THE GOOD, THE BAD, and THE UGLY

Francis Jeffry Pelletier

University of Alberta

Three Types of Vagueness:

Many different kinds of items have been called vague, and so-called for a variety of different reasons. Traditional wisdom distinguishes three views of why one might apply the epitaph "vague" to an item; these views are distinguished by what they claim the vagueness is due to.

One type of vagueness, The Good, locates vagueness in language, or in some representational system -- for example, it might say that certain predicates have a *range of applicability*. On one side of the range are those cases to which the predicate clearly applies and on the other side of the range are those cases where the negation of the predicate clearly applies. But there is no sharp cutoff place along the range where the one range turns into the other. Most examples of The Good are those terms which describe some continuum -- such as *bald* describes a continuum of the ratio of hairs per cm^2 on the head. But not all work this way. Alston (1968) points to terms like *religion* invoking a number of criteria the joint applicability of which ensures that the activity in question is a religion and the failure of all to apply ensures that it is not a religion. But when only some middling number of the criteria are fulfilled, the term *religion* neither applies nor fails to apply. Some accounts of "family resemblance" and "open texture" might also fit this picture. Such a view is often called a "representational account of vagueness".

Another conception of vagueness, The Bad, locates vagueness as a property of discourses, of memories, and of certain philosophers and their papers, etc. This sort of vagueness occurs when the information available does not allow one to tell, for example, that a certain sentence is true, but also does not allow one to determine that it is false. It occurs when the information available does not allow one to claim that a predicate applies to a name, but also does not allow one to claim that the negation of the predicate applies to that name. The underlying intuition here is that the sentence either *is* true or *is* false, that either the predicate *does* apply to the name or its negation does, but the vagueness results in it being impossible to determine which. Such a view is often called an "epistemological account of vagueness". Wheeler (1975: 369) characterizes these as

> cases where there is an answer as to whether a predicate is true of a [situation in which] the speaker 'doesn't know what to say', but this answer depends on data such that, if it were in, the speaker *would* know what to say....Epistemological vagueness...is a matter of not having determined whether the object conforms to the features of the concept sufficiently to fall under it or not, not having 'total information' about the case.

Besides The Good and The Bad, there is The Ugly. This conception of vagueness locates it "in the world". So, as opposed to being vaguely described, and as opposed to being unclear as to whether a situation actually obtains or not, The Ugly claims that certain *objects* in the world just plain *are* vague. Few writers have advocated -- or even explained, except with contempt -- such a view. Russell (1923) however assures us that The Ugly used to be a common view: "a case of the fallacy of verbalism -- the fallacy that consists in mistaking the properties of words for the properties of things." And more recently Heintz (1981) points to fictional entities, such as Hamlet, as neither having nor lacking certain properties -- such as having a 5mm wart on the great toe of the left foot. What makes fictional objects vague, he says, is precisely that there is a property which they neither have nor lack. Thus, a realistic theory of vagueness (as this view might be called) would claim that for vague actual objects, there is some property which it neither has nor lacks. In the course of arguing against The Ugly, Evans (1978) claims that such views must invoke the claim that, for certain names a and b the sentence 'a=b' is neither definitely true nor definitely false. In explaining Evan's explanation, van Inwagen (1986) -- the only hard-core believer in The Ugly that I know of -- proposes that the crucial test would be a situation in which the question "in talking about x and y, how many things are we talking about?" has the features that "none" is definitely a wrong answer; "three", "four", etc., are all definitely wrong answers; and neither "one" nor "two" is either definitely a right or definitely a wrong answer to it.

So: those are three views of what vagueness is. We shall investigate each of them in more detail shortly, but first some preliminaries. It is commonly believed that The Good is good, that The Bad is bad, and that The Ugly is ugly. Although I shall not defend this assessment of philosophical judgment in detail, a few representative quotations will give the flavour of (what I perceive to be) the popular philosophical view. In his influential textbook (influential because it was the only introductory text in philosophy of language for such a long time), Alston (1964: 85-86) says that The Good is good and The Bad is bad.

> [A] confusion that has infected many theoretical discussions is that between vagueness as a semantic feature of a term [i.e., The Good]...and vagueness as an undesirable feature of a certain piece of discourse.... Vagueness in the first sense is not always undesirable. There are contexts in which we are much better off using a term that is vague in a certain respect than using terms that lack this kind of vagueness. One such context is diplomacy....There would be grave disadvantages in removing this kind of vagueness...We *need* vague terms for situations like this.... There are also theoretical advantages to vagueness. Often our knowledge is such that we cannot formulate what we know in terms that are maximally precise without falsifying the statement or going far beyond the evidence.

Austin (1962: 125) too makes the latter claim

> ...we speak of people 'taking refuge' in vagueness -- the more precise you are, in general the more likely you are to be wrong, whereas you stand a good chance of *not* being wrong if you make it vague enough.

Both are following Russell (1923: 91): "...a vague belief has a much better chance of being true than a precise one, because there are more possible facts that would verify it...Precision diminishes the likelihood of truth...". Quine thinks that vagueness in natural language is inescapable (1960: 125)

> Vagueness is a natural consequence of the basic mechanism of word learning. The penumbral objects of a vague term are the objects whose similarity to ones for which the verbal response has been rewarded is relatively slight.

And this kind of vagueness, The Good, is good (1960: 127)

> Good purposes are often served by not tampering with vagueness. Vagueness is not incompatible with precision. As Richards has remarked, a painter with a limited palette can achieve more precise representations by thinning and combining his colors than a mosaic worker can achieve with his limited variety of tiles, and the skillful superimposing of vagueness has similar advantages over the fitting together of precise technical terms...Also, vagueness is an aid in coping with the linearity of discourse. An expositor finds that an understanding of some matter A is necessary preparation for an understanding of B, and yet that A cannot itself be expounded in correct detail without, conversely, noting certain exceptions and distinctions which require prior understanding of B. Vagueness, then, to the rescue! The expositor states A vaguely, proceeds to B, and afterward touches up A, without ever having to call upon his reader to learn and unlearn any outright falsehood in the preliminary statement of A.

Dummett (1975: 314-315) points to the non-transitivity of the relation 'not discriminably different' as "explaining the feeling we have that vagueness is an indispensable feature of language -- that we could not get along with a language in which all terms were definite", and as "providing us with a firm reason for saying that vague predicates are indispensable."

The Ugly, when it is mentioned at all, is characterized as ugly. The classic judgment is Russell's (1923: 85)

> Apart from representation,...there can be no such thing as vagueness or precision; things are what they are, and there is an end of it.

Or as Margalit (1976: 213) says, without mention of Russell

> Things are what they are. They are not what they are in grades, shades, or degrees. It is only relative to our ways of classifying them that they are subjected to gradations. The ascription of vagueness to objects may yield the quantities-turn-qualities kind of 'logic' ('dialectical' or otherwise), which commit the fallacy of *verbalism*, i.e., the ascription of properties of words to objects.

And finally, Dummett (1975: 314) says: "the notion that things might actually *be* vague, as well as being vaguely described, is not properly intelligible."

Representing Vagueness in Logic:

The challenge of vagueness is to give some coherent account of the underlying logic which is presupposed by the use of vague terms. Clearly, the different conceptions of vagueness call for different styles of logical representation. The natural understanding of The Good,[1] for example, partitions the predicates of the language into two sorts: the completely precise (as perhaps "has at least 25 hairs/cm^2 throughout the head") and those which describe a *range* of the properties denoted by the precise predicates. For each object, every precise predicate is such that either it or its negation is true of that object. With respect to the non- precise predicates, The Good claims that under some ways of "understanding" or "precisifying" them in terms of the precise predicates, they might be true of an object but under other "precisifications" they might be false of that same object. Nonetheless, some versions of The Good might maintain, there can be sentences containing only non-precise predicates which are true under all ways of precisifying the predicates. The natural way to regiment such a conception would be to treat the sublanguage containing the completely precise predicates as entirely classical and invoke some unusual semantic evaluation procedure for sentences containing the vague predicates. For this, such methods as supervaluation spring to mind.

Now, The Bad considers vagueness to be due to a lack of knowledge or information. Each sentence *is* actually either true or false but there is not sufficient grounds at the time of utterance for the hearer to determine which -- possibly because exactly which proposition is being expressed is not determinable. So far as the hearer/reader is concerned, on the basis of available information s/he can imagine it true and also imagine it false. Given a context where sentence *p* is known to be true (or known to be false), then *p* is not vague in that context; however, that context might have sentence *q* which is true (or is false) but not so known and *q* would then be vague in that context. On this picture, logical truths and falsehoods come out being the only non-vague sentences when there is no context. This explanation ought to suggest some sort of modal logic interpretation of vagueness. Indeed, given the empty context where only logical truths and falsehoods are non-vague, it ought to suggest a modal logic with contingency and non-contingency operators.[2]

Finally, the presuppositions of The Ugly's conception of vagueness commit such a theorist to the use of a many-valued logic of some sort. Just which one is a detail for the inner workings of the theory, of course, but that such a person must have one is determined by the view that vagueness is a property of *reality* rather than (say) due to our ignorance of reality [The Bad] or (say) due to "a deficiency in preciseness" of the terms describing reality [The Good]. The realistic conception of vagueness [The Ugly] views vague objects as if they had the property: being neither F nor not-F, for some property F. The question is: how can such a theory represent this? Let *Fa* be a proposition which asserts of an object named *a* that it has a property *F*. According the The Ugly, either the object has the property (definitely), or the object lacks the property (definitely), or else it neither has nor lacks it -- and this last fact is, in its own way, just as definite as the former two. It is a feature of the object, and as such is to be represented as being on a par with them. If *a* has property *F*, then *Fa* is true; if *a* lacks *F*, then *Fa* is false; so in the case where *a* neither has nor lacks *F*, *Fa* must take on some other truth value.[3] To be sure, there is a sense in which there are vague objects in The Good. Leslie, the person with an intermediate number of head hairs, is vague on this approach because some ways of precisifying 'bald' make it be true that Leslie is bald while other precisifications make it be false. But from the point of view of The Ugly, this is an uncritical and unjustified semantic descent. For, The Ugly would insist, according to The Good, the only "basic, real" properties are the completely precise ones, and there are no objects vague with respect to these. The Good counts 'bald' and the like as being derivitive -- defined or explained in terms of groups of other properties, and not as direct, primitive properties of objects. As the example about Hamlet was to illustrate, The Ugly believes that even a completely precise predicate might neither apply nor not apply to a given object, and that's what makes the object vague.

