ELSEVIER

COGNITION

Cognition 57 (1995) 129–191

Prototype theory and compositionality

Hans Kamp[a,*] Barbara Partee[b]

[a]*Institut für Maschinelle Sprachverarbeitung, Universität Stuttgart, Azenbergstrasse 12, D 70174 Stuttgart, Germany*

[b]*Department of Linguistics and Department of Philosophy, Box 37130, University of Massachusetts at Amherst, Amherst, MA 01003-7130, USA*

Received August 31, 1993, final version accepted December 6, 1994

Abstract

Osherson and Smith (1981, Cognition, *11*, 237–262) discuss a number of problems which arise for a prototype-based account of the meanings of simple and complex concepts. Assuming that concept combination in such a theory is to be analyzed in terms of fuzzy logic, they show that some complex concepts inevitably get assigned the wrong meanings. In the present paper we argue that many of the problems O&S discovered are due to difficulties that are intrinsic to fuzzy set theory, and that most of them disappear when fuzzy logic is replaced by supervaluation theory. However, even after this replacement one of O&S's central problems remains: the theory still predicts that the degree to which an object is an instance of, say, "striped apple" must be less than or equal to both the degree to which it is an instance of "striped" and the degree to which it is an instance of "apple", but this constraint conflicts with O&S's experimental results. The second part of the paper explores ways of solving this and related problems. This leads us to suggest a number of distinctions and principles concerning how prototypicality and other mechanisms interact and which seem important for semantics generally. Prominent among these are (i) the distinction between on the one hand the logical and semantic properties of concepts and on the other the linguistic that between concepts for which the extension is determined by their prototype and concepts for which extension and prototypicality are independent.

1. Introduction

In an important article, Osherson and Smith (1981) (henceforth; O&S) raise doubts about prototype theory as a theory of concepts. Their doubts

* Corresponding author; fax 711 121 1366.

SSDI 0010-0277(94)00659-8

arise from the concern that an adequate theory of concepts should eventually account for compound concepts and the mechanisms of conceptual combination as well as for simple concepts. Since prototype theory does not by itself contain any explicit theory of how complex concepts or judgments are built up from simpler ones, Osherson and Smith appeal to fuzzy-set theory (Zadeh, 1965) as the most likely means for extending prototype theory to deal with conceptual combination. O&S then succeed in clearly showing that fuzzy-set theory cannot support a compositional semantics whose input consists of prototype concepts. In their words, "It is possible that principles other than those provided by Zadeh can better serve prototype theory in accounting for conceptual combination but no suitable alternative has yet been suggested (so far as we know)."

The first aim of this paper is to point out some well-known defects of fuzzy-set theory which render it unsuitable in connection with *any* theory of vagueness, and to draw attention to an alternative technique for dealing with the compositional semantics of vague terms known as the method of *supervaluations* (Van Fraassen, 1969; Fine, 1975; Kamp, 1975). The superiority of the method of supervaluations over fuzzy-set theory for several of the kinds of cases discussed by Osherson and Smith has been generally appreciated within formal semantics for some time but is not well known outside of formal semantics and logic, so the aim of this part of the paper (section 4) is to fill that gap. The conclusion drawn from this part is that, at least for the range of cases to which the method of supervaluations gives appropriate results, the failures of prototype theory-*cum*-fuzzy-set theory say nothing directly about the adequacy of prototype theory, since they are readily predictable from the inadequacies of fuzzy-set theory alone.

Some of the problematic data discussed by Osherson and Smith can be resolved if fuzzy-set theory is replaced by supervaluation theory. In particular, the "logical" combinations such as *fruit which is or is not an apple* receive appropriate accounts with just that replacement. But not all the problems O&S discuss have such a simple solution, and when we explore the problem they illustrate with the example of *striped apple*, we find that we are involved with general issues of vagueness and context dependence that crucially affect the interpretation of the data adduced both for and against prototype theory and theories of vagueness more generally. For these issues we do not know of any tools or strategies available in the existing formal semantics repertoire, but we offer here the results of our preliminary struggles with them from a formal semantics perspective. While we do not attempt to settle either the theoretical or the empirical disputes concerning the adequacy of prototype theory as part of a theory of conceptual combination, we hope to contribute to the interdisciplinary investigation of vague concepts and their manipulation by arguing for the central relevance of a number of additional factors and distinctions that we believe call for deeper investigation.

Among the factors we discuss are the separation of questions of the semantics of English expressions from questions of the mechanism of conceptual combination; in section 3 we illustrate both the importance of the distinction and the difficulty of drawing it, while laying out our general views on the principle of compositionality and its application to theories of adjective–noun combinations. Following the discussion of fuzzy-set theory versus supervaluation theory in section 4, we return in section 5 to the *striped apple case*, which on first inspection appears to be a counterexample not only to the fuzzy logic treatment of the compositionality of prototypical concepts, but also to the supervaluation approach which we want to propose in its stead. We diagnose the case as crucially involving the dynamics of context dependence, and argue that once the linguistic and non-linguistic factors that affect the dynamic "recalibration" of predicates in context have places provided for them in an enriched framework, the supervaluation approach can survive.

