Jim McCloskey,
La Jolla,
April 16th. 1981

Two theories about adjectives[1]

J. A. W. KAMP

1.

I will discuss two theories about adjectives. The first theory dates from the late 1960s. It is stated in Montague (1970) and Parsons (1968). According to this theory the meaning of an adjective is a function which maps the meanings of noun phrases onto other such meanings; e.g. the meaning of *clever* is a function which maps the meaning of *man* into that of *clever man*, that of *poodle* onto that of *clever poodle*, etc. Predicative uses of adjectives are explained as elliptic attributive uses. Thus *This dog is clever* is analysed as *This dog is a clever dog* – or as *This dog is a clever animal*, or perhaps as *This dog is a clever being*. Which noun phrase ought to be supplied in this reduction of predicative to attributive use is in general not completely determined by the sentence itself, and to the extent that it is not, the sentence must be regarded as ambiguous.

The main virtue of this doctrine is that it enables us to treat, within a precise semantical theory for a natural language – as e.g. that of Montague – adjectives in such a way that certain sentences which are, or might well be, false are not branded by the semantics as logically true. Examples of such sentences are:

(1) Every alleged thief is a thief
(2) Every small elephant is small
(3) If every flea is an animal, then every big flea is a big animal

Each of these sentences would come out logically true in Montague's model theory if it were to treat adjectives as ordinary predicates, so that the logical form of (1), for example, would be $(\forall x)(A(x) \wedge T(x) \rightarrow T(x))$.

[1] Since I presented the outline of this paper at the Cambridge conference I – and, I hope, this paper – have profited from discussions with and comments by Michael Bennett, Richard Grandy, Hidé Ishiguro, David Lewis, Richmond Thomason and, in particular, George Lakoff. I was equally fortunate to hear Sally Genet's paper on comparatives at the summer meeting of the Linguistic Society of America in Ann Arbor, which proposed an approach similar to that taken here. Only after the present paper had already been given its final form did I become acquainted with Kit Fine's article 'Vagueness, truth and logic' which expresses on the topic of vagueness, which is the central theme of the second part of my paper, views very similar to those which can be found here. I know that I would have been able to offer a better contribution to this volume if I had known about Fine's work earlier.

Moreover, the theory allows us to express in very simple mathematical terms some important semantical features which some, though not all, adjectives possess. In order to give precise formulations of such features, it is necessary to make some assumptions about the comprehensive semantical theory in which this particular doctrine about adjectives is to be embedded. These assumptions can all be found in Montague (1970). I regard them as basically sound, but would like to point out to those who have strong qualms about possible world semantics that the distinctions drawn by the definitions below do not depend on these assumptions as such.

The assumptions are the following:

(a) Each possible interpretation (for the language in question) is based upon (i) a certain non-empty set W of possible worlds (or possible situations, or possible contexts) and (ii) a set U of individuals.
(b) A property relative to such an interpretation is a function which assigns to each $w \in W$ a subset of U (intuitively the collection of those individuals which satisfy the property in that particular world (or context) w).
(c) The meaning of a noun phrase in such an interpretation is always a property.

Thus the meanings of adjectives in an interpretation of this kind will be functions from properties to properties.

We may call an adjective *predicative in* a given interpretation if its meaning F in that interpretation satisfies the following condition:

(4) there is a property Q such that for each property P and each $w \in W$, $F(P)(w) = P(w) \cap Q(w)$.

Once we have singled out a given class $\mathscr{K}$ of admissible interpretations, we can also introduce the notion of being *predicative* simpliciter: an adjective is *predicative* (with respect to the given class $\mathscr{K}$) if and only if it is predicative in each interpretation (belonging to $\mathscr{K}$).

Predicative adjectives behave essentially as if they were independent predicates. If for example *four-legged* is treated as predicative then any sentence *If every* N_1 *is an* N_2 *then every four-legged* N_1 *is a four-legged* N_2, where N_1 and N_2 are arbitrary noun phrases, will be true in each admissible interpretation in all the worlds of that interpretation.

Predicative adjectives are, roughly speaking, those whose extensions are not affected by the nouns with which they are combined. Typical examples are technical and scientific adjectives, such as *endocrine*, *differentiable*, *superconductive*, etc.

We may call an adjective *privative in* a given interpretation if its meaning F in that interpretation satisfies the condition

(5) for each property P and each $w \in W$ $F(P)(w) \cap P(w) = \phi$

Again, an adjective will be called *privative* if (5) holds on all admissible interpretations.

A privative adjective A is one which, when combined with a noun phrase N produces a complex noun phrase AN that is satisfied only by things which do not satisfy N. If A is a privative adjective then each sentence *No* AN *is an* N will be a logical truth. Adjectives that behave in this way in most contexts are e.g. *false* and *fake*. I doubt that there is any English adjective which is privative (in the precise sense here defined) in all of its possible uses.

An adjective is *affirmative in* a given interpretation if its meaning satisfies

(6) for each P and w,
$F(P)(w) \subseteq P(w)$

It is *affirmative* if (6) holds in all admissible interpretations.

Clearly all predicative adjectives are affirmative. But there are many more. In fact the vast majority of adjectives are affirmative. Typical examples of affirmative adjectives which are not predicative are *big*, *round*, *pink*, *bright*, *sharp*, *sweet*, *heavy*, *clever*.

Finally, an adjective is *extensional in* a given interpretation if

(7) there is a function F' from sets of individuals to sets of individuals such that for every P and w $(F(P))(w) = F'(P(w))$

and *extensional* if (7) holds in all admissible interpretations.

Thus a predicative adjective is in essence an operation on extensions of properties: if two properties have the same extension in w then the properties obtained by applying the adjective to them also have the same extension in w.

Clearly all predicative adjectives are extensional. Non-extensional adjectives are for example *affectionate* and *skilful*. Even if (in a given world) all and only cobblers are darts players, it may well be that not all and only the skilful cobblers are skilful darts players;[1] and even if all men were fathers the set of affectionate fathers would not necessarily coincide with the set of affectionate men.[2]

[1] This example was given at the conference by Professor Lewis.
[2] This example was given to me by Dr Hidé Ishiguro of University College, London.

It is an interesting question whether there are any adjectives which are extensional but not predicative. It has been suggested[1] that in particular such adjectives as *small*, *tall*, *heavy*, and *hot* belong to this category. Indeed these adjectives are evidently not predicative, whereas their extensionality follows from a certain proposal according to which they derive from their comparatives in the following way. Let A be an adjective of this kind, and let $\mathscr{R}$ be the binary relation represented by the phrase *is more A than*. The function $\mathscr{A}$ from properties to properties which is associated with A is then characterized by

(8) for any property P and world w $\mathscr{A}(P)(w) = \{u \in P(w)$: for most $u' \in P(w) \langle u,u' \rangle \in \mathscr{R}(w)\}$

It will soon be evident why I have not much sympathy for analyses of positives in terms of comparatives generally. At this point, however, I only want to express some reservations which concern (8) in particular. That (8) cannot be right is brought out by the fact that it logically excludes the possibility that, for any property P, most Ps are small Ps and only a few are large Ps. Thus what we usually call a small car in England would according to (8) not be a small car; for we call most English cars small. (One might perhaps reply that this only shows that by *small car* we mean *car of a small model*. But that does not quite do. After all, it is the *individual* cars we call small.) In this case the conflict between usage and the consequences of (8) arises from the fact that cars are naturally divided into categories; and it is to these, if anything, that (8) applies.

There is yet another reason why (8) might fail for *small*. We might have a clear concept of what is the normal size of objects satisfying a certain property, even if objects of that size which have the property do not or only rarely occur. It is conceivable to me that we would then call almost all members of a species S small members of S if there was strong biological evidence that only accidental and abnormal circumstances C prevent the majority from growing into a height which most members of the species would reach under conditions we would regard as normal. Yet we might still be unwilling to call them small members of S-under-the-circumstances-C; for as objects falling under that second description they should be expected to have the size they do have. If, moreover, S and S-under-C had precisely the same extensions the case would tend to show that *small* is not purely extensional. But that it is so difficult to come up with a concrete and convincing example of this sort is perhaps an indication that for all practical purposes *small*, and similar adjectives, are indeed extensional.

[1] Cf. Bartsch and Vennemann (1972), part II.

2.

This theory of adjectives is of course not new. The observation that *John is a good violinist* cannot be analysed as *John is good and a violinist* is probably too old to be traced back with precision to its origin.[1] What is perhaps new in the doctrine as I have stated it here is the emphasis on the fact that what has for a long time been observed to be a feature of certain adjectives is a common feature of them all. *All* of them are functions from noun phrases to noun phrases. Some adjectives (expressed in (4)), however, possess a certain invariance property that makes them behave as predicates, which when combined with a noun phrase give a complex equivalent in meaning to the conjunction of the predicate represented by the adjective and that represented by the noun phrase.

Even if this theory does accomplish the rather simple-minded tasks for which it was designed, one may feel dissatisfied with it for a variety of reasons. Here I will mention only one (although I believe there are other grounds for dissatisfaction as well): The theory is incapable of providing an adequate treatment for the comparative and superlative. For reasons of convenience I will concentrate on the comparative and leave the superlative aside; but the theory which will emerge from our considerations will handle the superlative as well.

From a naïve point of view the comparative is an operation which forms out of an adjective a binary predicate. I believe that this naïve point of view is correct: that when we learn a language such as English we learn the meanings of individual adjectives and, moreover, the semantic function which this comparative-forming operation performs *in general*, so that we have no difficulty in understanding, on first hearing, the meaning of the comparative of an adjective of which we had thus far only encountered the positive. If this is so then the meaning of an adjective must be such that the comparative can be understood as a semantic transformation of that meaning into the right binary relation.

It is quite obvious that if adjectives were ordinary predicates no such transformation could exist. How could we possibly define the relation *x is bigger than y* in terms of nothing more than the extension of the alleged predicate *big*?

[1] A clear exposition of a view about the adjective *good* which is essentially what is here proposed for adjectives in general can be found in Geach (1956). Notice, however, that not only does *good*, as Geach makes clear, fail to be a predicate; it is not even extensional (cf. *skilful* above).