As I have said, the property of being neither F nor not-F is a real, primitive and basic property according to The Ugly. Some objects, the vague ones, have such properties, and to describe this situation we must invoke a truth functional logic (many-valued perhaps, but still truth functional). Now since Quine, Strawson, Geach, and others, we are accustomed to thinking of identity as crucial to individuating objects. The Ugly also follows this dictum. An object named *a* is vague if and only if some sentence like "it is indefinite whether *a* is identical to *b*" is true (for some name *b*).

The remainder of this paper will investigate possible logics which might be employed by The Good, The Bad, and The Ugly. It will be my contention that no many-valued logic can account for The Ugly's conception. But since The Ugly is committed to some truth functional, many valued logic, it follows that this conception is incoherent -- as Russell, Evans, Dummett, and Margalit had claimed. I will then present a series of logics suitable for The Bad, and urge one as being preferable to the others. This shows that while The Bad might be bad, it is not incoherent, and can be studied logically. Finally, I will sketch a method of evaluation which seems suitable for The Good. The ideas are similar to supervaluation techniques, but the details and overall results are quite different. The reason for wanting different results is due to some psychological data collected by Len Schubert[4] and myself which show that the predictions made by supervaluation theories just plain do not accord with how we ordinarily assign truth values to compound sentences which employ vague terms.

So let us start with The Ugly.

Many-Valued Logics for Vagueness:

There are a lot of many-valued logics. Such logics might be distinguished by the number of truth values, by how many of these truth values are counted as "designated" (i.e., "really true"), and by what the interpretations of the connectives are, amongst other things. Rather than try to go through them all one-by-one, I shall give a general argument which holds against any logic that obeys certain principles. All the well-known, finitely many-valued logics obey these principles. So I conclude that none of them can be used by The Ugly. Afterwards, I turn my attention to infinite-valued logics ("fuzzy logics") to indicate why these too are unsuitable for The Ugly.

We shall start this argument against The Ugly by looking at the 3-valued case and then generalizing it to any finitely-many-valued logic. In a 3-valued logic there are three truth values: 1 ("completely true"), 2 ("intermediate"), and 3 ("completely false"). In any version of this logic suitable for discussing vague objects, we must be able to "talk about" individuals and ascribe properties to them; so the logic must have names, quantifiers, variables, predicates, etc. In addition, and following those philosophers who have tried to develop many-valued logics for vagueness[5], we wish to *assert*, in the language, that an object is vague, or that a proposition is definitely true or definitely false. We therefore will want to introduce some method for doing this; I here adopt J-operators.[6] For every possible truth value *n* (here: '1', '2', or '3') there is a symbol J_n which is a two-valued, truth-functional operator on sentences. Intuitively, a sentence of the form $J_n p$ says that *p* has exactly the value *n*. Such a sentence is definitely true (= has value 1) if and only if *p* actually does have the value *n*. And it is definitely false (= has value 3) otherwise. These J-operators have the following properties:

(1) $[J_i p]$ is either 1 or 3, for any i (I use [ ] to indicate the semantic value of an expression)

(2) $[J_2\ J_i p] = 3$, for any i (since $J_i p$ must be either 1 or 3)

(3) if every sentential part of p is in the scope of a J-operator, then either $[J_1 p] = 1$ or $[J_3 p] = 1$,

and hence $[\sim J_2 p] = 1$

(4) $[J_1 p \vee J_2 p \vee J_3 p] = 1$ (every formula takes one of these values)

(5) $[\sim(J_i p \wedge J_k p)] = 1$, when $i \neq k$ (no formula takes more than one of these values)

(6) $[J_1 p \rightarrow p] = 1$ (which gives the interpretation of J_1 as being "most true")

(7) if $[p \leftrightarrow q] = 1$, then $[J_i p \leftrightarrow J_i q] = 1$, for any i

Of course, not every 3-valued logic with J-operators will obey these conditions. For example, a logic with no 'v' can hardly be expected to be able to state condition (4) much less obey it. So really we are just talking about 3-valued logics with J-operators in which (1)-(7) hold. Even though this may not be all the 3-valued logics, it contains all the interesting ones, at least if we are willing to broaden our outlook so that (for instance) we needn't have an 'v' symbol, but just a way to *define* a connective for which property (4) holds. Similarly, not every 3-valued logic has a '~', or a '∧' , or a '→', or a '↔' which obey the conditions stated. Still, if there is a combination of the available connectives which do obey properties (1)-(7), then we are talking about that logic and will show that it is not suitable as an account of vagueness.

Besides properties (1)-(7), we want the J-operators to interact with quantifiers in the following way, so that these formulas take value 1:

(8) $J_1(\forall x)Fx \leftrightarrow (\forall x)J_1Fx$

(9) $J_3(\forall x)Fx \leftrightarrow (\exists x)J_3Fx$

(10) $J_2(\forall x)Fx \leftrightarrow ((\exists x)J_2Fx \wedge \sim(\exists x)J_3Fx)$

Now, since The Ugly wishes to talk about objects, and since the identity of objects is crucial to this talk, we will add on some principles for identity, namely a statement that self-identity is definite ("reflexivity of identity", or Ref for short) and a statement of "Leibniz's Law" (LL for short): The following formulas have the value 1

(Ref) $J_1\ a=a$

(LL) $(a=b \leftrightarrow (\forall F)(Fa \leftrightarrow Fb))$

I have stated (LL) in a "second order" form -- a and b are identical just in case they share the same "real" properties. What are real properties? According to The Ugly they include statements such as "neither has nor lacks G", or "is vaguely identical to c", and the like. Now, if we were working in a modal logic, one might argue that certain formulas that can be constructed (e.g., "Kim believes that x is a spy") do not designate "real" properties -- ones directly true of objects. That this is so is claimed to follow from the fact that even if it is true when 'the shortest spy' is substituted for 'x', it might not be true if 'Orcutt' is so substituted, even if Orcutt *is* the shortest spy. (Or phrased in other words, the claim is that "Kim believes that x is a spy" is not a transparent context.) The underlying reason for all this is the fact that the truth value of the entire sentence does not depend solely on the references of the terms embedded in the sentence. But in the case under consideration -- representation of realistic vagueness using J-operators -- the truth of each sentence *does* depend entirely upon the references of the embedded parts. Many-valued logics with J-operators are truth functional, albeit with more than two truth values. Thus, if we assume that the atomic sentences contain only predicates which describe "real properties" (such as "is vague") then it is legitimate to allow J-operator formulas to describe "real properties" also, and to use them in (LL). The importing of three values and J-operators does not change anything: their use must be granted by the vagueness-in-nature-theorist.

Recall now that The Ugly holds that identity can be vague. In a 3-valued framework, this amounts to saying that (J_2 a=b) might be definitely true, that is [J_2 a=b] = 1, for some vague object. Therefore, let us suppose this and use our previous principles.

a. $a{=}b \leftrightarrow (\forall F)(Fa \leftrightarrow Fb)$	(LL)
b. $(J_2\, a{=}b \leftrightarrow J_2(\forall F)(Fa \leftrightarrow Fb))$	from a, and (7)
c. $J_2(\forall F)(Fa \leftrightarrow Fb)$	from b, and supposition
d. $(\exists F)J_2(Fa \leftrightarrow Fb) \wedge \sim(\exists F)J_3(Fa \leftrightarrow Fb)$	from c, and (10)
e. $\sim(\exists F)J_3(Fa \leftrightarrow Fb)$	from d
f. $(\forall F)[J_1(Fa \leftrightarrow Fb) \vee J_2(Fa \leftrightarrow Fb)]$	from e, quant.neg, and (4)
g. $J_1(J_1 a{=}a \leftrightarrow J_1 a{=}b) \vee J_2(J_1 a{=}a \leftrightarrow J_1 a{=}b)$	from f, instantiate F to $(\lambda x)J_1 a{=}x$
h. $J_1(J_1 a{=}a \leftrightarrow J_1 a{=}b)$	from g, and (3)
i. $J_1 a{=}a \leftrightarrow J_1 a{=}b$	from h, and (6)
j. $J_1 a{=}b$	from i, and (Ref)
k. $\sim J_2 a{=}b$	from j, and (5)

Which contradicts our assumption J_2a=b. Of course contradictions in many valued systems are not what they are in two-valued systems, but here it is a legitimate contradiction, since J_2a=b takes either the value 1 or the value 3 (being a J-formula). Hence the negation of the formula does actually contradict it, and shows that (with our 10 principles governing J-operators and quantifiers, and with LL and Ref) a three-valued system cannot admit vague identity -- and therefore cannot admit vague objects.

The argument as stated uses only (LL), (Ref), and principles (3), (4), (5), (6), (7), and (10). Indeed, it can be stated even more simply this way:

> If a and b are vaguely identical, then for each property F, either they definitely both have F or both lack F, or else it is indefinite whether they both have F or

> both lack F. (This is what step (f) says). This is therefore true for the property of being definitely identical to a. But by principle (3), for such a property it cannot be indefinite whether they both have or both lack it. Since a *is* definitely identical to a (by Ref), it follows that b must be also. And hence it is not indefinite whether a and b are identical.

It would be extremely difficult for a vagueness-in-reality-theorist to deny any of the principles which describe the use of J-operators. I suppose they might try to deny (Ref) or (LL) -- but if they do, I urge that we just ignore them. A possible strategy (see van Inwagen 1986) is to deny the move from (f) to (g) in the argument; deny, that is, that $(\lambda x)J_1 a=x$ is a legitimate substitution instance for the property variable (or equivalently, to deny the equivalence of "a has the property $(\lambda x)Fx$" and "Fa" -- which is in fact van Inwagen's 1986 tactic, and also employed by Chierchia 1982 and by Schubert & Pelletier 1987a, but for different reasons). I have already tried to explain why such a move is not available to The Ugly, and still stand by it. The property "$(\lambda x)J_1 a=x$" is, according to The Ugly, a "real and basic" property which holds "directly" of individuals. It is not an intensional property (or any other type of modal property) in this theory -- if it were treated as such, the resulting theory would be The Bad, not The Ugly. van Inwagen (1986) says that the "friends of vagueness-in-nature will be understandably hostile" towards identification of "a has the property $(\lambda x)Fx$" with "Fa". But other than for purposes of avoiding anti-Ugly arguments, I fail to see why -- they are, after all, firmly committed to this principle, at least for "real and basic" properties F. Invoking new methods of semantic evaluation of "a has the property $(\lambda x)Fx$" different from the direct ascription of a property to an individual, as in "Fa", merely admits that the theory is not suitable for The Ugly. Perhaps it can be appealed to by The Good or by The Bad, but certainly not by The Ugly.