One distinction which we find important but problematic is the distinction between two different possible construals of "characteristic functions" or *c*-functions in connection with prototype theory: a *c*-function as a measure of the degree to which an object falls in the extension of a given concept and a *c*-function as a measure of an object's degree of closeness to the prototype of that concept, or its "typicality". As we note in section 2, Osherson and Smith are concerned with versions of prototype theory in which the same function is meant to play both roles; we question the generality of such a view and note throughout the subsequent sections various points where we feel the conflation has serious consequences. Our conclusions concerning the possible existence of cases for which the same function does in fact play both roles are found in sections 5.8 and 6.1.

In the penultimate section 6 we open up some broader speculations on the variety of kinds of concepts and concept combinations, including concepts with and without prototypes, complex concepts which do have prototypes even though their component concepts lack them and vice versa, concepts whose prototypes do determine extension and ones which do not, and other related issues. We raise the question of whether prototype theory does in principle cover a natural domain of cases or whether it might turn out to be a not quite natural subtheory of a more comprehensive theory of vagueness.

In our conclusions, we agree with Osherson and Smith that compositionality is a fundamental requirement for any fully general theory of concepts, and we agree with their conclusions concerning the impossibility of providing such a theory on the basis of prototype theory together with fuzzy-set theory. But we end up suggesting that the question of whether a given version of prototype theory can or cannot be compatible with a compositional theory of conceptual combination is one whose resolution will require both theoretical and empirical research that goes well beyond what has been explored so far.

2. The Osherson and Smith challenge

How does a concept determine which things fall under it and which do not? As Osherson and Smith summarize it, prototype theory "construes membership in a concept's extension as graded, determined by similarity to the concept's 'best' exemplar (or by some other measure of central tendency)" (O&S, 1981, p. 35). Consider, for example, the concept *chair*. To determine whether a given object is a chair the object is to be compared to the *prototype* for *chair*. If the object is close enough to the prototype it will pass as a chair. If it is sufficiently different it will fail to qualify as one. For objects at some intermediate distance from the prototype it may be indeterminate whether they should count as chairs. Chairhood is a matter of degree, and the degree to which something is a chair is analyzed in terms of its similarity to the prototypical chair.

It should be noted right away that not all work on prototype theory supports this version of it. For instance, among the concepts that Rosch (1973) reports as having prototypicality effects are superordinate notions such as *bird* and *fish*. She states that a number of different psychological tests support the claim that, for example, robins are "more prototypical birds" than penguins. But there is no evidence that penguins do not count unequivocally as birds, or that subjects who rate penguins low on the typicality scale for *bird* would have any doubts on this point. Armstrong, Gleitman, and Gleitman (1983) draw attention to the membership/prototypicality distinction and point out the failure of Rosch and other contributors to be careful about making it. They observe some of the classic prototypicality effects even among such concepts as *odd number*, a concept which is certainly as sharp as any with respect to membership.

It seems therefore that a version of prototype theory intended to analyze degrees of membership for concepts in terms of prototypicality cannot be applicable to all concepts, including many concepts that would count as prototypical according to various of Rosch's criteria. In much of this paper we can ignore this point, but it is an issue that becomes important in several places later on, and which we discuss explicitly in sections 5.7, 5.8 and 6.5.

O&S explicitly acknowledge that the version of prototype theory they present for discussion should not necessarily be expected to apply to all concepts. However, as they point out, it should apply not only to "simple" concepts like *apple*, *red*, *stupid*, *pet*, and *fish*, but also to complex concepts such as *striped apple*, *apple which is not an apple*, *pet fish*, etc. They see the following problem for a prototype theory whose domain contains such complex concepts. Intuitively it would appear that we understand a concept such as *striped apple* by combining the component concepts, *striped* and *apple*, in some systematic fashion. It is a central task of a theory of concepts and concept formation to provide an account of how such combinations can be formed and understood.

Now suppose a concept is given as a complex consisting of a prototype and a function which measures the extent to which objects differ from that prototype. Thus *apple* will be characterized by a pair (p_a, d_a) and *striped* similarly by a pair (p_s, d_s). The composition procedure should synthesize these two into a pair (p_{sa}, d_{sa}) consisting of the prototype of *striped apple* and the associated distance function. Osherson and Smith argue that there is apparently no general procedure that will produce intuitively correct prototype–distance function pairs for arbitrary compounds from the pairs that characterize their components.

Strictly speaking, the version of prototype theory O&S discuss differs superficially from the one just described. (The differences are immaterial to what we have to say in this paper, but as we will be mostly discussing O&S's views, it is better to stick with the formulation they give.) In that version, a concept is a quadruple $\langle A, d, p, c\rangle$, where p is the concept's prototype; A, *the conceptual domain*, is the field of "readily envisionable objects (real or imagined)" which can be compared with that prototype; d is a distance metric on A, that is, a binary function which assigns to each pair of elements of A a numerical value indicating the extent to which they differ from each other; and c is a function which assigns to each element of A a number in the real interval [0,1], signifying the "degree" to which that element falls under the concept. This degree must be a monotone decreasing function of the distance to p, a condition expressed in (1) below.