Could functions from properties to properties serve as the basis for such a transformation? This is a more problematic question. One might for example characterize the transformation as follows:

For any adjective A with meaning $\mathscr{A}$ in a given interpretation we have for any $u_1, u_2 \in U$ and $w \in W$

(9) u_1 is more A than u_2 in w iff

(a) for every property P such that u_1 and u_2 both belong to $P(w)$ if u_2 belongs to $\mathscr{A}(P)(w)$ then so does u_1

(b) there is a property P such that $u_1, u_2 \in P(w)$, $u_1 \in \mathscr{A}(P)(w)$, and $u_2 \notin \mathscr{A}(P)(w)$.

This definition is in the right direction. But I doubt that it will do. In particular, I doubt whether (b) is a necessary condition. Take *tall*, for example. According to (9) u_1 is taller than u_2 only if there is a property P that applies to both of them and such that u_1 is a tall P while u_2 is not. But suppose that u_1 is taller than u_2 by a tiny bit. Can we then find a P which satisfies this condition? The question is not easy to answer. Let us suppose for the sake of argument that *tall* can be correctly defined by (8) as *taller than most*. Then the question depends on whether we can find a property P such that u_1 is taller than most *P*s while u_2 is not. But this can only be the case if there are enough things in the extension of P which have heights intermediate between those of u_1 and u_2. And perhaps there are no such things at all. Now, in our discussion of extensional adjectives we found that (8) is probably not adequate in any case. So there may be after all a property P which satisfies the condition. In this manner we might succeed in saving (9) by imploring the assistance of some bizarre property whenever we need one. But I find this solution ad hoc and unsatisfactory. What underlies the possibility of making comparative claims is that adjectives can apply to things in various degrees. It is my strong conviction that when we learn the meaning of an adjective we learn, as part of it, to distinguish with greater or lesser precision to what degree, or extent, the adjective applies to the various entities to which it applies at all. Once we have learned this we are able to understand the comparative of the adjective without additional explanation, provided we understand the function of the comparative in general.

In order to give my view on the primacy of positive over comparative an adequate foundation I will develop a semantical framework in which the idea of a predicate being true of an entity to a certain degree can be made coherent and precise. This specific problem is closely related to such general features of natural languages as vagueness and contextual dis-

ambiguation; indeed I hope that the theory which I will outline will provide an adequate framework for the treatment of these problems as well.

Before stating what at this point I believe to be the most promising framework for our purpose, I first want to make some remarks on a theory of formal logic which has been often proposed just for the solution of the problems with which I want to deal, viz. multi-valued, or many-valued logic.

Most systems of multi-valued logic available in the literature are systems of propositional calculus. In view of our purpose, our interest should lie with multi-valued predicate logic. But for reasons of exposition I will consider the simpler propositional logics.

Multi-valued logics differ from ordinary two-valued logic in the first place by their model theories. Indeed many such systems are syntactically indistinguishable from standard formulations of ordinary propositional calculus; and I will consider for the time being only such systems of multi-valued logic, all of which have the same syntax, based upon an infinite set $q_1, q_2, q_3, \ldots$ of propositional variables. Starting with these we can recursively construct complex formulae, i.e. $\neg(\phi)$, $\wedge(\phi,\psi)$, $\vee(\phi,\psi)$, $\rightarrow(\phi,\psi)$, $\leftrightarrow(\phi,\psi)$ from already constructed formulae ϕ and ψ. (I will write $(\phi \wedge \psi)$ for $\wedge(\phi,\psi)$, etc.) Let us call this language of propositional logic L_0.

A multi-valued semantics for L_0 will provide for this language a model theory based upon some set TV (of 'truth values') the cardinality of which is $\geqslant 2$. Two-valued propositional calculus emerges as a case where TV contains exactly two elements. A ***model*** for L_0 according to such a model theory based upon TV is a function which assigns to each variable q_i an element of TV. Such a function uniquely determines the (truth) values of the complex formulae of L in virtue of another component which specifies for each t_i in TV what the value is of $\neg\phi$ given that t_i is the value of ϕ; for each t_i, t_j in TV what the value is of $(\phi \wedge \psi)$ given that ϕ has t_i and ψ has t_j; and similarly for the other connectives.

The definition of logical truth requires a third component of the theory which singles out a proper non-empty subset TV_t of TV, the set of 'designated' truth values. A formula will be regarded as *logically true* if in each model it has a value belonging to TV_t. *Logical consequence* can be defined in an analogous manner.

Thus we come to the following formal definition:

A *multi-valued model theory* (in short m.m.t.) *for* L_0 is a triple $\langle TV, TV_t, F\rangle$ where (i) TV is a set of cardinality $\geqslant 2$;

(ii) TV_t is a proper, non-empty subset of TV;

(iii) F is a function which maps each n-place connective of L_o onto a n-place function from TV into itself.

A *model for* L_o *relative to the m.m.t.* $\langle TV, TV_t, F\rangle$ is a function from $\{q_1,q_2,q_3,\ldots\}$ into TV.

Let $\mathcal{M} = \langle TV, TV_t, F\rangle$. The truth value of a formula ϕ of L_o in a model M relative to $\mathcal{M}$, $[\phi]_{,\mathcal{M}}$ is defined by the clauses:

(i) $[q_i]_M = M(q_i)$;

(ii) $C(\phi_1,\ldots,\phi_n)\,[^{\mathcal{M}}_{M}] = F(C)\,([\phi_1]^{\mathcal{M}}_{M},\ldots,[\phi_n]^{\mathcal{M}}_{M})$ for any n-place connective C of L_o.

ϕ is *logically true in* $\mathcal{M}$ iff $[\phi]^{\mathcal{M}}_{M} \in TV_t$ for all models M relative to $\mathcal{M}$.

Clearly the classical semantics for propositional calculus is the model theory $\langle\{0,1\},\{1\},F_c\rangle$ where $F_c(\neg),F_c(\wedge),\ldots\ldots, F_c(\leftrightarrow)$ are the functions defined by the usual two-valued truth tables for these connectives.

It is natural to require of a model theory for L_o based upon a truth value set of cardinality $\geqslant 2$ that it has the feature:

(10) there are two particular elements of TV – let us call them 0 and 1 – such that
(i) $1 \in TV_t$;
(ii) $0 \notin TV_t$;
and (iii) for each connective C, the restriction of $F(C)$ to $\{0,1\}$ is the usual two-valued truth table for C (i.e. there are among the truth values two, which we might think of as 'absolute falsehood' and 'absolute truth', with respect to which the connectives behave in the ordinary classical manner).

It appears that those model theories for L_o which have been seriously proposed in the literature do indeed satisfy (10).[1] The vast majority of these theories assume, moreover, a linear ordering of the members of TV with respect to which 0 is the smallest and 1 the largest element. The formal properties of the theory, as well as its philosophical relevance, then depend on the characterization of the functions $F(C)$. Clearly there are, whenever the cardinality of TV is greater than 2, various such functions which do not violate (i). The question is which of these 'correctly capture' the function of the connectives *not, and, or,*...given a particular interpretation of what the truth values in TV really represent. Let us consider the case where $TV = \left\{0, \frac{1}{n-1}, \frac{2}{n-1}, \ldots, \frac{n-2}{n-1}, 1\right\}$ and where these numbers represent 'degrees of truth' – the higher the number the higher the degree.

[1] For references see Rescher (1969).

What function F would adequately reflect our intuitions about the semantic behaviour of these connectives in this case? I think there are no such functions. The reason is that the connectives *not, and, or,*...are not functions of degrees of truth. This becomes evident almost immediately when one reflects upon the definition of $F(\neg)$. The natural suggestion here is that $F(\neg)\left(\frac{k}{n-1}\right) = 1 - \frac{k}{n-1}$, i.e. that the negation of a proposition is true exactly to the degree that that proposition itself fails to be true. And this indeed is a definition of $F(\neg)$ which is commonly accepted.

Now let us assume that n is odd so that one of the truth values is $\frac{1}{2}$. What value would $F(\wedge)$ assign to the pair of arguments $(\frac{1}{2},\frac{1}{2})$? It is plausible that the value should be $\leqslant\frac{1}{2}$. For how could a conjunction be true to a higher degree than one of its conjuncts? But which value $\leqslant\frac{1}{2}$? $\frac{1}{2}$ seems out because if $[\phi]^{\mathcal{M}}_{M} = \frac{1}{2}$ then, if we accept our definition of $F(\neg)$, $[\neg\phi]^{\mathcal{M}}_{M} = \frac{1}{2}$. So we would have $[\phi\wedge\neg\phi]^{\mathcal{M}}_{M} = \frac{1}{2}$, which seems absurd. For how could a logical contradiction be true to *any* degree? However, if we stipulate that $(F(\wedge))(\frac{1}{2},\frac{1}{2}) = 0$, we are stuck with the even less desirable consequence that if $[\phi]^{\mathcal{M}}_{M} = \frac{1}{2}$, $[\phi\wedge\phi]^{\mathcal{M}}_{M} = 0$. And if we choose any number between 0 and $\frac{1}{2}$, we get the wrong values for both $\phi\wedge\neg\phi$ and $\phi\wedge\phi$.[1]

This argument indicates why we cannot represent the connectives accurately within the narrow framework of multi-valued semantics based upon linearly ordered truth-value sets. The reason can be expressed thus: the truth value of a complex formula – say $\phi\wedge\psi$ – should depend not just on the truth values of the components – i.e. ϕ and ψ – but also on certain aspects of these formulae which contribute to their truth values but cannot be unambiguously recaptured from them.