Many authors have recognized that it is implausible to try to account for vagueness by replacing one sharp cutoff (= either has F or lacks F) by two sharp cutoffs (= either has F or lacks F or neither has nor lacks F). Dummett (1975) says:

> Thus "hill" is a vague predicate, in that there is no definite line between hills and mountains. But we could not eliminate this vagueness by introducing a

> new predicate, say "eminence", to apply to those things which were neither definitely hills nor definitely mountains since there would still remain things which were neither definitely hills nor definitely eminences, and so ad infinitum.

Some might try to use this argument directly against The Ugly, by claiming that this theory does not have the resources to talk about such "in between", "higher-level" cases of vagueness. And perhaps this is true if they were to stick to the "either has F or else lacks F or else neither has nor lacks F" trichotomy. But The Ugly need not stick to this three-valued system. At least if they do not take Dummett's "ad infinitum" seriously, they might simply use more truth values -- say, 7 ± 2 or 300,000. So, let us consider finitely-many-valued systems in the abstract. For reasons just like those given above in discussing three-valued systems, even in this more general case, The Ugly is still committed to truth-functionality and to the legitimacy of using J-formulas as values of predicate variables. That is, since it views such statements as "is a hill to the eminence degree" as direct and basic properties of items of reality, it must allow lambda-abstraction over such predicates to directly apply to individuals (and with no unusual evaluation techniques).

So the question is: are all finitely-many-valued logics susceptable to the type of argument given above against a three-valued logic? The answer is yes. Given a many valued logic, with truth values from 1 ("most true") through some finite n, if it obeys the following analogues of the above principles, such an argument can be constructed.

(1´) $[J_i\, p]$ is either 1 or n, for all i

(2´) $[J_k\, J_i\, p] = n$, for any i and for any k other than 1 and n

(3´) if every sentential part of p is in the scope of a J-operator, then either $[J_1 p] = 1$ or $[J_n p] = 1$, and hence $[\sim J_i p] = 1$ for $1<i<n$

(4´) $[J_1 p \vee J_2 p \vee \ldots \vee J_n p] = 1$ (every formula takes one of these values)

(5´) $[\sim(J_i p \wedge J_k p)] = 1$, when $i \neq k$ (no formula takes more than one of these values)

(6´) $[J_1 p \rightarrow p] = 1$ (which gives the interpretation of J_1 as being "most true")

(7´) if $[p \leftrightarrow q] = 1$, then $[J_i p \leftrightarrow J_i q] = 1$, for any i

Just as in the three-valued case, we need use only one of the principles concerning the interaction of quantifiers with J-operators, namely the following set of axioms (different ones for different values of i and k).

(10´) $(J_k(\forall x)Fx \leftrightarrow ((\exists x)J_kFx \wedge \sim(\exists x)J_iFx))$, for any k other than 1 and n, and (given a k) for each i>k

Intuitively, (10´) says that it is true to degree k that everything is F just in case there is something of which it is true to degree k that it is F and furthermore there is nothing of which it is true that it is F to some degree which is "more false" than k. This, I think, accords well with our understanding of how J-operators should interact with quantifiers. Finally, we retain (Ref) and (LL) as stated above. Now quite obviously the same argument will go through. I state it here informally, following the informal statement of the three-valued argument. The "friends of vagueness in nature" will claim that there is some "vague identity" statement: that is, that a sentence like J_i a=b is true for some value of i between 1 and n. Let it be any such value. Then principle (10´) says that

> for every property F, (Fa ↔ Fb) cannot be true to any degree less true than i. Therefore, $(\forall F)[J_1(Fa \leftrightarrow Fb) \vee J_2(Fa \leftrightarrow Fb) \vee ... \vee J_i(Fa \leftrightarrow Fb)]$. Now instantiate F to the property of being definitely identical to a. By principle (3´), none of the disjuncts of this formula can be true, except for the first one. Thus we have $J_1(J_1a{=}a \leftrightarrow J_1a{=}b)$. By principle (6´) we have $(J_1a{=}a \leftrightarrow J_1a{=}b)$, and by (Ref) we have $J_1a{=}b$. But then by principle (5´) we have $\sim J_ia{=}b$, a contradiction.

The above result shows that no finitely-many valued logic can hope to account for The Ugly's ontological views about vagueness. But there are still infinitely-many valued logics to consider ("fuzzy logics"). In such logics, we have an infinity of possible truth values a statement might take, from 1 ("most true") to 0 ("most false"). Which infinity shall we choose? Most theorists, following the work of Zadeh, use all real numbers. However, this choice is not available for those theorists who want to talk about vagueness within the logic, since they could not make use of J-operators (or any other such methods of talking about "intermediate" truth values). For, with a non-denumerable number of values, we would require a non-denumerable number of such operators, *per impossible*.

We therefore will use the rational numbers between 1 and 0, and introduce a denumerable number of J-operators. But there still remains need for some other argument against such logics than a simple generalization of the one given above. That we require a different treatment is most easily seen by looking at principle (10´), relating universal quantifiers and J-operators of values other than 1 and 0. It is possible in fuzzy logic for there to be a predicate F which is not exemplified by any particular object, and yet the objects become asymptotically closer and closer to being F. For example, F might be the property defined on the positive integers: 1/x = 0. Now, there is no natural number x which satisfies this property, although one might argue (as fuzzy logicians do) that as x increases the semantic value of [Fx] gets closer and closer to 1. That is, [F(3)]<[F(20)]<[F(1000)], etc. What, then, about the semantic value [(∃x)Fx]? Most fuzzy logicians wish this to be 1, on the grounds that the limit as $x \to \infty$ of 1/x equals 0. Generally, such logicians say (recall that they are working with real numbers between 0 and 1), the semantic value of an existentially quantified formula should be the least upper bound of the values of all possible arguments; and the semantic value of a universally quantified formula should be the greatest lower bound. But we cannot use this in our fuzzy logics for vagueness because not every such bound is a rational number. However, if the least upper bound *is* a rational number, then $[J_i(\exists x)Fx]$ can be 1 without $[(\exists x)J_iFx]$ being 1. With respect to principle (10´) therefore, we can have $[J_k(\forall x)Fx]=1$ without $[(\exists x)J_kFx]$ =1 -- or put differently, we can have the values of Fx approach k in the limit without any of them actually being equal to k. The question is: what should those who wish to use fuzzy logics for accounts of vagueness (and therefore require the rational truth values) assign to such formulas? There seem to be two choices: they could deny that formulas like (∃x)Fx have a truth value in general (whenever the value it would have is irrational then say it is undefined), or they could say that its value is 0 (on the grounds that there is no argument for which the predicate holds).

It doesn't matter which option our theorists choose. Recall that we are interested in some kind of version of (10´); the entire strength of (10´) is not needed for the analogue of the above argument against vague identity to go through. All we need is this set of axioms:

(10´´) $[J_k(\forall x)Fx \to \sim(\exists x)J_iFx] = 1$ for any rational $k \neq 0$, and (given a k) for each rational $i<k$

That is, if a universally quantified claim is true to (rational) degree k, then there can be no object of which the claim is true to a (rational) degree less than k. This principle is correct no matter which of the above alternatives is chosen.

The analogues of the other principles (1´)-(7´) necessary for the argument to go through are straightforward (keep in mind that all possible truth values, and hence all possible subscripts for J-operators, are rationals between 0 and 1 inclusive):

(1´´) $[J_i\ p]$ is either 1 or 0, for all i

(2´´) $[J_k\ J_i\ p] = 0$, for any i and for any k other than 1 and 0

(3´´) if every sentential part of p is in the scope of a J-operator, then either $[J_1 p] = 1$ or $[J_0 p] = 1$, and hence $[\sim J_i p] = 1$ for $0<i<1$

(4´´) $[J_i p] = 1$, for some $0\leq i\leq 1$ (every formula takes one of these values)

(5´´) $[\sim(J_i p \wedge J_k p)] = 1$, when $i\neq k$ (no formula takes more than one of these values)

(6´´) $[J_1 p \rightarrow p] = 1$ (which gives the interpretation of J_1 as being "most true")

(7´´) if $[p \leftrightarrow q] = 1$, then $[J_i p \leftrightarrow J_i q] = 1$, for any i

(Ref) $J_1\ a=a$

(LL) $a=b \leftrightarrow (\forall F)(Fa \leftrightarrow Fb)$

Let us present the argument formally: The friends of vagueness in nature think that $J_i\ a=b$ is true for some i strictly between 0 and 1. Consider any such i; then

a. $a=b \leftrightarrow (\forall F)(Fa \leftrightarrow Fb)$	(LL)
b. $(J_i\ a=b \leftrightarrow J_i(\forall F)(Fa \leftrightarrow Fb))$	from a, and (7´´)
c. $J_i(\forall F)(Fa \leftrightarrow Fb)$	from b, and supposition
d. $\sim(\exists F)J_k(Fa \leftrightarrow Fb)$ (for all k<i)	from c, and (10´´)
e. $(\forall F)[J_1(Fa \leftrightarrow Fb) \vee .... \vee J_i(Fa \leftrightarrow Fb)]$	from d, quant.neg, and (4´´)
f. $J_1(J_1 a=a \leftrightarrow J_1 a=b) \vee .... \vee J_i(J_1 a=a \leftrightarrow J_1 a=b)$	from e, instantiate F to $(\lambda x)J_1 a=x$
g. $J_1(J_1 a=a \leftrightarrow J_1 a=b)$	from f, and (3´´)
h. $J_1 a=a \leftrightarrow J_1 a=b$	from g, and (6´´)
i. $J_1 a=b$	from h, and (Ref)
j. $\sim J_i a=b$	from i, and (5´´)

And once again we find a legitimate contradiction with the supposition that there might be vague identity. Therefore the friends of vagueness in reality cannot take refuge in fuzzy logics any more than they could in finitely-many valued logics.[7] But since they are required, by their ontological outlook on vagueness, to adopt *some* many-valued logic, I conclude that the notion of vagueness *in re* -- vague objects and their ilk -- is incoherent.

Modal Logics for Vagueness:

The Bad considers vagueness to be due to a "lack of the total information required" for an agent to determine the truth value of a statement. As mentioned above, the intuition here is that every statement either *is* true or *is* false, but the agent sometimes cannot determine which. In this section we will look at logics which seem suitable for capturing this intuition.

Accounts of The Bad presume a "conversational context". In such a context, some statements are *known* (by the participants) to be true or are *known* to be false, that is (as we shall say), these statements are *definite*. In this context however, other statements are *indefinite* (or *vague*) -- not known to be true and not known to be false. However, certain statements are definite in *any* context. Certainly the logical truths are definite, and certainly the logical falsehoods are definite. It is the logic induced by "the empty conversational context" -- that context in which only the logical truths and logical falsehoods are definite, and all other statements are considered vague -- in which we shall be interested. This will give the underlying logic behind The Bad.[8]

We shall introduce a sentence operator, ˚, which will mean "it is definite whether". We start an investigation of this logic by considering certain well-known modal formulas to see whether or not they are true with the present interpretation of ˚. I will divide these formulas into three types: Type A formulas are those for which the interpretation clearly holds, Type B formulas are those for which the interpretation clearly fails (sometimes a short justification will be given), and Type C formulas are those about which one might have no clear feelings. These last will be discussed in a little more detail later, and I will make pronouncements upon them. After giving this classification we will look at a semantic characterization of the logic(s) generated, with an eye towards showing what kinds of possible world structures are involved in modal logics for vagueness.