(1) $(\forall x, y \in A)(d(x, p) \leq d(y, p) \rightarrow c(y) \leq c(x))$

O&S refrain from specifying any general functional that determines (for arbitrary concepts) the function c in terms of the function d and the prototype p.

As O&S (1981, p. 37) make clear, on the conception of prototype theory they are concerned with, the c-function is understood to measure *both* the degree to which an object falls in the extension of a given concept *and* its degree of closeness to the prototype of that concept. We already note, however, that these two notions cannot always be equated. A relatively atypical bird, for instance, a pelican, may yet be an unequivocal member of the set of birds. If Robbie is a pelican then c_{bird}(Robbie) ought to have the value 1 if xc stands for membership, but a value significantly less than 1 when c reflects prototypicality.[1] In general, therefore, one cannot make do with just one function to serve both purposes. We will use "c^e" to refer to degree of membership in the extension of a concept and "c^p" for degree of prototypicality. One important question, then, is if and how these two functions are connected, and in particular for which concepts, if any, they coincide. We would like to note in this connection that some of the counterexamples O&S present to the theory they criticize are problematic precisely because no distinction is being made between what we are calling c^e and c^p. However, since parallel examples can be found involving concepts

[1] A low typicality rating for *pelican* for the category *bird* was indeed reported by Malt and Smith (1982).

for which c^e and c^p (arguably) do coincide, many of the points raised by O&S are unaffected, and we can ignore this complication in this and the next two sections. In particular we will, until further notice, go along with O&S's representation of a concept as a quadruple $\langle A, d, p, c \rangle$, and we will use c without superscript whenever the potential discrepancy between c^e and c^p is not at issue.

Having defined concepts in this way, O&S then ask what general mechanism could determine the quadruples characterizing such compound concepts as *striped apple* and *apple which is not an apple* in terms of their component concepts. One assumption they make initially and which we will dispute is that any such mechanism must compute the c-function of the compound directly from the c-functions (and from the c-functions only) of the parts. We shall discuss this assumption in section 4. They then argue, for a number of different types of compound concepts, that their c-functions cannot be obtained from the component c-functions by the algorithms fuzzy logic has to offer.

One of O&S's arguments concerns the concept *apple which is not an apple*. Evidently this concept unequivocally excludes everything.[2] Thus its c-function should assign to every object a number indicating that the concept definitely does not apply to it. Now consider an object a for which $c_{apple}(a) = r$, where $0 < r < 1$. Then, according to fuzzy logic, a's degree of non-applehood, $c_{non\ apple}(a)$, will be $1 - r$. According to the most familiar versions of fuzzy logic the degree to which a satisfies the conjunctive concept *apple which is not an apple* is then the minimum of the c-values of the conjuncts, and thus greater than 0.[3] Clearly this is not the right result. O&S conclude that prototype theory can apparently not meet the test of compositionality. As we indicated in the Introduction, we accept the substance of this argument, but want to resist the conclusion that prototype theory cannot meet the test of compositionality. We argue that it is not prototype theory that is at fault here, but the particular logical theory of vague concepts that was combined with it in O&S's reconstruction. But this is not the only possible choice for the logic, and we will propose another partnership, involving not fuzzy logic but supervaluation theory. Before we address this issue, however, we wish to say first a number of things about compositionality.

3. Compositionality

Before we present our arguments for the inadequacy of fuzzy-set theory and the merits of supervaluation theory for the logic of vague concepts we

[2] We agree on this matter with O&S's judgment, which has been challenged by some writers. We discuss the basis for apparent intuitions to the contrary in section 4.3.

[3] The particular case where $r = .5$ is discussed in section 4.1.

will review some of the basic issues surrounding the principle of compositionality in semantics, illustrate these issues with a brief discussion of the semantics of adjective–noun combinations, and finally enter a caveat concerning the difficulty of distinguishing between questions about the compositional semantics of English and questions about principles of conceptual combination.

3.1. The principle of compositionality

The principle of compositionality, often referred to as Frege's principle, is the principle that the meaning of a whole is a function of the meanings of its parts. Let us state it in a slightly more complete form, and then point out three crucial parameters that must be filled in to make it precise:

(PC) The principle of compositionality: The meaning of a complex expression is a function of the meanings of its parts and of their syntactic mode of combination.

The principle as stated in this form is nearly uncontroversial, since some version of it seems essential to any account of the ability to understand the meaning of novel utterances on first hearing. But in order to make the principle precise, one must give a specification of at least the following:

(a) the nature of the meanings of the smallest parts – that is, a theory of lexical semantics;
(b) the relevant part–whole structure of each complex expression – that is, a theory of the semantically relevant level or levels of syntax;
(c) the "functions" in question – that is, a theory of what combinatorial semantic operations there are, and how the rules for combining meanings operate on lexical meanings and syntactic structure to produce the meaning of the whole – in short, a theory of compositional semantics.