The possibility of treating the connectives truth-functionally in two-valued model theory (while not in model theories based on larger linearly ordered sets of truth values) is a reflection of the fact that two-element sets are the only linearly ordered truth-value sets which can be regarded as Boolean algebras in the following sense:

If we define the Boolean operations $\cap$, $\cup$, $-$ in the usual manner in terms of the ordering relation $\leqslant$ of a linearly ordered set TV (i.e. if we put $C\cap C' = df$ the largest C'' under $\leqslant$ such that $C''\leqslant C$ and $C''\leqslant C'$; $C\cup C' = df$ the smallest C'' such that $C\leqslant C''$ and $C'\leqslant C''$; and $\bar{C}$ as the largest C' such that $C\cap C' = \phi$), then, if TV consists of two elements, 0 and 1, we obtain the two-element Boolean algebra $\langle\{0,1\},\cap,\cup,-\rangle$. But as soon as TV

[1] This argument is certainly not new. It can be found e.g. in Rescher (1969). Yet it seems to have failed to discourage people from trying to use multi-valued logic in contexts where the argument shows it to be inadequate.

contains more than two elements the resulting algebra is not a Boolean algebra. In particular the equation $C \cup \bar{C} = 1$ will no longer be satisfied by all elements C.

3.

This last observation suggests in which direction a solution to our difficulty might be found: we should choose as truth-value sets not linear orderings, but rather sets which, like the two-valued system, display the structure of the propositional calculus, viz. Boolean algebras. We may then, if we want to, further 'reduce' these Boolean algebras to linearly ordered systems; in this reduction different Boolean values may be assigned the same element of the linear ordering. But this will no longer affect our semantic characterization of the connectives, as these will now be defined in the Boolean truth-value space and thus not directly on the linearly ordered, 'ultimate' truth values themselves. It may now happen that even if ϕ and ψ have the same ultimate value, $\phi \cap \chi$ has a different ultimate value from that of $\psi \cap \chi$ (viz. in certain cases where the Boolean values of ϕ and ψ which reduce to the same ultimate value are nevertheless distinct).

This idea is by now quite familiar to logicians. Yet it is surprising that its use for the specific problems with which we are here concerned occurred as late as it did; for there is a branch of mathematics, viz. probability theory, in which it has been accepted as the standard solution to what is in many ways the same problem as the one we are facing here. I am referring to the theory of probability in the definitive mathematical form that Kolmogorov (1970) gave it in the 1930s. In this theory one associates with a proposition in first instance a certain set. With this set is associated in turn a real number in the closed interval [0:1]. This number gives the probability of the proposition. Now while the set associated with the conjunction of two propositions is a simple function of the sets associated with the conjuncts, viz. their intersection, there is no way of telling in general the probability of the conjunction on the basis of just the probabilities of the conjuncts; and this is as it should be, for when p and q each have probability $\frac{1}{2}$, the probability of $p \wedge q$ could, intuitively, be anything between 0 and $\frac{1}{2}$.

Perhaps the main philosophical problem which this approach raises is that of giving a plausible interpretation of the sets with which propositions are associated. I will consider here only one doctrine, according to which the elements of the sets are regarded as possible worlds, or possible situations. Thus the probability of a proposition is measured in terms of the set of those possible worlds in which it is true. Seen in this light, probability theory is closely connected with the possible-world semantics for modal and other types of non-extensional logic. Both theories associate with a given sentence in any particular interpretation a set (of possible worlds, points of reference, contexts, etc.).

Of course, probability theory and intensional logic are concerned with different sorts of problems. In intensional logic we are primarily concerned with the analysis, in terms of the set of all possible 'worlds' (as well as, perhaps, various structural properties of this set), of the semantical function of certain non-truth functional operators, such as *it is necessarily the case that*. In probability theory one does not consider such intensional operators, but concentrates on the probability function which associates real numbers with the sets, and investigates how the probabilities of certain complex expressions depend on the probabilities of their components – often under certain assumptions about these components, such as independence or disjointness.

There is another theory of formal semantics which fits within the general frame which we are now discussing, viz. the theory of partial interpretations and supervaluations. This theory is, in its simplest form, a generalization of ordinary two-valued model theory, which allows for the possibility that in a given interpretation for a certain formal language (say, of ordinary first order logic with description operator) some sentences of the language are neither true nor false. Yet, in order to avoid the – from a certain standpoint undesirable – consequence that whenever p is without truth value, so are, among others, $p \wedge \neg p$ and $p \vee \neg p$, one considers the collection of all interpretations which extend the given interpretation by filling out its truth gaps in a consistent manner. If a formula comes out true in each of these completions it will be regarded as true in the interpretation even if it is not assigned a truth value directly by the (incomplete) recursive definition of truth. Similarly it will be counted as false in the interpretation if it is false in each of its completions.

One may view this process again as one of assigning, in a given interpretation, sets to sentences: to each sentence is assigned the set of all completions in which it is true. Sentences already true in the given interpretation, and also such sentences as $p \vee \neg p$ where p itself is not assigned a truth value directly, will be assigned the set of all completions, those already false as well as sentences such as $p \wedge \neg p$ will be assigned the empty set; only sentences which neither have a truth value in virtue of the recursive truth definition nor have the form of a logical identity or contradiction may be assigned intermediate sets.

The theory, which was first introduced by Van Fraassen (1969) suggests in what way the framework under discussion might be used in an analysis of vagueness. Vagueness is one of the various reasons why certain sentences may be without truth value. Thus if we regard the world, or any specific speech situation in it, as providing an interpretation for English, what it provides is at best a *partial* interpretation.[1] For such a partial interpretation we may consider the various completions in which all instances of vagueness are resolved in one way or another. The quantity of such completions in which a certain sentence is true ought then to be in some sense a measure for the degree to which the sentence is true in the original interpretation. Such considerations would of course apply not only to adjectives but to other parts of speech as well, in particular to those grammatical categories which, like adjectives, are usually treated as 1-place predicates in simple-minded predicate-logic-symbolizations of English sentences, viz. common nouns and intransitive verbs. (I will later try to say something about systematic differences between the semantic behaviour of adjectives and that of these other two categories.)

These considerations naturally lead to a modification of the model theory for formal or natural languages which I will exemplify for a rather simple case, viz. first order predicate logic. The example will make it clear enough how one could adapt in a similar fashion more complicated model theories – such as those for intensional logics or for fragments of natural languages.

Let us consider the language L for predicate logic, the logical symbols of which are $\neg$, $\wedge$, $\exists$, and the variables v_1 v_2, $v_3, \ldots$, and the non-logical symbols of which are the *n*-place predicate letters Q_i^n ($n = 1,2,3,\ldots$; $i = 1,2,3,\ldots$).

A *classical model for L* is a pair $\langle U,F\rangle$ where (i) U is a non-empty set and (ii) F assigns to each Q_i^n an *n*-place relation on U.

The *satisfaction value of* a formula ϕ of L *in* $M = \langle U,F\rangle$ *by* an assignment a of elements of U to the variables (in symbols $[\phi]_{M,a}$) is defined by the usual recursion:

(i) $[Q_i^n(v_{i_1}\ldots v_{i_n})]_{M,a} = 1$ iff $\langle a(v_{i_1}),\ldots,a(v_{i_n})\rangle \in F(Q_i^n)$
(ii) $[\neg\phi]_{M,a} = 1$ iff $[\phi]_{M,a} = 0$
(iii) $[\phi\wedge\psi]_{M,a} = 1$ iff $[\phi]_{M,a} = 1$ and $[\psi]_{M,a} = 1$
(iv) $[(\exists v_i)\phi]_{M,a} = 1$ if for some $u \in U$ $[\phi]_{M,[a]^u_{v_i}} = 1$

[1] Of course there are other factors whose effect is that a situation of speech will in general determine only a partial interpretation; for example, many predicates are not applicable to individuals of certain kinds – and yet not every statement which attributes a predicate to such a semantically improper object should be regarded as illformed. Such other sources of interpretational incompleteness, however, will here not concern us.

A sentence ϕ of L is *true in* M if $[\phi]_{M,a} = 1$ for some a.

For any model M, Tr_M will be the set of sentences of L which are true in M and Fa_M, the set of sentences of L false in M.

A *partial model* for L is a pair $\langle U,F\rangle$ where (i) U is a non-empty set and (ii) F assigns to each letter Q_i^n an ordered pair $\langle F^+(Q_i^n),F^-(Q_i^n)\rangle$ of disjoint *n*-place-relations on U.

N.B. I will assume throughout that M is a model and is of the form $\langle U,F\rangle$.

The *satisfaction value* of a formula ϕ in M by an assignment a is now defined by:

(i) (a) $[Q_i^n(v_{j_1},\ldots,v_{j_n})]_{M,a} = 1$ if $\langle a(v_{j_1}),\ldots,a(v_{j_n})\rangle \in F^+(Q_i^n)$
(b) $[Q_i^n(v_{j_1},\ldots,v_{j_n})]_{M,a} = 0$ if $\langle a(v_{j_1}),\ldots,a(v_{j_n})\rangle \in F^-(Q_i^n)$
(ii) (a) $[\neg\phi]_{M,a} = 1$ if $[\phi]_{M,a} = 0$
(b) $[\neg\phi]_{M,a} = 0$ if $[\phi]_{M,a} = 1$
(iii) (a) $[\phi\wedge\psi]_{M,a} = 1$ if $[\phi]_{M,a} = 1$ and $[\psi]_{M,a} = 1$
(b) $[\phi\wedge\psi]_{M,a} = 0$ if $[\phi]_{M,a} = 0$ or $[\psi]_{M,a} = 0$
(iv) (a) $[(\exists v_i)\phi]_{M,a} = 1$ if for some $u \in U$ $[\phi]_{M,[a]^u_{v_i}} = 1$
(b) $[(\exists v_i)\phi]_{M,a} = 0$ if for all $u \in U$ $[\phi]_{M,[a]^u_{v_i}} = 0$

Again a sentence ϕ is said to be *true in* M if $[\phi]_{M,a} = 1$ for some a; and ϕ is *false in* M if for some a $[\phi]_{M,a} = 0$

Again Tr_M is the set of true, and Fa_M the set of false sentences of L in M. But now it is clearly possible that certain sentences are neither true nor false in M, so that $Tr_M \cup Fa_M$ does not coincide with the set of all sentences of L.