So let us start by giving some Type A principles. The operator ◊ is to be interpreted as "it is vague whether" or "it is not definite whether". I give them their popular name (as listed in Chellas

1980); for those with no popular name, I invent one. Keeping in mind that $^\circ p$ means "p is definite" but does not make any claim about p's truth, we have the following principles.

(Def◊) $\Diamond p \leftrightarrow \sim{}^\circ{\sim}p$

(I) $^\circ p \leftrightarrow {}^\circ{\sim}p$

(I´) $\Diamond p \leftrightarrow \Diamond{\sim}p$

(I´´) $^\circ p \leftrightarrow \sim\Diamond p$

(I), (I´), and (I´´) are equivalent in the presence of (Def◊), and will hereafter be referred to as (I). If p and q are logically equivalent then one is definite/vague just in case the other is. The two principles given are equivalent (given the above) and will be referred to in the sequel as (RE)

(RE$^\circ$) if $\vdash (p \leftrightarrow q)$ then $\vdash ({}^\circ p \leftrightarrow {}^\circ q)$

(RE◊) if $\vdash (p \leftrightarrow q)$ then $\vdash (\Diamond p \leftrightarrow \Diamond q)$

If a formula is a logical truth then it is definite

(RN) if $\vdash p$ then $\vdash {}^\circ p$

If each conjunct is definite then the conjunction must be also

(C) $({}^\circ p \wedge {}^\circ q) \rightarrow {}^\circ(p \wedge q)$

Some Type B principles of this logic -- principles which fail -- are:

(T) $^\circ p \rightarrow p$

(P) $p \rightarrow \Diamond p$

(D) $^\circ p \rightarrow \Diamond p$

(T) fails because a statement can be definite without being true (can be "definitely false"); (P) fails because some true statements are not vague; and (D) obviously fails with the intended meanings of $^\circ$ and ◊. Two other Type B principles are

(M) $^\circ(p \wedge q) \rightarrow ({}^\circ p \wedge {}^\circ q)$

(K) $^\circ(p \rightarrow q) \rightarrow ({}^\circ p \rightarrow {}^\circ q)$

Principle (M) fails if, for example, p=A and q=~A, for atomic A. Here $^\circ(A \wedge {\sim}A)$ is true (that is, $(A \wedge {\sim}A)$ is definite) but neither A nor ~A is definite. (K) fails if $p=(A \wedge {\sim}A)$ and q=B, for atomic B. Here $^\circ((A \wedge {\sim}A) \rightarrow B)$ is true, $^\circ(A \wedge {\sim}A)$ is true, but $^\circ B$ is not true.

So far, what I have said about a modal logic for vagueness yields the logic ECNI. (Using Chellas's notation: it is the logic containing RE, C, RN, and I). Since it contains (RE) and (Def◊) it is *classical* in Segerberg's (1971) sense, and therefore can be given an analysis in terms of possible worlds. Of course, since principle (K) is not in the logic, the resulting logic is not *normal* and hence there is no normal, relational, possible worlds semantics. But there can nonetheless be a possible world semantics done by the "Montague-Scott" (or "neighborhood" or "minimal model") method. We first consider the logic as axiomatized by the propositional logic, (C), and (I), with the rules of inference (RE) and (RN). A model is a triple M=<**W**,**N**,**P**> such that **W** is a set of indices ("possible worlds"), **P** is a mapping from natural numbers to subsets of **W** (i.e., **P**(n)   **W**) for each natural number n -- telling us for each atomic proposition **P**(n) which subset of **W** it is true in), and **N** is a mapping from **W** to sets of subsets of **W** (i.e., **N**α   ℘(**W**) for every world α∈**W** -- that is, what propositions (subsets of worlds) are necessary at α). Define ˚p to be true at an index α in M if and only if the set of indices at which p is true (denoted as [p]) is a member of **N**α; and ◊p to be true at α if and only if (**W**-[p])∉**N**α. It is well known that propositional logic, (RE) and (Def◊) are valid in any class of such models. It remains only to find that subclass determined by (RN), (C), and (I). It is again well-known that (RN) holds when M *contains the unit*, i.e.,

(n) $\mathbf{W}\in\mathbf{N}\alpha$

for every α∈**W** (anything true at all worlds is an element of the necessitation of any world); and that (C) holds if M is *closed under intersections*, i.e.,

(c) if $X\in\mathbf{N}\alpha$ and $Y\in\mathbf{N}\alpha$ then $(X\cap Y)\in\mathbf{N}\alpha$

for every α∈M and all sets of indices X and Y. I dub the property which validates (I) as *contrariety* (if something is necessary so is its opposite)

(i) $X\in\mathbf{N}\alpha$ iff $((\mathbf{W}-X)\in\mathbf{N}\alpha)$

Standard methods (cf. Chellas 1980, Chap. 7) would clearly suffice to show that ECNI is determined by the class of contrary models that are closed under intersections and contain the unit. It is also obvious that principles (K), (M), (T), and (P) are not universally valid in this class of models. In the scheme of modal logics we find ECNI located:

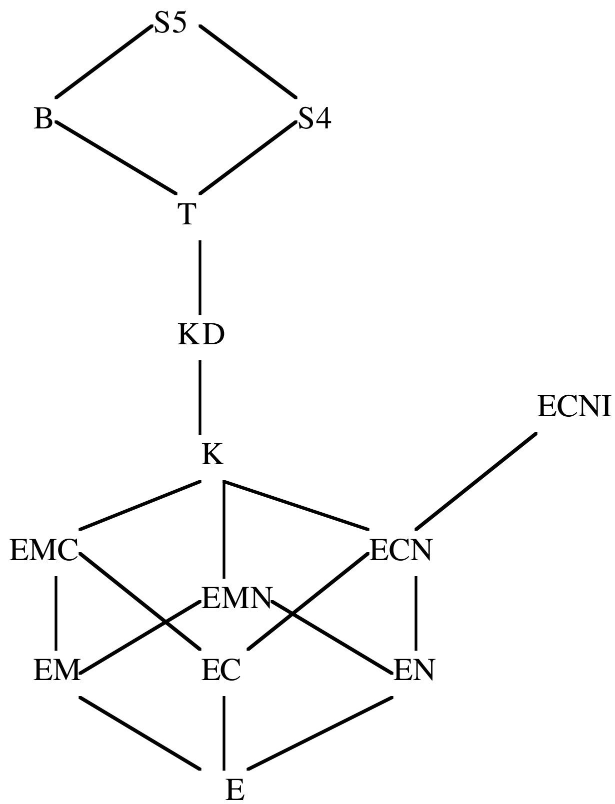

Fig. 1: A map of some modal logics.

What other theses might one suggest for a logic of indeterminacy? We've seen that (M) does not hold, but a closely related one does seem plausible:

(M*) $(^{\circ}(p\wedge q) \wedge (p\wedge q)) \rightarrow (^{\circ}p\wedge {}^{\circ}q)$

("if a conjunction is definitely true, then each conjunct must be definite"). The semantic condition corresponding to (M) is called *supplementation*

(m) if $X\in\mathbf{N}\alpha$ and X   Y then $Y\in\mathbf{N}\alpha$

For our weaker (M*), I recommend the name *partial supplementation*

(m*) if $X\in\mathbf{N}\alpha$ and $\alpha\in X$ and X   Y then $Y\in\mathbf{N}\alpha$

Since the logic ECNM* (without the I) is a sublogic of K and a superlogic of ECN, it falls on the line between K and ECN in Fig. 1. Our logic of vagueness is now ECNM*I, still independent of K (due to the presence of (I)), but a superlogic of ECNM*.

Some Type C principles -- other popular principles which might be considered for our logic of vagueness -- are

(4) ˚p → ˚˚p

(B) p → ˚◊p

(G) ◊˚p → ˚◊p

(5) ◊p → ˚◊p

(U) ˚(˚p → p)

If one takes the view that all indices are "accessible" to any other, so that ˚ means "is either true at all indices or false at all indices" and ◊ means "is true at some index and false at some index", then one will have these four principles (which are all equivalent, given (I) and (Def◊).

(V) ˚˚p

(V´) ~◊˚p

(V´´) ˚◊p

(V´´´) ~◊◊p

i.e., whether p is definite or vague is itself always definite. For, if p is definite then p is either true at all indices or false at all, and hence ˚p is true at all indices -- i.e., ˚˚p. On the other hand, if p is vague then it is true at each index that p is true at some index and false at some index -- i.e., ˚◊p. Given then one of the (V) principles, we can see why the other Type C principles would hold: their consequent is a (V) principle.[9]

The logic which contains all these Type C principles is axiomatized as ECNM*IV. The semantic condition on models for (V) is *all-pervasiveness* (the necessity of each proposition is necessary):

(v) $\{\beta : X \in \mathbf{N}\beta\} \in \mathbf{N}\alpha$

and the logic is determined by the class of all-pervasive, contrary, partially supplemented models that are closed intersections and contain the unit. Of course we could withhold (V) and instead add (4) to ECNM*I, or we could add

(B´) p → ˚(˚p ∨ p)

These considerations yield the following group of modal logics:[10]