A theory dealing with (c) presupposes a theory dealing with (b). In fact, these issues are often accounted for in tandem. One example of a joint account of both of them is the following: the syntax is a context-free grammar, and for each syntactic rule of the form (2) below, the grammar contains a corresponding semantic rule of the form (3):[4]

[4] The notation $\|A\|$ in (3) is to be read "The semantic value of the expression dominated by the node A." We can remain neutral at this stage about what sort of thing a semantic value should be taken to be : a model-theoretic construct, an expression in some language of thought, a concept, etc. The basic issues of compositionality are independent of many of these foundational differences.

(2) $A \rightarrow BC$

(3) $\|A\| = F_n(\|B\|, \|C\|)$

Thus on this conception, the grammar must specify for each syntactic combination rule a semantic interpretation rule including a specification of a particular function F_n that tells how the meanings of the parts in that syntactic construction are to be used in determining the interpretation of the resulting whole. (We will give examples of particular such functions below.)

There are several potential complications to the simple picture of compositionality suggested by this sketch. The first is that some theorists believe there are multiple levels of syntax that contribute to semantic interpretation; a second and related complication is that a number of theorists believe that semantic interpretation proceeds by way of a mapping from one or more syntactic representations to an intermediate level of representation which is then compositionally interpreted. (Such an intermediate level is what is often meant by "logical form".) A third is that some theorists would urge a "modest" compositionality, arguing that the meanings of complex expressions are not in fact always fully determined by the meanings of their parts. Idioms and other lexicalized complex expressions are generally regarded as semantically atomic, not to be interpreted by compositional rules; but whether a given expression is to be regarded as idiomatic or lexicalized may be controversial. ("Pet fish", a case discussed by O&S and a number of their critics, seems to be a prototypical controversial example!)[5]

3.2. *The semantics of adjective–noun combinations*

As a case study in compositionality, and one which is directly relevant to many of the examples discussed by O&S and their critics, we will review progress over the past several decades in the semantic analysis of adjective–common noun combinations.

3.2.1. *Semantic features*

The semantic features approach of Katz and Fodor (1963) can be summed up in Fig. 1. That is, meanings were conceived as bundles of features and the semantic interpretation of an adjective-plus-noun construction was taken to be the sum of the features of its parts. This view of meaning is hard to extend beyond single one-place predicates, and feature theories of meaning have typically had difficulty dealing with such phenomena as quantification.

[5] Our anonymous reviewer disagrees about "pet fish", and offers the intuition that "pet snake" and "pet raccoon", which are surely non-frozen expressions, clearly raise the specter to which O&S were alluding. The reviewer notes that the head primary principle discussed below in section 5.3 seems to deal well with these cases in any event.

$$\begin{array}{ccc} \text{Adj} & \text{N} & \text{Adj + N} \\ \begin{bmatrix} F_1 \\ F_2 \\ F_3 \end{bmatrix} & \begin{bmatrix} F_4 \\ F_5 \\ F_6 \end{bmatrix} & \begin{bmatrix} F_1 \\ F_2 \\ \vdots \\ F_6 \end{bmatrix} \end{array}$$

Fig. 1.

But even within the limited domain where it is readily applicable, the feature theory of Katz and Fodor leads to difficulties. What it says about adjective–noun semantics is equivalent to the intersection hypothesis which we discuss next, and will fail for the same kinds of examples as that hypothesis.

3.2.2. The intersection hypothesis

One simple and appealing hypothesis about adjective–noun semantics is that adjectives and nouns are both simple one-place predicates which denote sets and that their combination denotes the intersection of the two sets. We can illustrate this hypothesis as in (4):

(4) $\|\text{carnivorous}\| = \{x \mid \text{carnivorous}\ (x)\}$

$\|\text{mammal}\| = \{x \mid \text{mammal}\ (x)\}$

$\|\text{carnivorous mammal}\| = \{x \mid \text{carnivorous}\ (x)\ \&\ \text{mammal}\ (x)\}$

$= \|\text{carnivorous}\| \cap \|\text{mammal}\|$

The cases to be considered in the next two subsections show that the intersection hypothesis does not hold for adjective–noun combinations in general; the adjectives for which it does hold are often called "intersective adjectives".

O&S clearly intend their discussion of conceptual combination for adjective–noun pairs to be restricted to intersective adjectives, and we will initially restrict our discussion similarly while assessing the ability of fuzzy-set theory and supervaluation theory to extend the intersection (or conjunction) operation to deal with vague predicates. But first we must make clear why the intersection hypothesis fails as a uniform account of adjective–noun semantics.