The partial model $M = \langle U,F\rangle$ is said to be *at least as vague as* the partial model $M' = \langle U,F'\rangle$ (in symbols: $M \subseteq M'$) if for each Q_i^n:

$$F^+(Q_i^n) \subseteq F'^+(Q_i^n) \text{ and } F^-(Q_i^n) \subseteq F'^-(Q_i^n)$$

(Thus $Tr_M \cup Fa_M \subseteq Tr_{M'} \cup Fa_{M'}$.)

To each classical model $\langle U,F\rangle$ for L corresponds a unique partial model, viz. the model $\langle U,F'\rangle$ where for each Q_i^n, $F'^+(Q_i^n) = F(Q_i^n)$ and $F'^-(Q_i^n) = U^n - F(Q_i^n)$. Classical models, as well as the partial models corresponding to them, will be referred to as *complete* models.

A classical model M is called a *completion* of a partial model M' if M' is at least as vague as (the partial model corresponding to) M.

4.

In the theory of supervaluation one considers partial models in conjunction with *all* their completions. What I want to do here is formally almost the same – but with one crucial exception. Rather than all completions of a given partial model, we consider only a certain subset of them. In addition we consider a probability function over a field of subsets of this set of completions,[1] which contains, in particular, for each formula and assignment of elements to its free variables the set of all completions in which the former satisfies that assignment. (This condition warrants that each sentence has a measure.) The complex consisting of the partial model, the set of completions, the field over that set and the probability function over that field I will call a *vague model*. Formally

A *vague model for* L is a quadruple $\langle M,\mathcal{S},\mathcal{F},p\rangle$ where
(i) M is a partial model for L;
(ii) $\mathcal{S}$ is a set of classical models for L which are completions of M;
(iii) $\mathcal{F}$ is a field of subsets over $\mathcal{S}$
(iv) for each $\phi\in L$ and assignment a in the universe of M, $\{M'\in\mathcal{S}: [\phi]_{M',a} = 1\}\in\mathcal{F}$; and
(v) p is a probability measure on $\mathcal{F}$.[2]

Let $\mathcal{M} = \langle\langle U,F\rangle,\mathcal{S},\mathcal{F},p\rangle$ be a vague model for L. For any formula ϕ of L and assignment a of elements of U to the variables the *degree of satisfaction of* ϕ *by* a *in* $\mathcal{M}$, $[\phi]_{\mathcal{M},a}$, is defined as $p(\{M'\in\mathcal{S}: [\phi]_{M',a} = 1\})$. Thus, in particular, if $\phi\in Tr_M$ then $[\phi]_{\mathcal{M}} = 1$ and if $\phi\in Fa_M$ then $[\phi]_{\mathcal{M}} = 0$.

The idea behind the notion of a vague model is this. At the present stage of its development – indeed, at any stage – language is vague. The kind of vagueness which interests us here is connected with predicates. The vagueness of a predicate may be resolved by fiat – i.e. by deciding which of the objects which as yet are neither definitely inside nor definitely outside its extension are to be in and which are to be out. However, it may be that not every such decision is acceptable. For there may already be semantical principles which, though they do not determine of any one of a certain group of objects whether it belongs to the extension or not, nevertheless demand that if a certain member of the group is put into the extension, a certain other member must be put into the extension as well. Take for example the adjective *intelligent*. Our present criteria tell us of certain people that they definitely are intelligent, of certain others that they definitely are not, but there will be a large third category of people about whom they do not tell us either way. Now suppose that we make our standard more specific, e.g., by stipulating that to have an I.Q. over a certain minimum is a necessary and sufficient criterion for being intelligent. Further, suppose that of two persons u_1 and u_2 of the third category u_1 has a higher I.Q. than u_2. Then, whatever we decide this minimum to be, our decision will put u_1 into the extension if it puts u_2 into it. Finally, let us assume for the sake of argument that any way of making the concept of intelligence precise that is compatible with what we already understand that concept to be is equivalent to the adoption of a certain minimum I.Q. Then there will be no completions in the partial model that reflect the present state of affairs and in which u_2 is put into the extension of the predicate but u_1 is not.

Formally, if Q^1_1 represents the adjective *intelligent* and the model $\langle M,\mathcal{S},\mathcal{F},p\rangle$ reflects the situation just described, and M, in particular, that which obtains before any of the possible precise definitions has been adopted, then u_1 and u_2 are both members of $U-(F^+(Q^1_1)\cup F^-(Q^1_1))$ and there is no model $M'\in\mathcal{S}$ such that $u_2\in F(Q_1)$ and $u_1\notin F(Q_1)$.

My original motivation in setting up this framework was to give a uniform characterization of the operation which transforms adjectives into their comparatives. Let us see if this is now possible.

The relation x *is more* A *than* y (where A is any adjective) can be defined in terms of the relation x *is at least as* A *as* y by

(11) x is more A than y if and only if x is at least as A as y and it is not the case that y is at least as A as x

Therefore a semantic characterization of this second relation will automatically give us one for the first as well. As there are minor but undeniable advantages in discussing the relation *at least as…as* I will concentrate on that concept.

[1] A *field of subsets* of a given set X (or: a *field over* X) is a set of subsets of X, such that (i) $X\in\mathcal{F}$; (ii) $\Phi\in\mathcal{F}$; (iii) if $Y,Y'\in\mathcal{F}$ then $Y\cap Y'$, $X-Y\in\mathcal{F}$.
A *probability function over* a field $\mathcal{F}$ over X is a function p whose domain is $\mathcal{F}$, whose range is included in the real interval $[0,1]$, and which has the properties: (i) $p(X) = 1$; (ii) if $Y\in\mathcal{F}$, then $p(X-Y) = 1-p(Y)$; and (iii) if $\mathcal{G}$ is a countable subset of $\mathcal{F}$ such that (a) whenever $Y,Y'\in\mathcal{G}$ and $Y\not\equiv Y'$ then $Y\cap Y' = \Phi$; and (b) $\cup\mathcal{G}\in\mathcal{F}$; then $p(\cup\mathcal{G}) = \sum_{Y\in\mathcal{G}} p(Y)$.

[2] From the mathematical point of view this notion is unproblematic only if the universe U is finite. In that case we do not really need to require that p satisfy the condition (iii) of n. 1, but only the weaker condition obtained by replacing the word *countable* in (iii) by *finite*. (iii) is necessary when U is denumerable; in that case, however, as well as when U is uncountable, it may happen that no intuitively correct models exist. The only way in which I can see how to cope with these cases involves non-standard analyses. I do not want to go into this here.

Let us assume that some of the one-place predicates of L represent adjectives, in particular Q_i^l. We add to L the operator symbol $\geqslant$. $\geqslant$ forms out of one one-place predicate Q_i^l a two-place relation $\geqslant(Q_i^l)$. $\geqslant(Q_i^l)(x,y)$ should be read as *x is at least as Q_i^l as y*. (What relation $\geqslant(Q_i^l)$ might represent when Q_i^l is not an adjective is of no concern to us now.) Let L' be the language resulting from the addition of $\geqslant$ to L. Let $\mathcal{M} = \langle M, \mathcal{S}, \mathcal{F}, p\rangle$ be a vague model for L. In order to expand $\mathcal{M}$ to a model for L' we must determine the positive and negative extensions of the relation $\geqslant(Q_i^l)$ in M as well as its extensions in all the members of $\mathcal{S}$. To begin we will consider just the positive extension in M. Two possible definitions come to mind. According to the first an element u_1 of U stands (definitely) in the relation to u_2 if for every member M' of $\mathcal{S}$ in which u_2 belongs to the extension of Q_i^l u_1 belongs to that extension as well. So we get, representing the positive extension of $\geqslant(Q_i^l)$ in M as $F^+(\geqslant(Q_i^l))$,

(12) for all $u_1, u_2 \in U$, $\langle u_1,u_2\rangle \in F^+(\geqslant(Q_i^l))$ iff $[Q_i^l(v_1)]_{\mathcal{M},a_2} \subseteq [Q_i^l(v_1)]_{\mathcal{M},a_1}$ (where a_1 and a_2 are any assignments with $a_1(v_1) = u_1$ and $a_2(v_1) = u_2$, respectively)

According to the second definition u_1 stands in the relation to u_2 if the measure of the set of completions in which u_1 belongs to the extension of Q_i^l is at least as large as that of the set of completions in which u_2 belongs to the extension. So we obtain

(13) for all $u_1,u_2 \in U$, $\langle u_1,u_2\rangle \in F^+(\geqslant(Q_i^l))$ iff $p([Q_i^l(v_1)]_{\mathcal{M},a_2}) \leqslant p([Q_i^l(v_1)]_{\mathcal{M},a_1})$ where a_1 and a_2 are as above

Before we consider the relative merits of these definitions let us first remove a flaw which they share. Neither (12) nor (13) allows for the possibility that the comparative relation holds between two objects for each of which it is beyond doubt that it satisfies the positive. For if both u_1 and u_2 belong to $F^+(Q_i^l)$ then $[Q_i^l(v_1)]_{M,a_1} = [Q_i^l(v_1)]_{M,a_2} = 1$ and so both (12) and (13) would exclude $\langle u_1,u_2\rangle$ from $F^+(\geqslant(Q_i^l))$.

It seems that the only way in which we could meet this difficulty without departing too much from our present format is this: Instead of a vague model, consisting of a partial model M, a field $\mathcal{F}$ over a set $\mathcal{S}$ of completions of M and a probability function p over that field, we need to consider models $\mathcal{M}$, in which the set $\mathcal{S}$ comprises besides completions of M also complete models which in certain ways conflict with M. Such models will

represent (hypothetical) situations in which the standards for a predicate are set so high that certain objects which already have that predicate in M now fail to have it – or else in which the standards are set so low that objects belonging to the negative extension of the predicate in M now fall in its positive extension. This leads us to the following modification of the notion of a vague model:

A *graded model for L* is a quadruple $\langle M, \mathcal{S}, \mathcal{F}, p\rangle$, where

(i) M is a partial model for L
(ii) $\mathcal{S}$ is a set of classical models for L with universe U
(iii) $\mathcal{F}$ is a field over $\mathcal{S}$
(iv) For each formula ϕ of L and each assignment a to elements of the universe of M, $\{M' \in \mathcal{S}: [\phi]_{M',a} = 1\} \in \mathcal{F}$
(v) $\{M' \in \mathcal{S}: M'$ is a completion of $M\} \in \mathcal{F}$; and
(vi) p is a probability function over $\mathcal{F}$.

