dot{-} y = x - y$ if $x \geq y$, and 0 otherwise.) Now suppose that α and β have the values 0.25 and 0.5, respectively. Then the value of $\alpha \rightarrow \beta$ is also 0.25. Hence, the conclusion of a *modus ponens* can have a higher value (i.e., be more false) than its premises.

Solutions of the kind that invalidate *modus ponens* typically monkey around with the truth conditions for connectives and qualifiers. However, they tend to leave identity alone. For example, in the continuum-valued approach, it is normal to take identity to be a non-fuzzy predicate such that the value of '$a = b$' is 0 if the denotations of 'a' and 'b' are identical, and 1 otherwise. If you are tempted by this kind of solution (as I am), then, in the light of the second kind of sorites paradox, this is likely to seem inadequate: it is natural to suppose that identity, too, must be "fuzzified." The rest of the paper takes a look at the consequences of this move.

3. *Fuzzy Identity*

Let us suppose that we are given a domain of objects, D, and that the objects come with a distance metric, d. Specifically, d is a non-negative real-valued function satisfying the conditions:

$$d(x, x) = 0$$
$$d(x, y) = d(y, x)$$
$$d(x, z) \leq d(x, y) + d(y, z)$$

where $x, y, z \in D$.[5] If the distance between x and y is small under the metric, then they may not be *completely* identical, but they will be *almost* so. That is, the degree of their identity will be nearly 0. This suggests taking the metric itself as providing the degree of truth of an identity statement (at least, when x and y are close).

For example, suppose that we take D to be the real numbers, and let $d(x, y)$ be the usual metric, $|x - y|$. Let 'a' and 'b' denote 1 and 1.01, respectively. Then the truth value of $a = b$ is .01, which is very true, though not completely so. Alternatively, let D be the real plane and let $d(x, y)$ be the area of the symmetric difference between x and y, $x\backslash y$. Let 'a' denote the unit circle, and let 'b' denote the unit circle together with a small disjoint circle of radius 0.1 and centered at (1, 1). Then the truth value of $a = b$ is $\pi/100$, which is, again, pretty nearly completely true.

It is an easy enough task to build this idea into a formal semantics for a language with fuzzy identity. The language is a standard first-order language. An interpretation is a triple $\langle D, d, I\rangle$, where D is a non-empty domain, d is a metric on D, and I assigns denotations to all the non-(logical constants). I assigns every constant a member of D, every n-place function symbol a function from D^n to D, and every n-place predicate a function from D^n to [0, 1]. Identity is a special predicate, however; in particular, $I(=)(x, y) = \min(1, d(x, y))$.[6]

Given an interpretation for a language, $\mathcal{A}$, a denotation $I(t)$ is assigned to every term, t, by the usual recursive clause: $I(ft_1 \ldots t_n) = I(f)(I(t_1), \ldots, I(t_n))$. A truth value, $v_{\mathcal{A}}(\alpha)$, in the interval [0, 1], is then assigned to every formula, α, by the standard Łukasiewicz conditions.[7] For simplicity, we assume that every member of D has a name:

$$v_{\mathcal{A}}(Pt_1 \ldots t_n) = I(P)(I(t_1), \ldots, I(t_n))$$
$$v_{\mathcal{A}}(\neg\alpha) = 1 - v_{\mathcal{A}}(\alpha)$$
$$v_{\mathcal{A}}(\alpha \wedge \beta) = \max(v_{\mathcal{A}}(\alpha), v_{\mathcal{A}}(\beta))$$
$$v_{\mathcal{A}}(\alpha \to \beta) = v_{\mathcal{A}}(\beta) \dot{-} v_{\mathcal{A}}(\alpha)$$
$$v_{\mathcal{A}}(\forall x\alpha) = \text{lub}\,\{v_{\mathcal{A}}(\alpha(x/c)); \text{for every constant, } c\}$$

$\wedge$, $\exists$ and $\leftrightarrow$ can be given the natural truth conditions, or defined in the usual way. In what follows, I will often omit the subscript on v when this should cause no confusion.