3.2.3. Non-intersective adjectives

An adjective like *carnivorous* is intersective, in that (5) holds for any noun N:

(5) $\|\text{carnivorous}\ N\| = \|\text{carnivorous}\| \cap \|N\|$

But not all adjectives are intersective. *Skillful* is an instance of a non-intersective adjective. As Parsons (1968) and Clark (1970) pointed out in the late 1960s, the invalidity of arguments like (6) is sufficient to establish this. For if (5) were true with *skillful* substituted for *carnivorous*, then (6) should be valid. But clearly it is not:

(6) Mary is a skillful surgeon
Mary is a violinist
―――――――――――――――
Therefore Mary is a skillful violinist

Since *skillful* does obey the principle expressed in (7) and instantiated in (7.1), it is called a "subsective adjective"; some adjectives, as we will see in the next subsection, are not even subsective.

(7) *Subsective*: $\|\text{skillful } N\| \subseteq \|N\|$

(7.1) *Subsective*: $\|\text{skillful surgeon}\| \subseteq \|\text{surgeon}\|$

3.2.4. Non-subsective adjectives

Adjectives like *former*, *alleged*, *counterfeit* are neither intersective nor subsective:

(8) (a) $\|\text{former senator}\| \neq \|\text{former}\| \cap \|\text{senator}\|$
(b) $\|\text{former senator}\| \not\subseteq \|\text{senator}\|$

That is, not only does the set of former senators fail to be the intersection of the set of former things (whatever that might mean) with the set of senators; moreover, as (8b) asserts, it is not even true that the set of former senators is a subset of the set of senators. Among the non-subsective adjectives we might further distinguish the subclass of "private" adjectives, those for which an instance of the adjective + noun combination is never an instance of the noun alone. The adjectives *conterfeit* and *fake* are of this sort, at least if it is agreed that a fake gun is not a gun, while *alleged* is not, since an alleged gangster, for instance, may or may not be a gangster. For some non-subsective adjectives it is not completely clear whether they are privative or not, and even the case of *fake* is not uncontroversial. We are unsure about the case of *former*; the question hinges on whether someone who was once a senator, then out of office for while, but is now a senator again, is now both a senator and a former senator. But insofar as these questions of subclassification are just details of the lexical semantics of the items involved, represented perhaps by semantic features whose content can be spelled out with meaning postulates or the like, it is not an issue of theoretical importance here how finely we draw the subclassification scheme or whether there are unclear cases with respect to the properties of individual adjectives.

3.2.5. Adjectives as functions

Parsons, Montague, and others argued that the simplest rule for the interpretation of adjective–noun combinations which is general enough to subsume all of the cases considered so far involves interpreting adjectives as *functions* which map the semantic value of any noun with which the adjective combines onto the value of the adjective–noun combination; and from the evidence of the non-intersective and non-subsective cases they argued that the relevant semantic values must be *properties* rather than *sets*, that is, must be *intensions* rather than *extensions*. We review these notions briefly below before restricting our attention, as do O&S, to intersective adjectives which can in principle be treated extensionally.

If adjectives are interpreted as functions from properties to properties, then for example the interpretation of *former* is a function from the property of being a senator to the property of being a former senator.

A full explication of this idea requires a review of the notions of *intension* versus *extension*. To a first approximation, we can take the extension of a predicate to be a *set* and its intension to be a *property*; intuitively, the extension of the predicate *senator* at a time t in a possible world w is the set of things that have the property of being a senator in w at t. More generally, the extension of a predicate in a given state of affairs is, by definition, the set of all those things of which the predicate is true in that state of affairs. This set is a reflection of what the predicate means; for given the way things are, it is the meaning of the predicate which determines which things belong to the set and which do not. But the extension is also a reflection of the facts in the state of affairs or possible world; the meaning and the facts jointly determine what the extension happens to be. Two predicates may therefore differ in meaning and yet have the same extension. To extend an earlier example, we could imagine a situation in which the set of surgeons was identical to the set of violinists – a situation in which the predicates *surgeon* and *violinist* had the same extension. The difference in meaning between the two predicates is reflected in the fact that there are many other possible states of affairs in which the predicates have different extensions. This aspect of the meaning of the predicates is what semanticists are focusing on in assigning *properties* as the semantic values of predicates.

The property a predicate stands for not only determines its *actual* extension, which consists of all the things that have the property in actuality, but also the extensions which the predicate would have in other possible circumstances: in a given set of circumstances, or *possible world*, w, the extension of the predicate consists of precisely all those individuals that have the property in that world. Moreover, it is commonly (though not universally) accepted that, conversely, the property a given predicate stands for is completely determined by this "spectrum" of actual and possible extensions. In other words, the property is completely identified by the function which assigns to each possible world w the extension of the predicate in w. Carnap, who was the first to recognize the importance of these functions, called them *intensions*. (Not all contemporary theories of properties identify properties with intensions in this sense;[6] we can do so since nothing here depends on the differences among different property theories.)