We may then define the degree of truth of a sentence of L just as before except that we now consider the conditional probability of a certain set of completions of M on the set of all completions of M in $\mathcal{S}$. On the other hand the characterization (13) of the comparative of Q_i^l is now no longer vulnerable to the objection which led us to the introduction of graded models.

Let us consider an example. Suppose that Q_1^l represents the adjective *heavy*; that all other predicates represent properties of, and relations between, material objects, and (for simplicity) that Q_1^l is the only vague predicate. Let U be the set of material objects and let $\mathcal{M} = \langle M, \mathcal{S}, \mathcal{F}, p\rangle$ be a graded representation (restricted to material objects) of the actual world. What should in this case $\mathcal{S}$ and p be? As regards $\mathcal{S}$ a simple answer seems possible in this special case: For each particular real number r there will be a member M of $\mathcal{S}$ in which the extension of Q_1^l consists of those objects whose weight (in grams) exceeds r.

It is not possible to say precisely what the function p should be. But this much seems beyond doubt: there should be a strictly monotonic function f from the set of all positive real numbers into the interval [0,1] so that for any object u with weight r, $p([Q_1^l(v_1)]_{\mathcal{M},a} = f(r)$ (for some a with $a(v_1) = u$). Thus, the greater u's weight, the larger the class of members of $\mathcal{S}$ in which Q_1^l is true of u, and the greater the measure (or 'intermediate truth value') of the formula $Q_1^l(v)$ under a_1.

A "patch"? Note that thinking in terms of extents does not trigger this problem.

5.

We should now compare (12) and (13). According to (12) u_1 is at least as heavy as u_2 just in case the set of models in which u_1 *is heavy* is true includes the class of those which render u_2 *is heavy* true; and this will be the case if and only if u_1 has greater or equal weight. Indeed, within the context of the present example, (12) is precisely the proposal that can be found in Lewis (1970), where it is attributed to David Kaplan.

According to (13) u_1 will be at least as heavy as u_2, provided u_2 *is heavy* is true in a set of models with measure greater than or equal to that of the set of models in which u_2 *is heavy* is true. Again this will be true if and only if u_1 has greater or equal weight. Thus for this special case the two definitions are equivalent.

But this need not always be so. Suppose for example that Smith, though less quick-witted than Jones, is much better at solving mathematical problems. Is Smith cleverer than Jones? This is perhaps not clear, for we usually regard quick-wittedness and problem-solving facility as indications of cleverness, without a canon for weighing these criteria against each other when they suggest different answers. When faced with the need to decide the issue, various options may be open to us. We might decide that really only problem-solving counts, so that after all, Smith is cleverer than Jones; or we might decide on a particular method for weighing the two criteria – so that Smith's vast superiority at solving problems will warrant that in spite of Jones's slight edge in quick-wittedness Smith is cleverer than Jones; or we might decide that only quick-wittedness counts; and this time Jones will come out as the cleverer of the two.

It is not clear how the probability function of a graded model $\mathcal{M}$ representing this situation should be defined. Yet, if we assume that the third decision is less plausible than either the first or the second, then we should expect members of $\mathcal{S}$ which are compatible with that decision to have no more weight than those which are compatible with other decisions. Further, relatively few models of the first sort will be such that Jones belongs to the extension of *clever* and Smith not; for Jones is not that much quicker in conversation. But, because of the disparity in problem-solving ability, many models compatible with the first decision, as well as a good many that are compatible with the second, will have Smith in the extension of *clever* but not Jones. Given all this, we would expect the measure of the set of members of $\mathcal{S}$ in which Smith belongs to the extension to be greater than that of those members where Jones belongs to the extension. So by (13) and (11) we would have to conclude that Smith is cleverer than Jones.

But do we want to say this? I think not. Before any decision has been made it is true neither that Smith is cleverer than Jones nor that Jones is cleverer than Smith. This intuitive judgement is in agreement with (12), according to which Jones and Smith are incomparable in respect of cleverness. Indeed, it is (12) which, in my opinion, captures the comparative correctly – at least to the extent that it gives a necessary and sufficient condition for *definite* membership in the positive extension of $\geqslant(Q_1^{\prime})$. That (13) cannot be right becomes even more evident when we realize that it implies that for any objects u_1 and u_2 and adjective A, either u_1 is at least as A as u_2 or u_2 is at least as A as u_1; and this should fail to be true in general whenever we have two, largely independent, criteria for applicability of the adjective, but no clear procedure for weighing them.

We saw that for *heavy* (12) and (13) are equivalent (provided p has been correctly specified). The same is true for a number of other adjectives which, like *heavy*, may be called 'one-dimensional'. With each such adjective is associated a unique measurable aspect. The (numerical) value of that aspect for a given object determines whether or not the adjective applies. For *heavy* the aspect is weight. Other examples are *tall* (associated with height) and *hot* (associated with temperature).

But such adjectives are rare. Even *large* is not one of them. For what precisely makes an object large? Its height? or its volume? or its surface? or a combination of some of these? Here we encounter the same phenomenon that has already been revealed by our discussion of *clever*. There is no fixed procedure for integrating the various criteria. Often it is the context of use which indicates how the criteria should be integrated or, alternatively, which of them should be taken as uniquely relevant.

This is one of the various ways in which contexts disambiguate. Formally, contextual disambiguation can be represented as a function from contexts to models less vague than the ground model. While incorporating this idea into the framework already adopted, I will at the same time eliminate a feature of vague and graded models which is unrealistic in any case but would be particularly out of place in the context-dependent models defined below: Thus far I have assumed that all the members of $\mathcal{S}$ are complete models. But this is unnatural if we want to think of these models as the results of semantical decisions that could actually be made. For most decisions will fail to render the relevant predicates completely sharp. They will only make them sharper. (Indeed, we may with Wittgenstein, doubt that we could ever make any concept completely sharp.) It therefore appears more natural to posit that the members of $\mathcal{S}$ are partial models. It is possible, moreover, that one of these contextually determined models is less vague than another, viz. when, intuitively speaking, the semantic decision reflected by the second goes in the same direction, but not as far as, that reflected by the first. Thus $\mathscr{S}$ will be partially ordered by the relation *as vague as*.

I just suggested that a context picks from this set a particular model – which functions, so to speak, as the ground model of the graded model which represents the speech situation determined by that context. The various sharpenings acceptable from the viewpoint of that context would then be represented by those members of $\mathscr{S}$ which are at most as vague as the new ground model. But I am not convinced that this is absolutely correct. For it could conceivably be the case that two different contexts specify for a given predicate two different criteria from the set of those which are prima facie plausible and which, though they happen to determine the same new ground model, will not permit exactly the same further sharpenings. So the context should select a certain subset of $\mathscr{S}$ of contextually admissible further sharpenings. In addition, the context must select a subset of admissible modifications. This set we could not even hope to reconstruct from the new ground model alone.

Thus far the probability function p was defined over a class of complete models. This would now seem to be impossible as we no longer require that $\mathscr{S}$ consists of – or even that it contains any – complete models. Yet the intuition behind the function p – which I tried to convey in the example concerning *heavy* – makes it appear unnatural to define p as a function over sets of *partial* models, especially as these sets may now be expected to contain models one of which is vaguer than the other; it is, so to speak, the number of possible *ultimate* results of repeated sharpening that p should measure, and not, the number of intermediate steps that one may take on the way to these ultimate complete models.

A solution to this dilemma can be found if we assume that all individual cases of vagueness can be resolved, though not all at once; and this assumption does appear to be unexceptionable. Thus we will impose on the set $\mathscr{S}$ the following condition:

(14) if $\langle U,F_1\rangle \in \mathscr{S}$ and $\langle u_1,\ldots,u_n\rangle \in U^n - (F_1^+(Q_i^n) \cup F_1^-(Q_i^n))$, then there is a member $\langle U,F_2\rangle$ in which Q_i^n is less vague than $\langle U,F_1\rangle$ and such that $\langle u_1,\ldots,u_n\rangle \in F_2^+(Q_i^n) \cup F_2^-(Q_i^n)$.

Under this assumption we may construct complete models as the unions of maximal chains in $\mathscr{S}$: Let $\mathscr{S}$ be a set of partial models for L which all have the same universe U. Then $\mathscr{S}$ is a *chain under the relation 'vaguer than'* if for any two of its members $\langle U,F_1\rangle$, $\langle U,F_2\rangle$ either

(i) for each predicate Q_j^n of L, $F_1^+(Q_j^n) \subseteq F_2^+(Q_j^n)$ and $F_1^-(Q_j^n) \subseteq F_2^-(Q_j^n)$ or
(ii) for each predicate Q_j^n of L, $F_2^+(Q_j^n) \subseteq F_1^+(Q_j^n)$ and $F_2^-(Q_j^n) \subseteq F_1^-(Q_j^n)$.

A subset $\mathscr{S}'$ of a set $\mathscr{S}$ of models with universe U, is a *maximal chain in* $\mathscr{S}$ if (i) $\mathscr{S}'$ is a chain (under the relation *vaguer than*) and (ii) for any $M' \in \mathscr{S} - \mathscr{S}'$, $\mathscr{S}' \cup \{M'\}$ is not a chain. The *union of* a chain $\mathscr{S}$ of models with universe U is the model $\langle U,F_\infty\rangle$ where for each $Q_j^n F_\infty^+(Q_j^n) = \bigcup_{\langle U,F\rangle \in \mathscr{S}} F^+(Q_j^n)$ and $F_\infty(Q_j^n) = \bigcup_{\langle U,F\rangle \in \mathscr{S}} F(Q_j^n)$

If U is countable then (14) entails that

(15) The union of each maximal chain of $\mathscr{S}$ is complete

However, (15) does not follow automatically from (14) when U is uncountable. Since it is property (15) in which we are primarily interested in connection with the function p, I will make it, rather than (14), one of the defining conditions of graded context-dependent models.