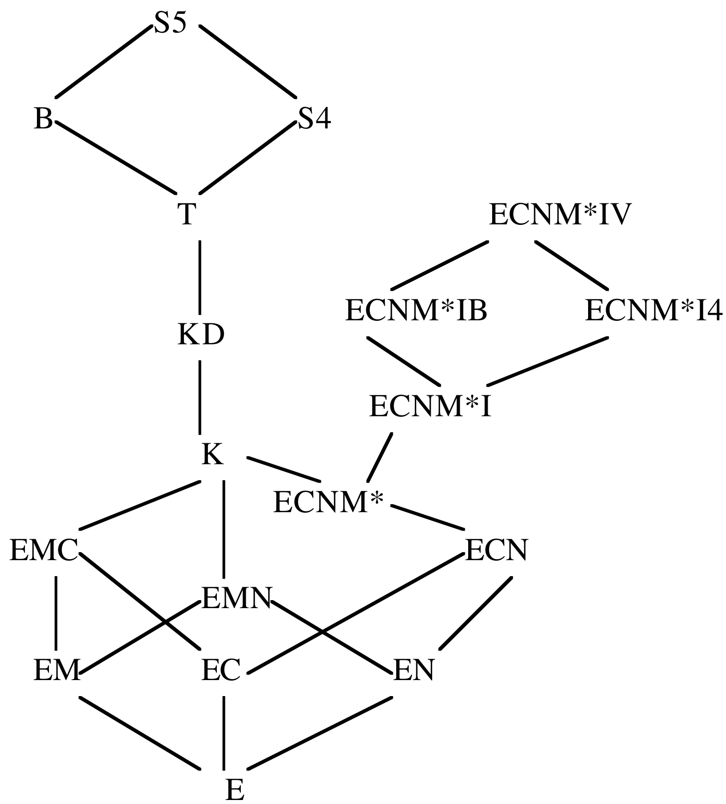

Fig. 2: Where logics of vagueness fit into the map of modal logics

I take the existence of "higher order vagueness" -- the possibility that a proposition might be definite, but not definitely so; that it can be vague whether a proposition is vague -- to count against any of (4), (B), (G), (5), (U), (V). Thus it seems to me that we should allow

˚p ∧ ~˚˚p

and also

◊p ∧ ~˚◊p

and so on, for any number of iterations of the operators ˚ and ◊. If we allow that all of these can happen, we shall want no reduction laws of the sort mentioned by (4), (5), etc. In light of this "higher order vagueness", I would recommend the logic ECNM*I as suitable for The Bad. I have not investigated the logic in detail, but mention here two theorems which might not have been obvious:

◊˚p ↔ ◊◊p

˚◊p ↔ ˚˚p

I think it is now time to turn our attention to the possible addition of names, predicates and quantifiers to this logic, and see whether The Bad can avoid arguments like the one used above to undermine The Ugly. The Bad wishes to retain (Ref) and (LL) in the forms

(Ref) $\;^{\circ}a=a$

(LL) $a=b \leftrightarrow (\forall F)(Fa \leftrightarrow Fb)$

But it is not obvious what principles governing the interaction of quantifiers with modal operators The Bad should adopt. The problem is in the reading of ˚ and ◊ as the *epistemic* operators "it is definite/indefinite whether". Consider for example a simple attempt to state some such principle

(α) $\;^{\circ}(\exists x)Fx \rightarrow (\exists x)\,^{\circ}Fx$

The principle fails because it might be quite definite that some ticket will win but there may be no ticket of which it is definite that it will win. Similarly

(β) $\Diamond(\exists x)Fx \rightarrow (\exists x)\Diamond Fx$

fails, since it can easily be indefinite whether there is an F, but there not be anything of which it is indefinite that that object is F. (Possibly because the person doesn't know whether all objects have been considered).[11]

However, for the purposes of constructing an argument akin to the anti-Ugly argument, what we need is a principle which tells us what $\Diamond(\forall x)Fx$ implies. Here is one possibility. I do not wish to put too much weight on it -- it seems to me that friends of The Bad might very well wish to reject it. Still, it has a certain plausibility and allows an anti-Bad argument to get started. Even so, The Bad emerges unscathed.

(γ) $\Diamond(\forall x)Fx \rightarrow (\forall x)(Fx \vee \Diamond Fx)$

To motivate this principle, let us note that if we were to inspect objects one by one, one way to show $\Diamond(\forall x)Fx$ false would be to come up with an object which was not-F and definitely so. The other way to show it false would be to check all objects, knowing that these were all the objects, and discover that they each were F and definitely so. But this last way cannot be stated in the language: $(\forall x)(Fx \wedge \,^{\circ}Fx)$ -- the "closest" we can come to saying this -- only says that all objects are in fact F, and for each of them it is definite that they are F. But it does not assert that we know this to be all the objects; so $(\forall x)(Fx \wedge \,^{\circ}Fx)$ might be true at the same time $\Diamond(\forall x)Fx$ is true. We therefore state the implication in one direction only. Let us now see how much of the anti-Ugly argument can be reconstructed against The Bad. We assume $\Diamond a=b$; then

1.	$a=b \leftrightarrow (\forall F)(Fa \leftrightarrow Fb)$	(LL)
2.	$\Diamond a=b \leftrightarrow \Diamond(\forall F)(Fa \leftrightarrow Fb)$	1 and (RE)
3.	$\Diamond(\forall F)(Fa \leftrightarrow Fb)$	2 and supposition
4.	$(\forall F)([Fa\leftrightarrow Fb] \vee \Diamond[Fa\leftrightarrow Fb])$	3 and (γ)
5.	$(^{\circ}a=a\leftrightarrow {}^{\circ}a=b) \vee \Diamond(^{\circ}a=a\leftrightarrow {}^{\circ}a=b)$	4, instantiate to $(\lambda x)^{\circ}a=x$

At this stage, in the anti-Ugly argument, we eliminated the second disjunct of 5 since no formula all of whose sentential parts were in the scope of J-operators could have any "intermediate" J-operator applied to it. The same would happen here if we were working in the logic embodying principle (V), in which "every index was accessible to every other". In that logic we could continue

6.	$^{\circ}a=a \leftrightarrow {}^{\circ}a=b$	5 and (V)
7.	$^{\circ}a=a$	(Ref)
8.	$^{\circ}a=b$	6 and 7
9.	$\sim\Diamond a=b$	8 and (I)

which is a contradiction. But even this logic has something they can object to in the argument. However, before getting to this, let me emphasize that a crucial step in the anti-Ugly arguments cannot be mirrored in The Bad's favoured logic ECNM*I: we cannot go from step (5) to step (6). The favoured logic allows iterated modalities -- unlike many-valued logics -- and can easily assign a meaning to them. Indeed, The Bad insists that this is precisely what happens. Note that after step (5) we could continue the argument in the following vein.

6´.	$\sim{}^{\circ}a=b$	from supposition and (I)
7´.	$^{\circ}a=a$	(Ref)
8.´	$\sim(^{\circ}a=a \leftrightarrow {}^{\circ}a=b)$	6´ and 7´
9´.	$\Diamond(^{\circ}a=a \leftrightarrow {}^{\circ}a=b)$	5 and 8´

So according to The Bad, even ignoring the further defense still to come, we should expect it to be indefinite whether a's definite self-identity is equivalent to it's definite identity with b, when a and b's identity is indefinite. And surely this is what is desired. Recall that The Bad thinks that either it *is* true or it *is* false that a=b. If we are epistemically unable to tell which, then we are epistemically unable to say whether the claim that a and b are definitely identitical is the same as saying that a and a are definitely identical.

As I said, there is yet another response available to The Bad. That response is to challenge the move from (4) to (5), the instantiation of (∀F) to (λx)a=x. Unlike The Ugly who held that such expressions as $(\lambda x)J_2a=x$ denoted real, basic, and primitive properties which held "directly" of objects, The Bad holds that such expressions describe instead an epistemic relation which an agent might have to descriptions of reality. That is to say, such expressions do not denote a property and therefore the alleged instance of (LL) is not a legitimate instance. This response is available even to those Bad theorists who adhere to the logic ECNM*IV. For whatever reasons one might hold that "John believes x to be a spy" or "necessarily, x<9" do not describe properties of objects (but perhaps describe an "opaque" relationship between a knowing agent and a description of reality), The Bad theorist can adopt these reasons in the present case.

I conclude therefore that the view of vagueness embodied in The Bad -- that vagueness amounts to inference under incomplete information -- is a coherent doctrine that can be logically investigated. The Bad might be bad, but it seems an unavoidable condition of imperfect knowers in a complex world. As such it deserves careful investigation from the logical point of view.

Extended Evaluation Techniques for Vagueness:

The Good, it will be recalled, locates vagueness in representation. Certain predicates (and possibly other types of representations, although we shall here consider only predicates) lack a kind of preciseness or specificity, so that they can hold of a range of objects (or: in a range of situations) but not have any precise or specific point at which they stop holding and their negations hold. Such a conception presupposes an underlying measure ("the range") and the notion of a predicate applying to a portion of this range. Slightly more precisely, The Good postulates that some predicates are "precise", denoting a specific location on the range; and that other predicates are "vague", denoting a portion of the range. The notion of there being "no sharp cutoff point", which is a distinctive characteristic of vague predicates under The Good's conception, is captured in the following way.

The negation of a vague predicate is also said to denote or describe a portion of the underlying range, but this portion is *not* the complement of that portion described by the unnegated predicate. Most accounts of The Good postulate a *gap*: a portion of the range where neither the predicate nor its negation applies. In seeking a logical account of such a purported phenomenon, these theorists will wish to give an account of the apparent failure of such laws as the excluded middle, and wish to give an account of how inference might work in cases where a statement apparently has no truth value. Some other accounts of The Good might postulate an *overlap*: an area of the range where both the predicate and its negation apply. In seeking a logical account of such a purported phenomenon, these theorists will naturally be concerned with the problem of giving a semantic account which would make it be coherent that a statement asserting that an object falls in the overlap could be true, and of giving an account of inference which does not allow *all* conclusions to follow from such a statement.

In general, one might say, The Good takes the view that the "deficiency of meaning" it finds in vague terms should be represented by some method which "reduces" the meaning of a vague term to the meanings of precise terms in specific contexts. We therefore expect an account of "contexts" to emerge from The Good.