In a many-valued logic, validity is standardly defined in terms of the preservation of designated values, these being the values of the "acceptable" sentences. In a fuzzy logic, an acceptable sentence is one that is "true enough". Now, what counts as true enough may well depend on the context; but anything at least as true as something true enough is true enough. Hence, designated values should at least be closed downwards under <. It is therefore natural to take an inference to be valid in a context with level of acceptability ε ($0 < \varepsilon < 1$) just if in any interpretation where the values of the premises are less than ε, so is that of the conclusion. In logic, it makes sense to abstract out from the vicissitudes of context, and define validity in terms of what works in all contexts. Hence we define validity as follows. For simplicity, I will restrict myself to the single-premise case. (This covers the finite-premise case, since, given the truth conditions for conjunction, an inference with finitely many premises can be reduced to one with a single conjoined premise. The infinite premise case is a simple generalisation.) $\pi \vDash \kappa$ iff $\forall \varepsilon \forall \mathcal{A}(\nu_{\mathcal{A}}(\pi) < \varepsilon \Rightarrow \nu_{\mathcal{A}}(\kappa) < \varepsilon)$. As is easy to check, this is equivalent to the following: $\forall \mathcal{A}(\nu_{\mathcal{A}}(\kappa) \leq \nu_{\mathcal{A}}(\pi))$. Valid arguments, then, are ones whose conclusions are always at least as true as their premises.

4. Local Validity

An inference not containing identity is valid iff it is valid in the intersection of all the Łukasiewicz systems (with variable ε).[8] But our concern here is with identity, so let us pass on to this. As is easy to check, identity is reflexive and symmetric. It is not, however, transitive. To see this, merely consider the model where $D = \{0, 1, 2\}$, and $d(x, y)$ is $|x - y| .\varepsilon/2$. Let the constants a, b, and c denote 0, 1 and 2 respectively. Then $\nu(a = b) = \nu(b = c) = \varepsilon/2$, whilst $\nu(a = c) = \varepsilon$.

Although the law of transitivity fails, there is a close approximation. If a is nearly (completely) identical to b, and b is nearly (completely) identical to c, then a is nearly (completely) identical to c. We can be sure, then, that the conclusion of an instance of transitivity is acceptable, provided that the premises are "true enough," that is, sufficiently close to 0. If we use 'ε' and 'δ' for positive reals, then, in the language of classical analysis: for every ε, there is a δ such that $\nu(a = b \wedge b = c) < \delta \Rightarrow \nu(a = c) < \varepsilon$. This follows from the properties of metrics. Since $d(x, z) \leq d(x, y)$

+ $d(y, z)$, we may simply take the required δ to be $\varepsilon/2$. More generally, consider an inference with premise π and conclusion κ. Let us say that it is *locally valid* (*LV*) iff $\forall\varepsilon\exists\delta\forall\mathcal{A}(v_{\mathcal{A}}(\pi) < \delta \Rightarrow v_{\mathcal{A}}(\kappa) < \varepsilon)$.[9] The notion of local validity captures the idea that an inference is "reasonably safe": if our premises are "true enough," the conclusion is going to be acceptable.[10]

It is clear that any inference that is valid is LV; but, as we have seen, the converse is not true: transitivity of identity is LV but not valid. Another example of this kind is *modus ponens*. We have already noted that this is invalid. To see that it is LV, suppose that for any v the value of $\alpha \wedge (\alpha \rightarrow \beta)$ (and so of each conjunct separately), is less than δ. If $v(\beta) \leq v(\alpha)$, $v(\beta) < \delta$. If $v(\beta) > v(\alpha)$ then $v(\beta) - v(\alpha) < \delta$. Hence, $v(\beta) < 2\delta$. In either case, $v(\beta) < 2\delta$. Thus, provided we choose δ to be less than $\varepsilon/2$, $v(\beta)$ will be less than ε.[11] The treatment of the two kinds of sorites along these lines is therefore essentially the same, as one would hope.

If an inference is locally valid, we know that its conclusion will be acceptable, provided that our premises have value close enough to 0. It should be noted that how close 'close enough' is, is not something that the notion of local validity itself tells us. That an inference is locally valid just assures us that there is *some* level of truth of the premises that will make the conclusion acceptable. What that level is, will depend on features of the particular inference in question, such as the "drop off rate" from premise to conclusion, and non-formal features of the situation, such as the level of acceptability required. That an inference is locally valid does not, therefore, on its own, answer the question of whether it is applicable on a particular occasion. Rather, the point of the notion of local validity is to circumscribe a class of inferences with the important property that there is an *a priori* guarantee that they are legitimately applicable under certain well-defined conditions.[12]