If properties are identified with predicate intensions then we see how adjective meanings become functions from intensions to intensions. This last identification allows us to express most of the semantic distinctions between different adjectives that have played a role in semantic theory to date. The distinctions among the various subtypes we have looked at in the preceding subsections can be characterized in terms of further restrictions on the kinds of functions that are expressed by the different classes of adjectives. Formally, these restrictions can be expressed as meaning postulates (see Carnap, 1952; Dowty, Peters & Wall, 1981); informally, we can think of classifications like "subsective" and "intersective" as semantic features on adjectives like *skillful* and *carnivorous* respectively, cashed out as restrictions on the corresponding functions requiring them to obey the respective conditions (7) and (5) above.

An important part of the strategy of semanticists in the early Montague tradition was to "generalize to the hardest case", in this instance the non-subsective adjectives. The argument was that if you want a uniform type of semantic value for all adjectives, that can only be achieved by making all of them functions from intensions to intensions. From there, the intuitively simpler subclasses, such as the intersective adjectives and the non-intersective but subsective ones, the privative ones, etc., can be treated as expressing functions from intensions to intensions which obey various further restrictive conditions making them behave as if they were of simpler types.

Why could one not proceed the other way around, and treat the simplest adjectives as the general case, adding complications only when needed? The answer lies in the goal of uniform compositionality – the challenge of finding a type of semantic value for adjectives and one for nouns so that one can give a single compositional semantic rule corresponding to the syntactic construction of adjective plus noun. This goal can be achieved if all adjectives are treated as functions from intensions to intensions: it is possible, for instance, to treat an apparently extensional, intersective adjective like *carnivorous* as a function from intensions to intensions of a highly restricted subtype (one that in fact ignores everything about the intension of the input except the extension it assigns in the given state of affairs). But the situation is not symmetrical; it is not possible to take the adjectives like *carnivorous* as the general case, treat adjectives in general as simple predicates and take the interpretation of the adjective–noun rule to be predicate conjunction (set intersection), because there is no comparable way to treat *former* or *skillful* as a restricted subcase of that. So if one wants a uniform compositional semantics for adjective–noun combinations, this

[6] See Chierchia, Partee, and Turner (1989) for some recent discussion of this issue.

appears to have to be done by treating adjectives as functions from intensions to intensions and the adjective–noun combination rule as function application.

To illustrate these last points, consider again a hypothetical state of affairs in which the nouns *surgeon* and *violinist* had the same extension. In such a state of affairs, the Belgian surgeons would be the Belgian violinists, but the skillful surgeons would not necessarily be the skillful violinists. This difference illustrates what is meant by calling an adjective like *Belgian* extensional and one like *skillful* intensional: the former function is sensitive only to the extension of its argument noun; the latter is sensitive to additional information connected with the noun – information that is part of its intension but not of its extension.

Thus within a uniform theory of adjectives as functions from intensions to intensions, different subtypes of adjectives can be accommodated as specifiable subcases of the general case; and one can single out the intersective adjectives as a particular highly constrained subclass. There is also another approach one can take to the diversity of subtypes of adjectives within the general framework of formal semantics – an approach that has been advocated by Partee and Rooth (1983) for the diversity of verb subtypes and by Partee (1987) for handling the diversity of subtypes of NP meanings. On this approach, known as "type-shifting", we associate each syntactic category with a set of semantic types rather than a single uniform type, enter each lexical item in the lexicon in its simplest type, and articulate general principles for assigning additional (predictable) interpretations of more complex types to those expressions which can have them. In the case of adjectives this would mean that the intersective adjectives, which intuitively and formally can be argued to have as their "simplest type" an interpretation as simple one-place predicates, would indeed be interpreted that way in the lexicon, and type-shifting rules would assign to them additional interpretations as functions from intensions to intensions in those constructions where meanings of that type were required. (See Partee, 1987, and Partee and Rooth, 1983, for details in the case of NPs and verbs; we will not try to explore the details of a parallel approach for adjectives here.) While the type-shifting approach is formally less elegant and more complex than the uniform approach of generalizing to the hardest case, there are not only some empirical arguments in its favor (see the works just cited), but it has the important appeal for potential processing models and theories of acquisition of treating the intuitively simplest cases as in fact simplest, and invoking the higher types (which remain the most general types for the given syntactic categories) only when an expression cannot be given an interpretation using the simpler types.

In either case, whether one adopts the uniform approach or the type-shifting approach, we can still say that on the one hand, the most general and comprehensive type of semantic interpretation for adjectives is as functions from intensions to intensions, while on the other hand, the intersective adjectives, which are singled out for analysis by O&S (and in many other works by linguists and psychologists alike), do indeed form a natural subclass which can be formally characterized as such on either approach.

With respect to the intension/extension distinction which is crucial for the analysis of the most general case as we have described it here, we believe that while O&S do not explicitly deal with intensions or intensional constructions, it seems to us perfectly consistent with O&S's use of the term *concept* to say that the concept a predicate stands for is a mental representation of the property the predicate expresses, that is, of the predicate's intension. However, since O&S's discussion is restricted to the subclass of intersective adjectives – the one subclass for which the extension/intension distinction can be safely ignored – the compositionality issue that is directly relevant is compositionality of concept *extension*. We will follow O&S in likewise confining our attention henceforth to questions of concept extension only, since we believe there are enough interesting and important issues within the purely extensional domain, focusing on the subclass of adjectives whose combination with nouns can be reasonably analyzed as predicate conjunction, that is, as intersection of the set which constitutes the extension of the adjective and the set which constitutes the extension of the noun.