A *graded context-dependent model for* L is a quintuple $\langle M,\mathscr{S},\mathscr{C},\mathscr{F},p\rangle$ where

(i) M is a partial model;
(ii) $\mathscr{S}$ is a set of partial models with the same universe as M;
(iii) The union of each maximal chain of $\mathscr{S}$ is complete;
(iv) $\mathscr{C}$ is a function the range of which consists of pairs $\langle M',\mathscr{S}'\rangle$ where (a) $M' \in \mathscr{S}$; (b) $\mathscr{S}' \subseteq \mathscr{S}$ and (c) the union of each maximal chain of $\mathscr{S}'$ is complete;
(v) $\mathscr{F}$ is a field over the set $\overline{\mathscr{S}}$ of unions of maximal chains of $\mathscr{S}$;
(vi) (a) for each formula ϕ and assignment a the set of members M' of $\mathscr{S}$ such that $[\phi]_{M',a} = 1$ belongs to $\overline{\mathscr{S}}$; (b) $\{M' \in \overline{\mathscr{S}}: M \subseteq M'\} \in \mathscr{F}$; (c) for each $\langle M',\mathscr{S}'\rangle$ in the range of $\mathscr{C}$ if $\overline{\mathscr{S}}'$ is the set of unions of maximal chains of $\mathscr{S}'$ then $\{M'' \in \overline{\mathscr{S}}': M' \subseteq M''\} \in \mathscr{F}$;
(vii) p is a probability function over $\mathscr{F}$.

We will refer to context-dependent graded models by means of the abbreviation cgm. Henceforth $\mathscr{M}$ will always be a cgm and will always be equal to $\langle M,\mathscr{S},\mathscr{C},\mathscr{F},p\rangle$; M will be called the *ground model* of $\mathscr{M}$; similarly if $\mathscr{C}(c) = \langle M'_c,\mathscr{S}'_c\rangle$ then M'_c is called the *ground model* (in $\mathscr{M}$) *with respect to c*.

Again we denote the set of members of $\overline{\mathscr{S}}$ in which ϕ is true under a as $[\phi]_{\mathscr{M},a}$, where $\overline{\mathscr{S}}$ is again the set of unions of maximal chains of $\mathscr{S}$.

Similarly, if $\mathscr{C}(c) = \langle M_c, \mathscr{S}_c \rangle$, $[\phi]_{\mathcal{M},c,a}$ is the set of members of $\overline{\mathscr{S}}_c$ in which ϕ is true under a.

The domain of $\mathscr{C}$ should be thought of as the set of contexts. Contexts may be more or less specific; correspondingly Dom $\mathscr{C}$ may contain elements c and c' such that $M_c \leqslant M_{c'}$ and $\mathscr{S}_c \subseteq \mathscr{S}_{c'}$; in this case c will be at least as specific as c'. Thus the members of Dom $\mathscr{C}$ are partially ordered by the relation $\leqslant$, defined by: $c \leqslant c'$ iff $M_c \subseteq M_{c'}$ and $\mathscr{S}_c \subseteq \mathscr{S}_{c'}$. One may wonder if for every member M' of $\mathscr{S}$ there should be a c such that M' is the ground model with respect to c. This would mean that for any possible sharpening of a predicate there is a context which indicates that the predicate should be understood in precisely *that* sharper way. I have no argument to show that this assumption is false; yet I see no gain from it; thus I prefer not to make it.

In a cgm it is possible that while the relation $\geqslant(Q^l)$ does not hold in the ground model it does hold in the ground models of certain contexts. Thus assume Q^l_l represents the adjective *clever*; further assume that c_1 represents a context in which *clever* must be understood as '*good at solving problems*'; that c_2 represents a context in which *clever* must be understood as '*quick-witted*'; and that c_3 represents a context on which both quick-wittedness and the ability to solve problems are to be regarded as constitutive of cleverness. Then we may expect that if $a_1(v_1)$ = Smith and $a_2(v_1)$ = Jones,

(a) $[Q^l_l(v_1)]_{\mathcal{M},c_1,a_2} \subseteq [Q^l_l(v_1)]_{\mathcal{M},c_1,a_1}$; and
(b) $[Q^l_l(v_1)]_{\mathcal{M},c_2,a_1} \subseteq [Q^l_l(v_2)]_{\mathcal{M},c_2,a_2}$

while nothing definite can be said about the relation between $[Q_l(v_1)]_{\mathcal{M},c_3,a}$ and $[Q^l_l(v_1)]_{\mathcal{M},c_3,a_2}$ until more is known about whether, and in what way, c_3 determines how the two criteria for *clever* are to be weighed. In order that (a) and (b) formally guarantee that in c_1 Smith is cleverer than Jones, while in c_2 Jones is cleverer than Smith, we must specify, parallel to (13)

(16) if $\mathscr{C}(c) = \langle\langle U,F_c\rangle, \mathscr{S}_c\rangle$ and $u_1,u_2 \in U$ then $\langle u_1,u_2\rangle \in F_c^{+}(\geqslant(Q^l_i))$ if and only if $[Q^l_i(v_1)]_{\mathcal{M},c,a_2} \subseteq [Q^l_i(v_1)]_{\mathcal{M},c,a_1}$.

Since not every member of $\mathscr{S}$ is necessarily the ground model with respect to some context, (16) may not define the positive extension of $\geqslant(Q^l_i)$ for some of these models. This is of little practical importance. If we insist on defining the extensions in these models as well, we may stipulate that for any such model $\langle U,F\rangle$, $\langle u_1,u_2\rangle \in F_1^{+}(\geqslant Q^l_i))$ if and only if for some $c, \mathscr{C}(c) = \langle M_c,\mathscr{S}_c\rangle$, $\langle U,F_1\rangle \in \mathscr{S}_c$, $M_c \subseteq \langle U,F_1\rangle$ and $\langle u_1,u_2\rangle \in F_c^{+}(\geqslant(Q^l_i))$.

What is the negative extension of $\geqslant(Q^l_i)$? It should consist in the first place of those pairs $\langle u_1,u_2\rangle$ of which it is definitely true that u_2 is more Q^l_i than u_1, i.e. in view of (12), those pairs for which

(17) $[Q^l_i(v_1)]_{\mathcal{M},a_1} \subsetneqq [Q^l_i(v_1)]_{\mathcal{M},a_2}$.

One might question this condition on the ground that it makes *u_2 is more Q^l_i than u_1* definitely true also when the difference between u_1 and u_2 is only marginal. But I do not believe that the objection is well-founded. However marginal the difference, if it is a difference in an aspect which is irrevocably bound to the predicate, so that no context can break this tie, then the relation definitely obtains irrespective of whether it is difficult, or even physically impossible, to observe this.

This leaves us with those pairs $\langle u_1,u_2\rangle$ such that neither $[Q^l_i(v_1)]_{\mathcal{M},a_1} \subseteq [Q^l_i(v_1)]_{\mathcal{M},a_2}$ nor $[Q^l_i(v_1)]_{\mathcal{M},a_2} \subseteq [Q^l_i(v_1)]_{\mathcal{M},a_1}$. Which of these should go into $F^{-}(\geqslant(Q^l_i))$? I think none. As long as there are some acceptable ways of sharpening Q^l_i which render u_1 at least as Q^l_i as u_2, the falsehood of *u_1 is at least as Q^l_i as u_2* cannot be definite.

I introduced the probability function to show how the notion of 'degrees of truth' can be made coherent. But so far the function has served to no good purpose. In particular it has proved useless for the characterization of the comparative: once more it turned out to be necessary to define the operation on the sets themselves rather than on the numerical values to which p reduces them.

However, there are expressions the analysis of which does seem to require the function p. Consider *rather*. *Rather* forms adjectives out of adjectives, e.g. *rather tall* out of *tall*, *rather clever* out of *clever*, etc. When is a person rather clever? Before I can discuss the really important aspects of this question, I should first settle a minor point. *x is rather clever* sometimes seems to deny that x is clever, while on other occasions it appears to be entailed by the fact that x is clever – just as e.g. *most x are F* sometimes seems to entail *not all x are F*, while on other occasions it seems to be entailed by *all x are F*. I think that both cases, as well as a great many similar ones, ask for an explanation involving Grice's theory of implicature: *most x are F* is a consequence of *all x are F*; but when uttered by a speaker whom the hearer assumes to know whether all x are F, it will convey that not all x are F – for if all x were F, why would not the speaker have said so? Similarly, *rather clever* is weaker than *clever*. But one would use the longer phrase only if one had doubts that the shorter applies.

Thus *x is rather clever* is weaker than *x is clever*. *x* is rather clever if a certain lowering of the standards for cleverness would make *x* clever, i.e. if the proportion of members of $\overline{\mathscr{F}}$ in which *x* belongs to the extension of *clever* is large enough. Indeed the closer *x* is to being truly clever, the smaller is the modification of the standards that is required, and thus the larger will be the class of those models where *x* is in the extension.

It should be noted that just as *x* may pass the test of cleverness for different reasons, so he may also pass that of being rather clever in a variety of ways. Thus it is possible that *x*, *y* and *z* are all rather clever (though not unambiguously clever); *x*, because he is remarkably quick-witted, while hopeless at mathematical problems; *y*, because he is good at such problems, though slow in conversation; and *z*, because he has both capacities to a moderate degree. For any two of *x*, *y*, *z*, there will be certain modifications of standards which will warrant membership in the extension of *clever* for one but not for the other. Thus it will be true of the set of members of $\overline{\mathscr{F}}$ in which, say, *x* is in the extension of clever and the set of those members of $\overline{\mathscr{F}}$ where the extension contains, say, *y*, that neither will include the other. Yet they both guarantee membership in the extension of *rather clever*, essentially because they are both large enough. It is this intuition concerning the largeness of sets which p tries to capture.

Thus if *clever* is again represented by Q_I^l, then we may put:

u_1 is rather clever if and only if $[Q_i^l(v_1)]_{\mathcal{M},a_1} \geqslant p_0$ (where p_0 is some number in (0,1).