It may not yet be clear why I have called the notion 'local validity'. To see this, consider what happens when we have a sorites argument. For example, let Σ be the set of its major premises (conditionals or identities), and let $\vdash_\Sigma$ be deducibility with respect to Σ. The sorites is a sequence of inferences: $\alpha_0 \vdash_\Sigma \alpha_1 \vdash_\Sigma \ldots \quad \alpha_{v-1} \vdash_\Sigma \alpha_n$, each of which is LV. It is not difficult to see that the resulting inference, $\alpha_0 \vdash_\Sigma \alpha_i$, is also LV. For any $\varepsilon > 0$, we can choose δ_{n-1} small enough so that if the truth values of α_{n-1} and all the members of Σ are less than δ_{n-1} the truth value of α_n is less than

ε. But then we can choose a δ_{n-2} such that if the truth values of α_{n-2} and all the members of Σ are less than δ_{n-2}, that of α_{n-1} is less than δ_{n-1}, and so on . . . we can choose a δ_0 such that if the truth values of α_0 and all the members of Σ are less than δ_0, that of α_1 is less than δ_1. Hence, if δ is the minimum of all the δ_is, and the values of α_0 and all the members of Σ are less than δ, the value of α_n is less than ε, as required.[13] Hence, the soritical inferences can be chained together with security, as long as the premises are "true enough," and we do not chain too many together. Chain too many, however, and things may go wrong. This seems to capture the sorites phenomenon *exactly*: sorites inferences are acceptable provided we use them locally (to take us a short distance down the sorites), but not globally.

Any solution to the sorites paradox must do more than show where the unacceptable argument breaks down. It must, just as importantly, diagnose the confusion that led us to suppose that it was sound in the first place—otherwise the solution would leave a problem as mysterious as the one it is meant to solve. The machinery at hand allows us to do this. Our mistake is simply thinking that inferences that can be used over short distances are reliable over long distances. That is, we confuse (global) validity with local validity.

5. *Evans's Argument*

Let us next look at the behaviour of the substitutivity of identicals in these semantics. Transitivity of identity is an instance of substitutivity. Substitutivity must therefore fail. For another example, suppose that our language is augmented with an operator, Δ, whose intuitive sense is 'it is definitely true that'. It is natural to take $v(\Delta\alpha)$ to be 0 if $v(\alpha) = 0$, and 1 otherwise. Now consider the inference: $a = b$, $\Delta a = a \vdash \Delta a = b$. Choose an interpretation where $D = \{0, 1\}$, $d(0, 1) = \eta < \varepsilon$, $I(a) = 0$ and $I(b) = 1$. Then the premises have value η and 0, respectively, whilst the conclusion has value 1. Moreover, for any δ, we can choose $\eta < \delta$. Hence the inference is not even locally valid.

I choose this example because it is of some importance philosophically. There is a well-known argument due to Evans, aimed at showing that there can be no vague objects: all true identities are definitely true. Evans's argument is essentially the inference we have just seen to be invalid (given that $\Delta a = a$). Hence, it fails.[14]

It is natural to defend Evans's argument as follows. According to the semantics given, the central inference of the argument is invalid, but this is only because the notion of identity involved in it is not true identity, but an *ersatz*. The correct notion of identity is one for which (*):

$$v(a = b) = 0 \text{ if } I(a) = I(b)$$
$$v(a = b) = 1 \text{ otherwise}$$

For such a notion of identity, substitutivity *is* valid, and so the argument works.

Such a defence would fail though: it is entirely question-begging. For what is at issue is *exactly* the correct truth conditions of identity statements for vague objects. And just because the truth conditions (*) are more familiar, it does not follow that they are correct. Plausibly, the fuzzy truth conditions do not define an *ersatz* identity relation, but the genuine thing for vague objects. Suppose, for example, that I take a picture of *a* against a beautiful background. *a* then moves. I want to take a picture of *b*, and I tell them to go and stand where *a* was standing, which they do. Even though *b*'s feet may be in a slightly different position from *a*'s, *b* is in the same place as *a*. That is, the place where *a* had been (a vague spot) is identical with the place where *b* is.

The objector does not even have a right to *assume* that the domain *D* is furnished with a relation of the kind required for classical identity. (Note that the semantics given make no use of identity as applied to objects in *D*. The metalanguage contains a crisp identity predicate, but it is only ever applied to truth values.) From the perspective of fuzzy identity, the conditions (*) are not even well defined. For a start, the definition assumes the Law of Excluded Middle, which is moot. Moreover, if the relation = between members of the domain is a fuzzy one, the conditions make the relation between formulas and their truth-values fuzzy too, which it is not.