One important issue in this regard is the difficulty of determining just which adjectives are in fact extensional in this sense; in the next subsection we bring the important phenomenon of context sensitivity into the discussion and show how it complicates the assessment of which adjectives belong to which subclasses.

3.2.6. Context dependence

In section 3.2.3 above we indicated that the inference pattern (6) was a test of whether an adjective was intersective. However, an adjective can be intersective but context-dependent, and may then appear to fail the test of (6) simply by virtue of the influence of the noun on the context. Consider *tall*, for example. We all expect a *tall mountain* to be much taller than a *tall man*, and we will generally accept that someone may be a tall man but not a tall basketball player. Does this establish that tall is non-intersective? No, as Kamp (1975) suggested and Siegel (1976a,b) argued at length, "relative" adjectives like *tall*, *heavy*, and *old* are context-dependent as well as vague, with the most relevant aspect of context a comparison class which is often, but not exclusively, provided by the noun of the adjective–noun construction. To see that the relevant contextual cues need not be provided by the noun alone, consider the difference in the most likely standards of height suggested by (9a) and (9b), with the same noun *snowman* in each case:

(9) (a) My 2-year-old son built a really tall snowman yesterday.
    (b) The D.U. fraternity brothers built a really tall snowman last weekend.

Siegel argues that the truly non-intersective subsective adjectives like *skillful* occur with *as*-phrases as in *skillful as a surgeon*, whereas context-dependent intersective adjectives take *for*-phrases to indicate comparison class: *tall for a* 12-*year old.*

One might argue that some or all of these other supposedly non-intersective adjectives like *skillful* might also be better analyzed as context-dependent intersective adjectives, differing from adjectives like *tall* only in the nature and extent of the contextual effects. In fact the two authors of this paper tend to disagree about *skillful*, and we have been unable to find or construct any fully conclusive arguments for either side. But we agree in wishing to emphasize that it is both difficult and important to try to sort out the effects of context dependence on the interpretation of different sorts of adjectives and nouns, both alone and in combination. And we agree that there are almost certainly some adjectives which are best analyzed as context-dependent intersective ones (probably including *tall*) and almost certainly some adjectives which are genuinely non-intersective and need to be treated as functions as described in section 3.2.5 (probably including *skillful*, almost certainly including *former*), although there may be many borderline cases and there may be cases which involve homonymous or polysemous doublets (as suggested by Siegel, for example, *clever*).

Vagueness and context dependence are in principle independent properties, although they often co-occur. *Left* and *right* are context-dependent but not (very) vague, whereas nouns like *vegetable* and *bush* are vague but not (very) context-dependent. (Probably almost every predicate is both vague and context-dependent to some degree.) We note in passing that not all vague predicates can be reasonably analyzed in terms of prototypes: for color terms like *red*, the vagueness can be seen in terms of distance from a prototype ("true red") but for scalar adjectives like *tall*, *hot*, *heavy*, etc., there is no "central value" determining maximal tallness but rather an open ended standard of comparison. We return to this and related points in section 6.

3.2.7. A note on compounds

The remarks above about adjective–noun combinations are intended to apply to all cases of modifier–head constructions, including cases where a noun is converted into an adjective (presumably lexically) and used to modify another noun, as in *oak table*, *cardboard box*, *brass ring*. But they are not intended to apply to compounds, either of the noun–noun or adjective–noun variety. Compounds can generally be recognized by their heavier stress on the first word:

(10)	*Modifier–head* (*compositional*)	*Compound* (*idiomatic*)
	(a) black board	(d) bláck board
	(b) brick factory	(e) bríck factory
	(c) toy store	(f) tóy store

The contrast in the case of adjective–noun combinations, as in (10a) and (10d), is familiar. The similar contrast in noun–noun combinations is less familiar but perfectly analogous. *Brick* as an adjective means "made of brick" and is intersective; *toy* as an adjective means something like "a toy version of a . . ." and is arguably non-subsective, although this is debatable. (See the discussion of the interpretation of *toy train* and *stone lion* in sections 5.3 and 5.4.) In compounds, on the other hand, there is no general rule for predicting the interpretation of the combination, intersective or otherwise. A tóy store is a store that sells toys, a tóy box is a box that holds toys, etc. The grammar does specify that it is the second part of a compound that serves as the head: a bríck factory is a factory and not a brick. The interpretive rule for a compound AB may be viewed as an extreme case of context dependence, specifying approximately "a B which has the contextually salient relevant property involving A". This clearly puts much of the burden of the interpretation of a compound onto the hearer's knowledge of the world and of the context of the utterance. Formal semanticists in general therefore do not expect a fully compositional semantics for compounds but do expect a compositional semantics for modifier–head constructions; their reason is that a native speaker cannot generally interpret a novel compound on first hearing on the basis of her knowledge of the language alone, but can do so for a novel modifier-noun construction.