Obviously p_0 should be less than $p(\{M' \in \overline{\mathscr{F}}: M \subseteq M'\})$; but not much more can be said about it. For of course p_0 is not fixed. If that were so, *rather clever* would be a sharp predicate, which evidently it is not.

The vagueness of *rather* could be represented in the following way. We associate with each $c \in \mathrm{Dom}\,\mathscr{C}$ a pair of real numbers r_c^-, r_c^+ between 0 and 1 such that whenever $c \leqslant c'$, then $r_c^- \leqslant r_{c'}^- < r_{c'}^+ \leqslant r_c^+$. The positive and negative extensions of *rather* Q_I^l in the ground model M_c *w.r.t.c* are then defined as the sets

$$\{u \in U: [Q_i^l(v_1)]_{\mathcal{M},c,a} > r_c^+ \cdot p(\{M' \in \mathscr{S}_c: M_c \subseteq M'\})\}$$
$$\{u \in U: [Q_i^l(v_1)]_{\mathcal{M},c,a} < r_c^- \cdot p(\{M' \in \mathscr{S}_c: M_c \subseteq M'\})\},$$

respectively; finally the intermediate value of *u is rather* Q_i^l in the ground model is given by $p(\{M' \in \overline{\mathscr{F}}: u$ belongs to the positive extension of *rather* Q_i^l in $M'\})$.

There are a number of words which, like *rather*, form adjectives out of adjectives and which can be analysed along similar lines. Another prominent example is *very*. The extension of *very* Q_i^l is again a function of $[Q_i^l(v_1)]_{\mathcal{M},a}$. The limit which $[Q_i^l(v_1)]_{\mathcal{M},a}$ must exceed in order that $a(v_1)$ belong to the extension of *very* Q_i^l must be larger, and not smaller, than $p(\{M' \in \mathscr{S}; M \subseteq M'\})$.

6.

For traditional logic adjectives, nouns and intransitive verbs are all of a kind – viz. one-place predicates. My second theory of adjectives tries to vindicate this view against the one expressed earlier which puts adjectives into a different category than verbs and nouns. Yet it is an undeniable fact about ordinary English that while the comparative is in general a natural operation on adjectives, similar operations on nouns are of relatively little importance, and on verbs they are virtually non-existent. This suggests a difference between adjectives on the one hand and verbs and nouns on the other hand. I will leave verbs out of consideration in the following discussion, as they present problems quite different from those with which this paper is concerned. But I will try to say something about the difference between adjectives and nouns. Why is it that comparisons involving nouns are in general so much more dubious than those which involve adjectives? *This is more a table than that* sounds awkward and is perhaps never unequivocally true, except in the cases in which it is evident that this is a table and that is not (but then we can say precisely this, and thus do not need the first phrase). Yet it appears that nouns too are vague, some of them just as vague as certain adjectives. Why does not their vagueness allow for equally meaningful comparatives? To discover the reasons, it is advantageous to reconsider 'one-dimensional' adjectives.

For any such adjective Q_i^l it will be the case that, for arbitrary a_1, a_2,

(18) either $[Q_i^l(v_1)]_{\mathcal{M},a_1} \subseteq [Q_i^l(v_1)]_{\mathcal{M},a_2}$ or
$[Q_i^l(v_1)]_{\mathcal{M},a_2} \subseteq [Q_i^l(v_1)]_{\mathcal{M},a_1}$;

and this ensures that *u_1 is more Q_i^l than u_2* always has a definite truth value.

We have already seen that most adjectives do not satisfy (18) unambiguously. *u_1 is cleverer than u_2* could remain without truth value in the ground model. Yet there should still be a fair proportion of pairs $\langle u_1, u_2 \rangle$ where u_1 and u_2 both lie in the extension gap of *clever*, but for which (18) holds (with $a_1(v_1) = u_1$ and $a_2(v_1) = u_2$). And, for the same reason, there are many contexts c in which (18) holds (with $[Q_i^I(v_I)]_{\mathcal{M},c,a_I}$ for $[Q_i^I(v_I)]_{\mathcal{M},a_I^1}$, etc.), so that in c each comparative sentence involving Q_i^I is either definitely true or definitely false. On the other hand, if Q_i^I is a noun then (18) will in general be satisfied for very few pairs of objects which both fall in the extension gap of Q_i^I.

A very rough explanation for this formal distinction is the following: In order for an object to satisfy a noun it must in general satisfy all, or a large portion, of a cluster of criteria. None of these we can promote to the sole criterion without distorting the noun's meaning beyond recognition. We cannot, therefore, compare the degrees to which two different objects satisfy the noun in terms of their degrees of satisfaction of just one of its many criteria; and in order to compare them by comparing their ratings with respect to a variety of these criteria we need a method for integrating these various ratings. And such a method is in general not part of the meaning of the noun.

There is another aspect to the difference between nouns and adjectives which is related to the one discussed above, but perhaps even more important in connection with the former's resistance to comparatives. Nouns, though potentially just as vague as adjectives, tend in actual practice to behave much more like sharp predicates. Take *cat*. In principle there could be all sorts of borderline cases for this predicate – but in actual fact there are very few at best. The same is true, be it in slightly varying degrees, of *table*, *rock* or *word*. Thus nouns often have very small, or no, extension gaps in the actual world – even if it is easy to think of possible worlds in which these gaps would be enormous. This gives an additional explanation of why comparatives involving nouns should be of relatively little use. For they are particularly important in those cases where neither of the two objects compared belongs unambiguously to the positive or to the negative extension of the predicate in question. And these cases will seldom arise when the predicate is a noun.

It is an interesting question how nouns 'manage' to be as sharp as they are. The explanation must be more or less along the following lines: Even if each of the several criteria for the noun may apply to actual objects in varying degrees, these criteria tend to be, with respect to the actual world, *parallel*: an object which fails to satisfy a few of them to a reasonable degree will generally fail to score well with regard to almost all of them. Consequently, it will either be recognized as definitely inside the extension of the noun, or else as definitely out. The nature of this parallelism is very much that of a physical law – it is a feature of our world, and thus in essence empirical. This is one of the ways in which the actual structure of

the world shapes the conceptual frame with which we operate, and one of the reasons why it is difficult to separate the empirical from the purely conceptual.

Where the simple comparative of a predicate is non-sensical, addition of certain special expressions can restore its meaningfulness. Examples of such expressions are: *in a sense*, *as far as function is concerned*, *with regard to shape*.[1]

Let us consider this last phrase. How should we analyse

(19) with regard to shape u_1 is more a table than u_2?

First we should determine the logical type of the expression *with regard to shape*. This is really a problem which does not belong in this paper. I will therefore give an answer which is convenient in connection with the issues which concern us here and does not distort them. I will treat *is more... than...with regard to shape* as an atomic, i.e. not further analysable, expression which stands for a new comparative operation – one which again forms binary relations out of predicates. This new comparative differs from the one considered thus far in the following way: The phrase *with regard to shape* places us so to speak in a context where shape is singled out as the only criterion for whatever the property is in respect of which the comparison is made. Let us suppose that there are such contexts – contexts in which those predicates to which shape is at all relevant are evaluated with respect to shape alone. Then (19) should be true in the ground model if and only if *u_1 is more a table than u_2* is true in each of these contexts.

I should like to make a brief comment at this point on the nature of contexts and the role which in my opinion they ought to play in semantic analysis of the sort of which I have tried to give instances in this paper. We could give an alternative, but evidently equivalent, account of (19) by stipulating that the phrase *with regard to shape transforms* the context in which (19) is used *into* one where shape is the only relevant issue. For our account of (19) it does not make much difference which line of explanation we choose. However, I believe that the solution to certain other semantical problems can be found only if we investigate not only the effect of the various aspects of context on the meanings of expressions used in those contexts, but also the mechanisms which *create*, or *modify*, contextual aspects. A proper understanding of these mechanisms seems essential to the

[1] An extensive discussion of such expressions can be found in Lakoff (1972). Lakoff calls such expressions 'hedges', a term I will adopt here too.

analysis of more extended pieces of discourse – such as told, or written stories.[1] Given that such understanding must eventually be reached in any case, an account of (19) along the lines of the second proposal may well ultimately be the more desirable. It seems, however, too early to pass judgement on this matter.

At any rate it is important to realize that contexts are made up of verbal and nonverbal elements alike. The same contextual aspect may on one occasion be manifest through the setting in which the utterance is made, while on another occasion its presence is signalled by a particular verbal expression. Exclusive preoccupation with shape, for example, can be evident to both speaker and audience either because they have been discussing shape and nothing but shape all along (think of a session about shape during a conference on industrial design); or because the previous sentence was *But let us now concentrate exclusively on shape*; or because the sentence itself contains the qualifying phrase *with regard to shape*. The three cases differ as regards the degree of permanence with which the feature in question is part of the context. In the first case, preoccupation with shape will last throughout the session; and a special verbal effort would be necessary to remove it; in the last the modification will be in force only during the evaluation of the particular phrase to which *with regard to shape* is attached; the second case is somewhere between the two. Indeed, without further information it is not possible to say whether the modification is valid just for the present sentence, for everything this particular speaker is going to say right now, or for the remainder of the entire discussion.

Another expression of the sort we have just been discussing is *in a sense*.[2] What is it to be clever in a sense? That depends on what are the various possible senses of the word *clever*. It will help to consider such related sentences as *Smith is clever in the sense that he is good at solving problems* or *Jones is clever in the sense of being quick-witted*. The expressions following *clever* in these two sentences have, again, the effect of transforming the context, viz. into one where *clever* is given a more specific sense. The truth value of the sentence should therefore be the same as it is in any of these contexts. The contexts in question are the same as those created by antecedent specifications like *Let us understand by 'clever': 'good at solving problems'*. Each such specification will single out a set of contexts in which *clever* is understood correspondingly. *x is clever in a sense* is then true if there is such a set of contexts such that *x is clever* is true in each of its members.

[1] Cf. Isard (1973). Others whom I know to have developed similar ideas are Thomas Ballmer of the Technische Universität, Berlin, and David Lumsden of University College, London.

[2] Cf. Lewis (1970:65).