6. *Tolerance*

As we see, then, substitutivity of identity, real identity, fails. But there is more to the story. It is a feature of vague predicates (though not crisp ones) that they are *tolerant* with respect to small changes in their arguments. That is, small changes in their arguments make relatively no change in the extent to which the predicate applies. Thus, if P is a predicate of this kind —let us keep matters simple and consider just monadic predicates—and f is its interpretation, then small changes to x make relatively little difference to $f(x)$. Doubtless, if the domain, D is the reals, this means that f should be continuous: $\forall \varepsilon \exists \delta (d(x, y) < \delta \Rightarrow |f(x) - f(y)| < \varepsilon)$. But this condition is not sufficient. For a start, if D is not the reals, and d is a discrete metric with a minimum difference between each pair of objects (and this can certainly happen with vague predicates: think of 'is small' on the natural numbers), then this condition is automatically satisfied, no matter what f is. (Take δ to be less than the minimum distance.) Even if f is continuous, this is compatible with it changing arbitrarily large amounts for arbitrarily small changes in its arguments. (Think of the exponential function.) A natural way of cashing out the idea that f represents a tolerant predicate, which avoids these problems, is the condition: $|f(x) - f(y)| \leq d(x, y)$ (which certainly implies continuity in the case of reals).[15]

We may express the tolerance of a predicate, in this sense, in the object language, if we add a new functor, *Tol*, such that if P is any predicate $Tol(P)$ is a formula, and give it the following truth conditions:

$$v(Tol(P)) = 0 \text{ if } \forall x, y \in D, |v(P)(x) - v(P)(y)| \leq d(x, y)$$
$$v(Tol\ (P)) = 1 \text{ otherwise}$$

If we do this, then the following form of substitutivity *is* LV:

$$Tol(P) \wedge a = b \wedge Pa \vdash Pb$$

To show this, we need to establish that $\forall \varepsilon \exists \delta \forall \mathcal{A} (v_{\mathcal{A}}(Tol(P) \wedge a = b \wedge Pa) < \delta \Rightarrow v_{\mathcal{A}}(Pb) < \varepsilon)$. Fix ε. If $v_{\mathcal{A}}(Tol(P)) = 1$, then δ can be any positive real < 1. So suppose that $v_{\mathcal{A}}(Tol(P)) = 0$, and write $I(P)$ as f, $I(a)$ as x and $I(b)$ as y. Then $|f(x) - f(y)| \leq d(x, y)$. So $f(y) \leq f(x) + d(x, y)$. Hence, $f(x)$, $d(x, y) < \varepsilon/2 \Rightarrow f(y) < \varepsilon$, i.e., $v_{\mathcal{A}}(a = b \wedge Pa) < \varepsilon/2 \Rightarrow v_{\mathcal{A}}(Pb) < \varepsilon$. Thus we may take δ to be $\varepsilon/2$.[16]

7. The Fuzzy Continuum

As one application of the theory of vague identity, let us consider briefly the fuzzy continuum. Let D be the real numbers, and let the metric

$d(x, y)$ be the standard distance function, $|x - y|$. Take a language that has a supply of constants (for simplicity, we will take these to be the members of D themselves), monadic function symbols and the binary predicates $=$ and $<$. Now consider the interpretation $\langle D, d, I\rangle$, where, for every constant, c, $I(c) = c$; for any function symbol, f, $I(f)$ is a (simply) continuous function; and $I(<)$ is the standard ordering relation on D.[17] The fuzzy continuum of degree of coarseness ε—or rather, its description—is the set of sentences whose values in this interpretation are less than ε.

The fuzzy continuum has some interesting properties. For example, suppose that $f(0) = 1$ is true (i.e., its truth value is less than ε) then if r is any real number such that $1 = r$ is "true enough," $f(0) = r$ is true. For example, suppose that $\varepsilon = 0.3$, and that the value of $f(0) = 1$ is 0.1. Consider the identity $f(0) = 1.1$. The value of this must be less than or equal to the sum of those of $f(0) = 1$ and $1 = 1.1$, i.e., 0.2, and so this equation is true.

Does this mean that f is not single valued? Not at all! Let t be any term. Consider the sentence $\forall y(y = t \rightarrow y = t)$. The value of this is 0. Hence the value of $t = t \wedge \forall y(y = t \rightarrow y = t)$ is 0, as, then, is that of $\exists x(x = t \wedge \forall y(y = t \rightarrow y = x))$, i.e., $\exists !x(x = t)$. In particular, take t to be $f(0)$; then we have $\exists !x(x = f(0))$. f *is* single valued; it's just that its value is a vague one: about 1.