One account of The Good is "supervaluations" or "supertruth" (Fine 1975, Kamp 1975). In this method, it is supposed that, along the underlying range, there is a gap -- an area where neither the vague predicate nor its negation applies. To evaluate a statement involving a vague predicate along this range, one is to postulate a *precisification* -- a context in which there is a specific point along the range where the change from applicability of the positive to the applicability of the negative occurs. With respect to this precisification, evaluation can occur as usual. For example, suppose "is (a) tall (man)" is true of the range of heights greater than 6´3´´, and that "is not (a) tall (man)" is true of that range less than 5´7´´. So there is a gap between 5´7´´ and 6´3´´. The sentences "Al is tall" and "Al is not tall" (said about 5´11´´ Al) might be claimed then to be neither true nor false. But if we choose a precisification -- any choice of a value to eliminate the gap, a precise value which makes "is tall" and "is not tall" be true contradictories -- then the sentence "Al is tall" will have a value and "Al is not tall" will have the opposite value in that precisification. So, a sentence *is true (false) in a precisified* *context* in just the usual, classical manner. A sentence is *true* just in case it is true in every precisified context; a sentence is *false* just in case it is false in every precisified context. A sentence is *indeterminate* (or *unspecified*, or *vague*) just in case it is neither true nor false. (Or equivalently, just in case it is true in some precisified context and false in some (other) precisified context.) Therefore, "Al is tall" is indeterminate, "Al is not tall" is indeterminate, "Either Al is tall or he is not tall" is true, and "Al is tall and he is not tall" is false. In general, according to the supervaluation theory, if we start with a specification of certain "axioms" as being true, then all logical consequences of these axioms will also be true (even ones which involve only vague predicates, such as the previous sentence about Al). Logical consequence amounts to "is a classical logical consequence in every precisification."[12]

Elegant as the supervaluation account is, it nonetheless seems wrong. I, for one, when faced with 5´11´´ Al and asked "Al is tall: true, false, or otherwise?" want to say "false." And when asked "Al is not tall: true, false or otherwise?" again want to say "false." Rather, I want to say that "Al is tall and he isn't tall" is true. To see whether this is a common feeling, a series of informal experiments was run on university undergraduates.[13] There were basically three experiments. In one, subjects were asked these questions about 5´11´´ Al (which height had earlier been empirically determined to be the height at which students are maximally non-committal about whether such a man is tall or not). In another, a series of colour chips ranging from clearly red to clearly pink were asked about; and in the third, a series of colour chips ranging from green to yellow were asked about (because in the second experiment perhaps pink might be viewed as a particular shade of red). Without getting into all the details, different versions of each experiment asked for explicit "True, False, Undecided", or "Clearly true, Clearly false, Neither clearly true nor clearly false", or asked questions in a game where you answered "true" if the subject felt the questioner had spoken truly, "false" if it was felt the questioner had spoken falsely, and remained quiet otherwise, etc.

Abstracting away from many details, let us call a judgment *definite* when a subject says "x is P" is true or false and "x is not P" has the opposite truth value.[14] Of the 533 answers given, 330 (61.8%) were definite in this sense. Many had to do with the colour experiments, especially on the "end chips" which were designed to manifest one colour and not the other. Below, when I give the results I will indicate what happens when these end-chip judgements are removed. An analysis of the results from the red/pink experiment shows that students do *not* perceive pink to be a shade of red. Although there were 4 places where subjects said of a particular chip that it was both red and pink, each of these subjects said of some other chip that it was pink but not red.

Here then are the data. They are lumped in the following way: it will be recalled that for each item being questioned, the subject can answer "True" ("T"), "False" ("F"), or "Otherwise" ("U"). For any particular item, we are interested in the correspondence between how the subject responds to "x is P", "x is not P", and "x is P and it is not P". A complete listing of the results would produce a 27-rowed "judgement table"; to simplify the data, I have grouped the data first into the following types. A *definite* judgement means that the subject called "x is P" either true or false and "x is not P" the opposite truth value. An *overlap* judgement means that the subject called both "x is P" and "x is not P" "T". A *gap* judgement means that the subject either called both "x is P" and "x is not P" "F" or called them both "U". A judgement is on the *true boundry* if one of "x is P" and "x is not P" is called "T" and the other is called "U"; it is on the *false boundry* if one of them is called "F" and the other is called "U"; and it is called a *boundry* judgement if it is either a true boundry judgement or a false boundry judgement.

type of judgement	response to "x is P and x is not P"	number of respondents	percentage of respondents within type	percentage of respondents of that type
gaps	T	81	55.5%	27.34%
	F	33	22.5%	
	U	32	22.0%	
overlaps	T	12	57.5%	3.93%
	F	3	14.0%	
	U	6	28.5%	
definites	T	32	10.0%	61.80%
	F	212	63.5%	
	U	86	26.5%	
true boundry	T	4	14.0%	5.43%
	F	5	17.0%	
	U	20	69.0%	
false boundry	T	4	57.0%	1.31%
	F	2	28.5%	
	U	1	14.0%	
boundry (sum of true and false)	T	8	22.0%	6.75%
	F	7	19.5%	
	U	21	58.0%	

Fig. 1: Results of Experiments, da ta lumped into groups.

As I remarked above, a number of the judgements were concerning the "end chips" on the experiments about colour. The experiments were designed so that these end chips would clearly manifest one but not the other of the colours being studied. (Of course, not every subject perceived it in this manner -- some for example thought an end chip was borderline -- but the majority of the students did answer in the expected manner.) These cases are common enough in the various experiments that they somewhat skew the data presented in Fig. 1. It would be more accurate to delete questions about the end chips when the students did perceive them as manifesting one but not the other of the colours; for these cases are not really judgements about vagueness. After deletion of end chip judgements where the subject saw them as manifesting one but not the other of the colours, we have the following data.

type of judgement	response to "x is P and x is not P"	number of respondents	percentage of respondents within type	percentage of respondents of that type
gaps	T	81	55.5%	
	F	33	22.5%	37.5%
	U	32	22.0%	
overlaps	T	12	57.5%	
	F	3	14.0%	5.5%
	U	6	28.5%	
definites	T	29	15.5%	
	F	93	49.5%	48.0%
	U	65	35.0%	
true boundry	T	4	14.0%	
	F	5	17.0%	7.5%
	U	20	69.0%	
false boundry	T	4	57.0%	
	F	2	28.5%	1.5%
	U	1	14.0%	
boundry	T	8	22.0%	
(sum of true	F	7	19.5%	9.0%
and false)	U	21	58.0%	

Fig. 2: Results of Experiments, da ta lumped into groups.
Definite end-chip judgements removed.

Note that, in Fig. 2, in 37.5% of the judgements the subjects found gaps; of these 55.5% judged that "x is P and x is not P" was true, and the remainder were about equally divided between calling it false and calling it indeterminate. Rather few (9%) thought that any particular case was on a boundry; if they did, the majority (58%) wanted to call the conjunctive judgement uncertain. Fewer still (5.5%) thought there was an overlap; most of them (57.5%) thought the conjunctive judgement was true, presumably using *and* in the usual sense. Almost half (48%) of the particular judgements were deemed to be definite; and of these about half (49.5%) thought the conjunctive judgement was false, again presumably using *and* in the usual sense. The remaining half of these definite judgements were about two-to-one in favour of calling the conjunctive judgement indeterminate.

If we remove all "classical judgements" -- those cases where the subject thought that "x is P" is either True or False and that "x is not P" had the opposite truth value, and from that decided that "x is P and it is not P" was false -- then we are left with the following data which shows subjects' judgements in only those cases where they found some sort of vagueness or indeterminacy:

type of judgement	response to "x is P and x is not P"	number of respondents	percentage of respondents within type	percentage of respondents of that type
gaps	T	81	55.5%	
	F	33	22.5%	45.6%
	U	32	22.0%	
overlaps	T	12	57.5%	
	F	3	14.0%	6.5%
	U	6	28.5%	
definites	T	32	27.1%	
	F	0	0.0%	36.6%
	U	86	72.9%	
true boundry	T	4	14.0%	
	F	5	17.0%	9.1%
	U	20	69.0%	
false boundry	T	4	57.1%	
	F	2	28.6%	2.2%
	U	1	14.3%	
boundry	T	8	22.2%	
(sum of true	F	7	19.5%	
and false)	U	21	58.3%	11.3%

Fig. 3: Results of Experiments, da ta lumped into groups.
"Classicalist Answers" removed.

Note that, when questioned about cases wherein subjects find some sort of vagueness or indeterminacy, only 43 of the 321 judgements (13.4%) fit the predication made by supervaluation theory: that "x is P and it is not P" is false. Nor, it might be added, do they fit the predications made by fuzzy logic theory that "x is P and it is not P" will receive the lesser value attributed to each conjunct, and that the only legitimate values of the conjuncts (being negations of each other) are (i) T, F (ii) F, T and (iii) U, U. Again removing the "classical judgements" on end-chips, we see that only

123 of the 414 judgements (29.7%) follow this pattern. (Of these 123 judgements, 93 were cases where subjects thought the judgements were "definite": one conjunct was either T or F and the other had the opposite truth value, and the subjects thought the conjunction was F. On cases where the subjects *did* find some indeterminacy or vagueness, only 30 of the 307 judgements (9.8%) obeyed the fuzzy logic prediction.) This sort of data suggests that supervaluations simply are not the way people think about sentences which employ vague predicates. "But," it might be rejoined, "so much the worse for the logical abilities of undergraduate students! How can *and* mean what it usually does, "x is P" and "x is not P" both be false, and yet "x is P and it is not P" be true?" Now, I agree that not every response made by naive respondents to logical questions should be taken seriously. For example, were it to turn out (as I think it would) that such informants were to say that the sentence "If 2+2=5 then Ronald Reagan will be elected to a third term" was false or meaningless or indeterminate, I would not take this as evidence that the material conditional should be scrapped. But I find the present case different. As I remarked before, if *I* were asked "Is 5´11´´ Al tall?" I would say "No"; and if I were then asked "Well then, is he not tall?" I would again say "No". Indeed, I would want to say "He's tall and he's not tall." Thus I find myself in the position of my informants,[15] and I wish to give an account whereby this makes logical sense.[16]

The answer to the question of how each conjunct of a vague predication can be F (or U) and yet the conjunction be T must be that the various sentences (each conjunct and then the conjunction) are being evaluated with respect to different contexts. Somehow the information being used to determine truth and falsity shifts between these evaluations. Of course, supervaluations allow change of context, but their only mechanism is that of precisifications, and we have seen that this by itself is insufficient to account for the judgements which are actually made. We need some other notion of a context, and some other notion of a change in context.

So...let us try again to specify what a "context" suitable for use by The Good might be. Again, it is supposed that some predicates are precise and some are vague, and that the latter describe a *range* of the former. Just what range is described is highly context-dependent. If the current discussion is about baseball players, "tall" might be viewed as "more than 6´2´´." In such a context, "not tall" might be viewed as "less than 5´11´´." Note that the negations of vague predicates also describe a portion of the underlying range given by the relevant precise predicates, but *not* (necessarily) the complement of the positive predicate's portion. For each vague predicate and for each context, we will call these *the positive extension* and *the negative extension* of the predicate in that context. We will impose the requirement that, within a given context, the positive and negative extensions are exclusive. If we now change the context so as to talk about basketball players, or about jockeys, the positive and negative extensions of "tall" will be different, and possibly the extensions of other vague predicates will be altered as well.

So there are three questions about contexts to be answered before tackling the analysis of our puzzle cases. (1) What information is contained in a context? (2) How can a context come to be changed? (3) When a context is changed, what gets altered?