But which are the acceptable specifications of a given noun or adjective? This is a question to which no definite answer can be given; for the notion of an acceptable specification of a given concept is itself subject to just that sort of vagueness with which this paper is concerned. Clearly not every logically possible definition is acceptable; for if this were so, then all statements of the form

(20) x is a...in a sense

would be true. But what is an acceptable specification can if necessary be stretched very far indeed. That is why it is so hard to establish that a particular sentence of the form (20) is false.

I want to conclude this discussion of hedges with a few remarks on the expression *to the extent that*. Let us consider Lakoff's example:

(21) To the extent that Austin is a linguist he is a good one

Once more I will leave questions concerning the ultimate logical form of the expression aside. It will be adequate for our present interests if we regard *to the extent that* as a two-place sentential operator which forms out of two formulas ϕ and ψ the compound formula

(22) to the extent that ϕ, ψ.

The semantical analysis of this connective brings into focus a problem connected with contextual disambiguation which I have so far failed to mention: to what extent does the sharpening of one predicate affect other predicates? Clearly the decisions concerning two different words cannot in all cases be independent. Sharpening of the noun *leg* will yield sharpening of the adjective *four-legged* as well. Yet there are many pairs of adjectives such that a sharpening or modification of one does not carry with it any perceptible semantic change in the other. This is true in particular of *linguist* and *good*. This is important for the following account of (21).

The truth conditions of (22) are essentially these: (22) is true (in its actual context of use c) if ψ is true in all contexts in which ϕ is true and which are as similar to c as is possible, given that they make ϕ true. In the case of (21) these contexts will be contexts in which we have modified the semantics for *linguist* in such a way that Austin is now definitely inside its extension, and have left the semantics otherwise as much the same as the modification of *linguist* permits. In particular *good* would, it seems to me, not be affected seriously by the modification. The truth of the main clause of (21) in such a context is to be understood in the usual manner.

It is interesting to compare (21) with the slightly more complicated

(23) To the extent that Austin and Russell are linguists, Austin is at least as good a linguist as Russell

This sentence will be true in c if in every maximally similar context c' in which *linguist* has been modified in such a way that both Austin and Russell are in its extension, it is true that Austin is at least as good a linguist as Russell. When is it true in c' that Austin is at least as good a linguist as Russell? This will be the case if the pair ⟨Austin, Russell⟩ belongs to the positive extension of $\geqslant$(good linguist) with respect to c', i.e. if the set of members of $\mathscr{S}_{c'}$ in which *Austin is a good linguist* is true includes the set of those members in which *Russell is a good linguist* is true. It is important that this account will give us the intuitively correct truth conditions for (23) only if the members of $\mathscr{S}_{c'}$, involve modifications of *good* but not of *linguist*.

It is clear from this brief discussion that a formal elaboration of such analyses within the framework provided by cgm's requires a great deal more structure on the set of contexts than I have given.

7.

I have claimed that vagueness is often reduced by context. This doctrine is void, however, unless it is accompanied by a concrete analysis of those contextual factors which contribute to such reduction of vagueness and of how they succeed in doing so. To provide such an analysis is a difficult task, the completion of which will perhaps forever elude us. Yet I feel I ought to say something on this topic, more, in fact, than I actually have to offer. But let me mention at least one contextual aspect which plays a central role in almost all cases where adjectives occur in attributive position. That aspect is the noun to which the adjective is attached. In a great many cases the noun alone determines, largely or wholly, how the adjective should, in the given context, be understood. Indeed, if we assume that the noun is the *only* factor, we are back with the first theory according to which adjective meanings are functions from noun-phrase meanings to noun-phrase meanings.

But of course the noun is not always the only determining factor. *Smith is a remarkable violinist* may be true when said in comment on his after-dinner performance with the hostess at the piano, and false when exclaimed at the end of Smith's recital in the Festival Hall – even if on the second occasion Smith played a bit better than on the first.

It would be desirable to give a general account of how the meaning of the noun determines that of the adjective that combines with it. Here I will mention just one aspect of this problem. One of the main purposes of the use of an adjective in attributive position is to contribute to the delineation of the class of objects that the complex noun-phrase of which it is part is designed to pick out – or, alternatively, to help determine the particular individual which is the intended referent of the description in which the adjective occurs. In order that the adjective can be of any use at all for these purposes, it should, in the presence of the noun in question, have an extension which, so to speak, cuts the extension of the noun in half – i.e. if we assume for the sake of this argument that both noun and adjective (in the presence of the noun) are sharp, and that $\bar{A}$ is the extension of the adjective, and $\bar{N}$ that of the noun, then both $\bar{N}\text{-}\bar{A}$ and $\bar{N} \cap \bar{A}$ should be substantial proportions of $\bar{N}$. Thus in order to be able to use the adjective profitably in combination with an unlimited number of nouns, we should let the noun determine the criteria and/or standards for the adjective in its presence in such a way that the above condition is in general fulfilled. (The proposal (8), p. 126, obviously meets this requirement.)

The distinctions between nouns and adjectives adumbrated in the previous section are of course far from absolute. *Four-legged*, for example, has virtually no extension gap – which is hardly surprising given the manner in which it is derived from the noun '*leg*'. And indeed it yields comparatives as infelicitous as those derived from most nouns. *This is more four-legged than that* would on most occasions sound positively non-sensical. *Blue*, though apparently not derived from a noun, also gives rise to rather strained comparatives. *This is bluer than that* is sometimes a meaningful statement, but would fail to be more often than not. So it seems that *heavy* and *four-legged* are really very far apart and that they will ultimately require analyses that are fundamentally different.

This brings us to a likely objection against the theory I have outlined. Does it not blur fundamental distinctions between different kinds of adjectives? Yes, undoubtedly it does. Still, I feel that what it reveals about adjectives in general is important. But this conviction should not bar the way to accounts that deal in detail with small provinces of the wide realm of all those concepts to which it claims to apply. It should be pointed out in this connection that the second theory itself can hardly be regarded as comprehending all adjectives. Is *alleged* a predicate, even in the most diluted sense? It seems not. Of course we can still maintain that in each particular context of use it behaves as a predicate, in so far as the accompanying (or tacitly understood) noun phrase determines to which objects in that context the adjective applies. But this is just a restatement of the first theory in slightly different terms. The original intuition which led to the second theory seems to be inapplicable to *alleged.* The same can be said to be true, to an almost equal degree, of adjectives such as *fake*, *skilful*, or *good.* Where precisely we should draw the boundaries of the class of adjectives to which the second theory applies I do not know. For example, does *skilful* belong to this class? Surely we must always ask 'skilful what?' before we can answer the question whether a certain thing or person is indeed *skilful*; this suggests that the theory is not applicable to the word *skilful.* Yet there appears to be some plausibility in the view that *having a good deal of skill* does function as a predicate – be it a highly ambiguous one as there are so many different skills. Here the question whether we face an expression that stands for a function from properties to properties or rather an ambiguous predicate which is disambiguated by accompanying expressions for properties has perhaps no definite answer. Both views appear to be equally plausible accounts of the same phenomenon. So it may be impossible to determine in a non-arbitrary manner how far the domain of our theory extends. But then it probably does not matter whether we can or not. This will certainly be unimportant once we have a complementary theory which deals specifically with such adjectives as *alleged, fake, skilful* and *good.* It is bad to be left with a semantic phenomenon that is explained by no theory; but it does no harm to have two distinct theories which give equally adequate, albeit different, accounts of those phenomena that fall within the province of both.

8.

To conclude, let me mention some of the questions which I should have liked to discuss and which I believe can be treated within the framework I have set up.

In the first place there are intransitive verbs. I have avoided them throughout, even though they too appear to be one-place predicates and to display a good deal of vagueness. In particular I have failed to give any account of what semantically differentiates verbs from adjectives, or, for that matter, from nouns. My excuse for this is that the proper understanding of these differences involves the consideration of tense, of the time spans during which a predicate is true of an object, and of similar issues

which seem to require for their formal elaboration a framework which incorporates a good deal of tense logic.

Secondly, I have given only the scantest attention to hedges. I think that my framework is basically suitable for their analysis, although more structure on the set of contexts will be needed than I have provided.

Thirdly, I have considered only the simplest kind of comparatives. Examples of comparatives which are considerably more difficult to treat, are

Jones is more intelligent than he is kind
This building is higher than that is long
Smith is much cleverer than Jones and
Smith is more cleverer than Jones than Jones is than Bill

(accepting this as English).
The last two sentences in particular, present problems of a rather different kind than those I have tackled in this article. Their analysis requires more mathematical structure than has been built into the models here considered. The difference between the formal framework needed there and the one I have presented is essentially that between metric and arbitrary topological spaces. These and other problems I hope to consider in some other paper.

REFERENCES

Bartsch, R. and Vennemann, T. (1972). *Semantic Structures.* Frankfurt: Athenäum.

Geach, P. (1956). 'Good and evil', *Analysis*, **17.** Also in P. R. Fort, ed. (1967). *Theories of Ethics.* Oxford: Oxford University Press.

Isard, S. (1973). 'Changing the context.' This volume.

Kolmogorov, L. (1970). *Foundations of the Theory of Measurement.* New York: Chelsea.

Lakoff, G. (1972). 'Hedges: a study in meaning criteria and the logic of fuzzy concepts', *Chicago Linguistic Society*, **8.**

Lewis, D. K. (1970). 'General semantics', *Synthèse*, **22.** Also in G. Harman and D. Davidson, eds. (1971). *Semantics of Natural Language.* Dordrecht: Reidel.

Montague, R. (1970). 'English as a formal language.' In *Linguaggi nella Societá e nella Technica.* Milan: Editione di Communitá.

—— (1974). *Formal Philosophy*, ed. with an introduction by R. H. Thomason. New Haven: Yale University Press.

Parsons, T. (1968). 'A semantics for English.' Unpublished.

Rescher, N. (1969). *Many-Valued Logic.* New York.

Van Fraassen, B. (1969). 'Presuppositions, supervaluations and free logic.' In K. Lambert, ed., *The Logical Way of Doing Things.* New Haven: Yale University Press.