Next, suppose again that $f(0) = 1$ is true. If r is any real number such that $0 = r$ is true enough, $f(r) = 1$ is also true. Since f is continuous, for any ε, there is a δ such that $|0 - r| < \delta \Rightarrow |f(0) - f(r)| < \varepsilon$. Now suppose that the value of $f(0) = 1$ is $\eta(< \varepsilon)$, and choose δ such that $|f(0) - f(r)| < \varepsilon - \eta$. Then the value of $f(r) = 1$ is less than or equal to the sums of those of $f(r) = f(0)$ and $f(0) = 1$, which is less than ε, as required.

More generally, suppose that the value of $f(a) = b$ is η. Choose $r > a$ such that $|r - a| < \varepsilon$ and $|f(r) - f(a)| < \varepsilon$. Then the value of $f(r) = b$ is less than $\eta + \varepsilon$. So, therefore, is that of $\exists x(x > a \wedge f(x) = b)$. Hence, the truth value of $f(a) = b \rightarrow \exists x(x > a \wedge f(x) = b)$ is less than ε. (Exactly the same argument works if we replace '$<$' with '$>$'.) Hence, whatever the argument of a function, there are greater and lesser numbers such that it is true to say that f applied to these takes the same value.

It seems to me that those properties of the fuzzy continuum may well have significant implications for the issue of higher-order vagueness, and so for the question of a general and adequate solution for the problems of vagueness. However, that is a matter for another paper.

8. Conclusion

We have seen that a semantics can be given for fuzzy identity which is natural, and which also brings out the similarity between the two different kinds of sorites in a natural way. It makes possible the notion of *local validity*, which provides for a simple and plausible diagnosis of the mistake in sorites arguments. We have also seen that the semantics shows the invalidity of the standard objection to fuzzy identity. As a theory of fuzzy identity, little more could be asked of it.[18]

Graham Priest

The University of Queensland, Australia

NOTES

1. The two kinds of sorities are discussed at greater length in Priest (1991).
2. Nor does appealing to dialetheism help at all. Suppose, for example, that one takes a to be a borderline case of P iff Pa is both true and false. Then in a sorites-style transition, there will still be a precise borderline between cases where Pa is true and not false, and cases where it is both true and false. Nor can one simply suppose that sorites arguments are veridical; for they give rise, not just to isolated contradictions, but to near-triviality. Let a be any object, and let P be any property such that for some b, Pb. Let b change, gradually and continuously, into a. A sorites argument will now establish that Pa. In particular, taking P to be the predicate $x = b$, it would follow that all things are one.
3. Strictly speaking, there is a third possibility: deny the truth of the minor premise. One cannot say that such solutions are impossible: in this area, everything is contentious. But this would be an act of extreme desperation.
4. See, e.g., Goguen (1968–69), Machina (1976).
5. In the mathematical theory of metric spaces, metrics are usually taken to satisfy the extra condition $d(x, y) = 0 \Rightarrow x = y$, ruling out the degenerate metric (see, e.g., Dieudonné (1969), p. 28). Whilst this condition could be added here also, it plays no significant part in what follows.
6. The minimisation is necessary since truth values must be no greater than 1. One way to avoid this would be to employ, instead, a normalised metric, whose value is always between 0 and 1.
7. The Lukasiewicz truth conditions are not the only possible ones for a fuzzy logic. However, they are ones that have very satisfactory consequences in the present context, as we shall see.
8. For a survey of some results concerning these, see Chang (1963).
9. In the language of non-standard analysis, this is equivalent to saying that whenever the value of π is infinitesimally close to zero ($v_{\mathcal{A}}(\pi) \simeq 0$), so is that of κ. The proof is as follows. (Thanks are due to Moshe Machover here.)

Suppose that the inference is locally valid. Then for all real ε there is a real δ such that $\forall \mathcal{A}(v_{\mathcal{A}}(\pi) < \delta \Rightarrow v_{\mathcal{A}}(\kappa) < \varepsilon)$ is true in the standard model. Hence it is true in the non-standard model. But now suppose that in that model $v_{\mathcal{A}}(\pi) \simeq 0$. Then $v_{\mathcal{A}}(\pi) < \delta$. Hence $v_{\mathcal{A}}(\kappa) < \varepsilon$. And since this is true for all real ε, $v_{\mathcal{A}}(\kappa) \simeq 0$, as required.