Let us start with the first question. In earlier papers (Schubert & Pelletier 1987a, 1987b) a context is described in such a way as to be a method of evaluation for connected discourse. In those papers we were mainly concerned with problems of anaphoric reference and of the specification of event times so as to allow correct temporal evaluation; and so a context explicitly had this information, and a method was given for alteration of contexts along these dimensions. Here I will extend the notion of a context to encompass vague predicates. A context

1. Supplies a classical *interpretation* for the precise terms
    - predicates are assigned an appropriate n-tuple of the domain **D**
    - constants are assigned some member of **D**
2. Fixes the references of "indexical constants" such as
    - I, you, now, here, this, that, yesterday, the actual world

Since a context was to be used in evaluating connected discourse, it also was supposed to supply values for pronouns (or rather, supply values for free variables, which was how the pronouns were translated into "logical form" notation). So, a context

3. Supplies possible referents for free variables. Such possible referents come from the evaluation of preceding sentences. For example,

• John is at the door. He is wearing a hat.

At-Door(John). Behatted(x) -- a context for evaluating the second sentence should supply John for x.

• A man is at the door. He is wearing a hat.

(∃x)At-Door(x). Behatted(x) -- for pronoun reference back to a previously-mentioned, quantified NP, we want the context to supply values for the free variables. Here we want the free x to be a value which made (∃x)At-Door(x) true.

4. Supplies (a set of) times or time intervals at that to evaluate sentences which mention events

• A cat was dropped to the ground. It landed on its feet. -- here the time of evaluation of the second sentence should be shortly after the event described by the first sentence.

5. Supplies an appropriate n-tuple of **D** to each vague predicate, and an appropriate n-tuple to their negations. These needn't be contradictories, but *are* contraries. (So, the positive and negative extensions of a vague predicate are defined relative to a context.)

The idea in the afore-mentioned papers was that as we process the discourse, evaluating sentences one-by-one, we are continually changing or "updating" the context we had from the previous portion of the discourse. In those papers we thought that the updates to the context could be specified by the "logical form" of preceding sentences; thus, we start the discourse with a "neutral" context and the processing of each sentence alters the context (from the one which was in use when processing the immediately preceding sentence) in a way which could be characterized by merely looking at the "logical form representation" of the last sentences. I think this is not entirely suitable as an account of how a context can be altered so as to account for the processing of vague predicates. Still, there are some parts of it which are relevant to the present discussion, and so perhaps a brief look at how we viewed the manner in which the presence of some previous sentence changed the context would be helpful.

I write a context as ˙ Ó. To say that sentence Φ is true in that context I write ˙ΦÓ=1. Variant contexts -- e.g., the result of (re-)interpreting *x* as some d∈**D** -- are written as ˙ $Ó_{x:d}$. Thus for instance, if Φ contained a free *x*, then ˙F$Ó_{x:d}$ would amount to the same thing as evaluating Φ (with *x* replaced by d) in the context ˙ Ó. Remembering that a context specifies w as the value of "the actual world", to say that Φ is true in possible world j, we can write ˙Φ$Ó_{w:j}$=1.

("Reinterpret the actual world as j"). In the previous papers we had in mind that a context is altered as a function of the "logical form representation" of preceding sentences. Therefore we could write ˙ $Ó^{A}$ to indicate the context which results when A is asserted in the context ˙ Ó; or write ˙ $Ó^{A \cdot B \cdots Z}$ to indicate the context resulting from having asserted A, B, ...Z in context ˙ Ó. Intuitively, ˙B$Ó^{A}$ means "the value of B, given that A has been asserted."

I have already indicated the essentials of our view of the "carrying forward" of possible values of bound variables so they can be available later as anaphors of pronouns. Thus if "A man is at the door" is asserted, translated as (∃x)At-Door(x), and evaluated with respect to the context ˙ Ó, then if we encounter the sentence "He is wearing a hat" (translated, with free variable, as Behatted(x)), we want to evaluate this new sentence with respect to the context ˙ $Ó^{(\exists x)\text{At-Door}(x)}$. The rules given in those earlier papers generate from this a context wherein x must satisfy the predicate At-Door, and so when Behatted(x) is evaluated in this context the only values to be considered are those that are relevant. Similar mechanisms were at work to "carry forward" event times for use in the evaluation of following sentences.

As we noted in those earlier papers, not only does a context need to have information generated by evaluating previous sentences, but also sometimes we need first to generate a new context *within* a sentence on the basis of some other clause that occurs in the sentence being evaluated. For example, in a sentence like "When Freud went to Vienna, his theories had not yet been accepted" we need to have evaluated the *when* clause before we can evaluate the main clause, if we are to get the proper anaphor for *his*. We also need a time so that the *yet* (and the tense/aspect) of the main clause can be evaluated. Similarly "donkey sentences" such as "If Pedro owns a donkey, he will ride it to town tomorrow" require us to first evaluate the *if* clause to find a value which satisfies the quantified NP *a donkey*, and then use this new context in evaluating the main clause. And in generic-*when* sentences such as "When cats are dropped to the ground, they usually land on their feet", we need to evaluate the *when* clause to find out first, a suitable (set of?) reference(s) for *they*; second, a suitable time at which to evaluate *land on their feet* (a time shortly after a time when the *when* clause is true); and third, a suitable set of situations (namely cat-droppings-to-the-ground) with which to evaluate the *usually* (after all, the sentence does not mean that usually cats are landing on their feet!).

In addition, not only do we need to construct these new contexts from "left-to-right" as we are processing a sentence, but sometimes, within a sentence, we need to first evaluate a following clause before we are in a position to evaluate the previous one. In "When he went to Vienna, Freud's theories had not yet been accepted" we see that before we will be in a position to completely evaluate the *when* clause, we must find Freud in the main clause. Furthermore, it should be noted that before we can completely evaluate the main clause (especially the *yet*), we must get the information about the event time referred to in the *when* clause. Far from being an isolated occurrence, this sort of "zig-zag", "bootstrapping" construction of the appropriate context seems quite common. Consider such sentences as "Usually, when it is dropped to the ground, a cat will land on its feet" and the like. To find a suitable referent for *it*, we first must partially evaluate the main clause; but to find the time appropriate to evaluate *will land on its feet* we first must find the event time appropriate to this from the *when* clause; and finally, in order to properly evaluate *usually* we need to have evaluated the *when* clause to find out the class of events to which *usually* is supposed to apply. Similar examples can be constructed at will, once one knows what to look for.

As I mentioned, however, this method of "looking at the logical form representations" of preceding sentences does not seem appropriate as an account of how contexts are altered, in the area of context-changing relevant to vague predicates. Consider sentences where a change of context is explicitly made: "John is tall for a baseball player, but not tall for a basketball player" does not seem to require that John be either a baseball or a basketball player in order to be true. Clearly the first clause somehow directs us to consider "the context of baseball players" while the second directs us to change this context to "the context of basketball players"; but apart from the presence of the words *baseball* and *basketball*, it is not obvious what clues are being employed to direct us to do so. And in any case, the mere presence of those words -- even in a similar linguistic setting -- does not necessarily make for a change of context. For example, had the sentence been "John is a baseball player but not a basketball player" there would be no change of context. The issue of just what causes a new context to come into existence (with respect to vague predicates such as *tall*) is a difficult issue about which I have little to say. Besides explicit changes of context ("Now let's talk about basketball players"), a change in context can be initiated by a noun being modified by a vague adjective. For example, it has been viewed as something of a puzzle how "Mickey is a tall mouse but not a tall mammal" could be true when "All mice are mammals" is true (cf. Kamp 1975, among others). The answer, from the current viewpoint, is that when a vague adjective (*tall*) modifies a noun (*mouse*, *mammal*), the noun calls into play its own new context. And in this new context the positive and negative extensions of the vague predicate are different from the ones in use in previous contexts.

When a new context is called into play (in whatever way -- e.g., by explicitly stating what is now to be considered, or by background knowledge of the current topic of discussion, or by modification of a noun by a vague predicate, etc.), the positive and negative extensions of *all* vague predicates (may) need to be altered from what they were before. For example, suppose the current context is one in which we are discussing North American men. This context determines *inter alia* the positive and negative extensions of *tall*, of *very good at basketball*, of *educated*, of *has dark skin colour*, etc. Now suppose we change the context (in whatever way) so that it is clear that we are talking about professional basketball players. In that new context, we will want to change the positive and negative extensions of each of the above vague predicates. Generally speaking, it is possible for *every* vague predicate to "change its meaning" with a change of context, and given a change of context it does not seem possible to predict which of the vague predicates will remain the same as they were in any preceding context.

Note that almost anything can be used to establish a new context. Again, suppose we are talking about North American men and perhaps evaluating their heights -- John is tall, Bill is not tall, etc. We can change context by saying "Let's consider <u>tall</u> North American men." This will now change the positive and negative extensions of *tall*, so that *not tall* will now include a portion of the range which used to be described by *tall* (in the old context). Speaking about the discourse, we might say that what was tall previously is no longer tall. This is no contradiction -- even though no person has changed his height -- but merely a "Cambridge change", involving an alteration in the standards of judging tallness.

I now wish to use all this machinery to account for our puzzle about "Al is tall" is False (or Undeterminate), "Al is not tall" is False (or Undeterminate), and yet "Al is tall and he is not tall" is True. Suppose ˙ Ó is a context in which tall men are more than 6´3´´ and not-tall men are less than 5´7´´, and Al is in the extension of the precise predicate "is a 5´11´´ tall man". Then ˙Al is a tall manÓ will be False (or Undeterminate, depending on the details of how we assign truth to such sentences) and ˙Al is a not-tall manÓ will also be False (or Undeterminate) in that context. But what of our conjunctive sentence? Recall from before that complex sentences often call for a kind of "zigzag", "bootstrapping" construction of new contexts with which to evaluate each of the clauses. With respect to conjunctions, the following rule of evaluation seems appropriate:

˙A ∧ BÓ = 1 iff ˙AÓ^{B} = 1 and ˙BÓ^{A} = 1

That is, each conjunct must be evaluated in a context which is a modification of the current context in the way described (however that is done) by the other clause. If each of the conjuncts are true in these different contexts, then the conjunction is true in the current context. So, with respect to "Al is tall and he is not tall", this is true (in the current context) if and only if

(1) "Al is tall" is true in a context in which it is given that Al is not tall

(2) "Al is not tall" is true in a context in which it is given that Al is tall

Now, given what the positive and negative extensions of *tall* were in the current context, the newly-generated contexts in which it is given that Al is not tall are those old contexts in which Al does not fall under the predicate "is tall"; that is, those contexts in which we are examining men under 6´3´´. What is the positive extension of "is tall" in such a context? Perhaps it is those men taller than 5´9´´ (and of course less than 6´3´´). In such a context, Al is tall. The newly-generated contexts in which Al is tall are those old contexts in which Al does not fall under the predicate "is not tall"; that is, those contexts in which we are considering men over 5´7´´. What is the extension of "is a not-tall man" in such a context? Perhaps it is those men less than 6´1´´ (and of course greater than 5´7´´). In such a context, Al is not tall. Since Al is tall, given a (natural) context in which he is not tall, and since Al is not tall, given a (natural) context in which he is tall, the sentence "Al is tall and he is not tall" is true.

Note that if the predicates involved in the two conjuncts are all precise, then the contexts do not get altered in this manner (although the contexts might still get altered by some of the other mechanisms having to do with pronouns and event times). The reason for this is that we do not get "predicate extension shift" with precise predicates from context to context. Thus, if all predicates in some sentence A are precise, the sentence (A∧~A) will be always False.

Supervaluation theory had an account of what vague predicates are true of: "*x* is a tall man" is true of those objects for which it is true in *every* precisification. Besides the different notion of what a context is (a context in supervaluation theory was merely a precisification, the current theory says that a context specifies a positive and negative extension of all vague predicates -- but no precisification), the theory being put forward here also differs in what it thinks vague predicates are true of. In the present theory, the predicate "*x* is a tall man" is true of those objects for which *some* (natural) context can be found in which the predicate would be true of them. Surely this is an account more in keeping with our intuitions -- in supervaluation theory the predicate "*x* is a tall man" would not hold of 6´5´´ Julius simply because there is a precisification involving basketball centers in which Julius was not tall. Here, we count Julius as tall because there is a (natural) context in which he is tall. But, it might be asked, do we also count him as not-tall? The answer is yes, on the assumption that the context of basketball centers is a natural context. Can a 7´5´´ man ever be counted as not-tall? Probably not, since it seems unlikely that there is a <u>natural</u> context in which such a person would be not-tall. (There are "artificial contexts", to be sure, such as "let's consider the group of men whose average height is 8´0´´"; but given ordinary usage as our guide, there are no natural contexts in which a 7´5´´ man would be called not-tall.)

I think the present account of The Good's notion of "semantic vagueness", involving, as it does, the notion of changing contexts, is quite interesting in its own right. Furthermore, it holds the tantalizing promise of fitting well with the independently motivated notion of context-shift used for an account of anaphora, event-times, and the evaluation of generic sentences as discussed in Schubert

& Pelletier 1987a, 1987b. Much work remains to be done, however, especially in giving a formal account of what causes a context to change. I hope that the present account will be suggestive enough to encourage other researchers to try to carry it further. In any case, I hope to have shown that The Good stands a good chance of really being good.

ACKNOWLEDGEMENTS

This paper began life as comments on Heintz (1981). Various versions (of various of the parts) have been read at: Univ. Nebraska (1981), Western Canadian Philosophical Association (Calgary, 1981), Conference on the Foundations of Logic (Waterloo, 1982), Society for Exact Philosophy (U. Georgia, 1984). I wish to thank the audiences at these occasions for various and sundry remarks and discussion. A paper based on the above section "Modal Logics for Vagueness" was published in *Logique et Analyse* (Pelletier 1984a). As can be seen from the discussion in the section on The Good, I am much in the debt of Len Schubert. I have discussed the issues brought up by the topic of vagueness with a number of people besides the above, and would like to take this opportunity to thank them: Ali Ahktar Kamzi, Kit Fine, Ray Jennings, David Lewis, Bernard Linsky, Mohan Matthen, Richard Routley Sylvan, Peter Schotch, Alisdair Urquhart, and Johan van Bentham. Additionally, I would like to thank the members of the University of Alberta Logical Grammar Study Group for discussions on this paper and the topic generally: Matthew Dryer (Linguistics), Bill Blackburn, Ron Clark, Bernard Linsky (Philosophy), and Len Schubert (Computing Science). I also gratefully acknowledge the aid of NSERC grant A5525.

FOOTNOTES

1. But not for Russell (1923) who thinks that *every* term of *every* representational system is vague. His only exception is *Principia Mathematica* which he thinks is completely precise.

2. Thanks to Richard Routley Sylvan for pointing this out to me. One should consult the series of papers in *Logique et Analyse* in the late 1960's by Montgomery & Routley on this.

3. Just which other truth value is a matter of detail of the inner workings of the system. There might for example be many other truth values. All I am insisting on here is that, due to the "realistic" nature of The Ugly's conception of vagueness, the "neither having nor lacking *F*" is a definite, real property of objects. And therefore there must be some truth value to describe this situation.

4. Of the Dept. Computing Science, Univ. Alberta. See "Acknowledgements" section.

5. Heintz (1981), Kearns (1974), van Inwagen (1986).

6. Following Rosser & Turquette (1952).

7. There are other reasons to eschew fuzzy logics, such as (i) even the propositional fragment has no complete and sound theory of argumentation, (ii) the monadic (and full) predicate logic has no normal forms and hence there can be no automated theory of deduction along the lines of "resolution", (iii) the full predicate logic is not recursively axiomatizable. See Morgan & Pelletier (1977).

8. This discussion of logics for The Bad follows Pelletier (1984a).

9. Actually, it takes a bit of an argument to show that (U) holds for this reason. See Pelletier (1984a), footnote 6.

10. In Pelletier (1984b) it is proved that, appearances to the contrary notwithstanding, the logic ECNM*I "is just logic T in disguise", and that adding (B), or (4), or (V) just yields the logics B, S_4, and S_5.

11. In the logic under consideration (a) and (b) are equivalent to

$$(\alpha')\ {}^{\circ}(\exists x)Fx \rightarrow \sim(\forall x)\Diamond Fx$$
$$(\beta')\ \Diamond(\exists x)Fx \rightarrow \sim(\forall x)\,{}^{\circ}Fx$$

respectively. From these we can more clearly see the force of the objections. With respect to (α´), it is certainly possible for it to be definite whether something is F but also for everything that there happens to be, be indefinite whether it is F (and still think that there's definitely *some* F). With respect to (β´), one can be unclear whether something is F, but of everything there happens to be nonetheless be quite definite whether it is F.

12. As an interesting feature of supervaluation theory, note that the "inductive premise" of sorites arguments is false, since it is false in every precisification. "For all n, if a man with n hairs is bald, then a man with n+1 hairs is bald" fails in each precisification.

13. Devised in large part by Len Schubert (see fn. 4) and partly by me. The tests were administered to beginning students in various Computing Science courses. Whether English was a native language was also noted, but did not discriminate answers; neither did whether the subject had taken any logic courses.

14. Different of the experiments will have different values for "x is not P". For instance, in the "tall" experiment, it is "Al is not tall"; in the "pink/red" experiment, "x is P" might be "This chip is

red" and "x is not P" might be "This chip is pink"; and in the versions where it is asked "Is it definitely true that Al is tall", the "x is P" answer is "yes" and the "x is not P" answer is "no".

15. And, I might add, of various colleagues in Philosophy, Linguistics, Psychology, and Computing Science.

16. There are further problems with using naive subjects in these kinds of experiments. For example, one might be curious as to how they would answer "x is P or it is not P" for these vague predications. Such questions were in fact asked at the beginning by us in our experiments, but unfortunately the answers are difficult to interpret. Part of the difficulty is that naive subjects in general are not willing to affirm a disjunctive statement when they *know* which one of the disjuncts is true -- they will about half the time say it is False, about a quarter of the time say it is True, and about a quarter of the time say it is Undecided. This information comes from another "experiment" carried out by Schubert. Subjects in two groups are told: "There are two marbles in the bag, one pink and one red." One group is then told: "The pink one is removed. Is the remaining marble either pink or red?" Here, more than 50% said "no", and about 25% each said "yes" or said "undecided". The other group was told: "One of the marbles is removed. Is the remaining marble either pink or red?" 100% of subjects said "yes" to this, showing that part of their judgements about *or* statements has to do with knowledge of which disjunct is true. So we should be suspicious of answers to such *or* questions when used to look into vague statements. I do not think anything like this is happening in the *and* questions.

BIBLIOGRAPHY

Alston, W. (1964) *Introduction to Philosophy of Language* (Englewood Cliffs: Prentice-Hall).

Alston, W. (1968) "Vagueness" in P. Edwards *Encyclopedia of Philosophy* (Macmillan: N.Y.)

Austin, J. (1962) *Sense and Sensibilia* (Oxford: Oxford UP).

Chellas, B. (1980) *Modal Logic* (Cambridge: Cambridge UP).

Chierchia, G. (1982) "Nominalization and Montague Grammar: a Semantics without Types for Natural Languages" *Linguistics and Philosophy* **5**, 303-354.

Dummett, M. (1975) "Wang's Paradox" *Synthèse* **30**, 301-324.

Evans, G. (1978) "Can There Be Vague Objects?" *Analysis* **38**, 208.

Fine, K. (1975) "Vagueness, Truth, and Logic" *Synthèse* **30**, 265-300.

Heintz, J. (1981) "Might There be Vague Objects?" Paper read to Canadian Philosophical Assoc. meetings, Halifax.

Kamp, J. (1975) "Two Theories About Adjectives" in E. Keenan *Formal Semantics of Natural Language* (Cambridge, Cambridge UP) 123-155.

Kearns, J. (1974) "Vagueness and Failing Sentences" *Logique et Analyse* **17**, 301-316.

Margalit, A. (1976) "Vagueness in Vogue" *Synthèse* **33**, 211-221.

Morgan, C. & F.J. Pelletier (1977) "Some Notes on Fuzzy Logic" *Linguistics and Philosophy* **1**, 87-121.

Pelletier, F.J. (1984a) "The Not-So-Strange Modal Logic of Indeterminacy" *Logique et Analyse* **27**, 415-422.

Pelletier, F.J. (1984b) "Six Problems in 'Translational Equivalence'" *Logique et Analyse* **27**, 423-434.

Quine, W. (1960) *Word and Object* (Cambridge, Mass.: MIT Press).

Rosser, B. & A. Turquette (1952) *Many-Valued Logics* (Amsterdam: North Holland).

Russell, B. (1923) "Vagueness" *Australasian Journal of Psychology and Philosophy* **1**, 84-92.

Schubert, L. & F.J. Pelletier (1987a) "Problems in the Representation of the Logical Form of Generics, Plurals, and Mass Nouns" in E. LePore *New Directions in Semantics* (London: Academic Press) 387-453.

Schubert, L. & F.J. Pelletier (1987b) "Generically Speaking: With Remarks on the Interpretation of Pronouns and Tenses" forthcoming in G. Chierchia, B. Partee, R. Turner *Property Theory, Type Theory, and Semantics* (Dordrecht: Reidel).

Segerberg, K. (1971) *An Essay in Classical Modal Logic* (Uppsala Univ.: Filosofia Series).

van Inwagen, P. (1986) "How to Reason About Vague Objects" Paper read to Central Meetings of the American Philosophical Association, St. Louis.

Wheeler, S. (1975) "Reference and Vagueness" *Synthèse* **30**, 367-379.