Conversely, suppose that the inference is not locally valid. Then in the standard model, for some real ε, $\forall \delta \exists \mathcal{A}(v_{\mathcal{A}}(\pi) < \delta \wedge v_{\mathcal{A}}(\kappa) \geq \varepsilon)$. Hence we can find a sequence, $\langle \mathcal{A}_i; i \in N \rangle$, such that the corresponding sequence $\langle v_i(\pi); i \in N \rangle$ converges to 0, whilst $\forall i\, v_i(\kappa) \geq \varepsilon$ is true in the standard model, and so in the non-standard model. Now let n be some non-standard integer. Then in the non-standard model $v_n(\pi) \simeq 0$, by the properties of convergence, whilst $v_n(\kappa) \geq \varepsilon$. Hence, in the non-standard model $\exists \mathcal{A}(v_{\mathcal{A}}(\pi) \simeq 0 \wedge \neg v_{\mathcal{A}}(\kappa) \simeq 0)$, as required.

10. Interestingly, Adams's probabilistic notion of *reasonable validity* is defined in a very similar way (1966), p. 274.

11. Using infinitesmals there is an even easier argument: suppose that $v(\alpha) \simeq 0$, and $v(\alpha \rightarrow \beta) \simeq 0$. Either $v(\beta) \leq v(\alpha)$, and so $v(\beta) \simeq 0$, or $v(\beta) \geq v(\alpha)$, in which case $v(\beta) \simeq v(\alpha)$, and hence again $v(\beta) \simeq 0$.

12. It is an interesting project to give a proof theory for the notion of local validity, but one for another occasion.

13. More generally, it is an easy exercise to show that any chain of LV inferences is LV.

14. Strictly speaking, Evans's argument is slightly different. He interprets Δ as 'it is definite that', gives the contraposed form of the argument, and also uses property abstraction. But these differences do not affect the point being made here. For a longer discussion of Evans's argument along essentially the same lines, see Copeland (1994).

15. To a certain extent, this is arbitrary. The condition $|f(x) - f(y)| \leq k.d(x, y)$, where k is some positive integer, would do just as well. However, this is no real generalisation, since its effect can be obtained simply by renormalising d.

16. Note that requiring P to have a continuous interpretation is not sufficient to give LV. To see this, fix ε. There is no δ such that $\forall \mathcal{A}(v_{\mathcal{A}}(a = b \wedge Pa) < \delta \Rightarrow v_{\mathcal{A}}(Pb) < \varepsilon)$. For choose an interpretation, $\mathcal{A}$, where D is the non-negative reals, the metric is as usual, $v_{\mathcal{A}}(a) = 0$, $v_{\mathcal{A}}(b) = \delta/2$, and $v_{\mathcal{A}}(P)(x) = 2x\varepsilon/\delta$ (which is certainly continuous). Then $v_{\mathcal{A}}(a = b \wedge Pa) = \delta/2$, but $v_{\mathcal{A}}(Pb) = \varepsilon$. The point is that the definition of LV requires δ to be uniform in $\mathcal{A}$. Simple continuity does not guarantee this: given δ, we can always make f rise more sharply.

17. It would be in the spirit of vagueness to fuzzify the relation, so that the interpretation of < is a tolerant approximation to the usual step-function. However, let us keep matters simple for the present.

18. Earlier drafts of this paper were read at the Universities of Notre Dame, Torùn, St. Andrews and Uppsala. I am very grateful to many people at these places for valuable comments and suggestions. I am also grateful to Tim Williamson and an anonymous referee for comments that further improved the paper.

REFERENCES

[1] E. Adams (1966), "Probability and the Logic of Conditionals," ch. 14 of J. Hintikka and P. Suppes (eds.), *Aspects of Inductive Logic*, North-Holland Publishing Company.

[2] C. C. Chang (1963), "Logic with Positive and Negative Truth Values," *Acta Philosophica Fennica* 16, 19–38.
[3] B. J. Copeland (1994), "On Vague Objects, Fuzzy Logic and Fractal Boundaries," *Southern Journal of Philosophy*, Supp. vol. 33, 83–96.
[4] J. A. Dieudonné (1969), *Foundations of Modern Analysis*, Academic Press.
[5] G. Evans (1978), "Can there be Vague Objects?," *Analysis* 38, 206.
[6] J. Goguen (1968–69), "The Logic of Inexact Concepts," *Synthèse* 19, 325–73.
[7] K. Machina (1976), "Vagueness, Truth and Logic," *Journal of Philosophical Logic* 5, 47–78.
[8] G. Priest (1991), "Sorites and Identity," *Logique et Analyse*, 135–36, 293–96.