3.2.8. Summary

O&S remark that "grammatical constituency can often serve as a guide to conceptual constituency". Part of the point of this discussion of the semantics of adjective–noun constructions has been to show that the determination of the semantic rules corresponding to particular grammatical constructions is often a difficult and delicate matter. In the sequel we will focus, as do O&S, primarily on adjective–noun combinations containing adjectives which we regard as uncontroversially intersective, but it should be noted that many of the arguments that have been raised over the years in discussions of treatments of conjunction of vague concepts may in fact be best understood as arguments against treating certain English adjectives (or perhaps all of them) as intersective, that is, as simple predicates. O&S were careful in choosing their examples to involve only relatively clear cases of intersective adjectives, but the discussions that have followed their original paper have not all been equally uniform on this point and this issue has mostly not been explicitly addressed. In section 5 we will return to a closer examination of the effects of context dependence mentioned briefly in section 3.2.6 above and make some suggestions towards the development of better theories than are presently available of the dynamics of contextual effects on interpretation and reinterpretation of both simple and complex expressions in their linguistic as well as their non-linguistic context.

3.3. Semantic compositionality and conceptual combinations

As the last remark hinted, we believe that it is important but difficult to try to separate problems that arise in giving an adequate semantics for English from problems in the characterization of concepts and conceptual combination. It is a difficult task because we have to talk about concepts in English, and because we assume that there must be some close connection between English expressions and the concepts they express. But it is an important one, since, for instance, no theory of conjunction for vague concepts can hope to explicate our understanding of the complex concept expressed by *skillful surgeon* if the mode of semantic combination expressed by the grammatical construction in this case is not conjunction (intersection) but function–argument application with a necessarily intensional argument; and just such examples as *skillful surgeon* have been subject to competing analyses of just these sorts over the past decades.

In the next section, where we compare fuzzy-set theory with supervaluation theory, we will restrict our attention initially to expressions built up from simple nouns with overt *and*, *or* and *not*, on the premise that in these cases we can be fairly sure that it is indeed some version of logical conjunction, disjunction, or negation that is the appropriate mode of combination to be explicated. In subsequent sections we return to combinations of nouns and modifiers, concentrating on cases where our best current semantic analyses likewise involve these basic logical modes of combination, but always bearing in mind that apparent counterexamples may reflect misidentifications of the semantic roles of certain English expressions rather than problems in the "logic of vague concepts".

4. Fuzzy-set theory versus supervaluation theory

4.1. Inadequacies of fuzzy-set theory

The appeal of fuzzy-set theory is probably due in part to the fact that it is an attempt to deal explicitly with the important and difficult problem of vagueness. We all have strong intuitions that the concepts encoded by many natural-language predicates are vague; whether something is a chair, or is red, does not seem to be an all-or-none matter but a matter of degree; there may be some clear positive cases and some clear negative cases, but there are many unclear cases in between. Thus there is no sharp dividing line between the positive and the negative cases (and there are no sharp dividing lines either between the unclear cases and the clear ones.) Standard set theory and standard first-order logic abstract away from the problem of vagueness and are applicable only to those predicates that can be regarded as definitely true or definitely false of any given object.[7] Now suppose we want a logic that deals with vague predicates like *tall* and *red* and *chair*, and with conjunctions, disjunctions, and negations of such predicates. For the sake of argument, let us start by adopting one of the assumptions that prototype theory shares with fuzzy-set theory, namely that for each atomic predicate we have a fuzzy characteristic function that assigns to any object in the domain of discourse a real number in the interval [0,1].[8] Such a characteristic function clearly corresponds to c^e (degree of membership) rather than to c^p (degree of prototypicality) as that distinction was introduced in section 2, and in this subsection and the next a simple c is always to be understood as c^e.

How should we think about conjunction, disjunction, and negation of vague predicates (or intersection, union, and complementation of fuzzy sets)? Fuzzy-set theory has proposed the following answers (see, for example, Zadeh, 1965):

(11) Let c_A, c_B be the characteristic functions of sets A and B.

Then

(i) $c_{A \cup B}(x) = max(c_A(x), c_B(x))$

(ii) $c_{A \cap B}(x) = min(c_A(x), c_B(x))$

(iii) $c_{not\ A}(x) = 1 - c_A(x)$

But this immediately runs into conflict with familiar logical principles, as illustrated by O&S's examples *apple that is not an apple* and *fruit that either is or is not an apple*. The first example was already mentioned in section 2. Let us go over both examples for the special case of an object a such that $c_{apple}(a) = .5$. If $c_{apple}(a) = .5$, then $c_{not\ apple}(a) = 1 - .5 = .5$, $c_{apple \cup not\ apple}(a) = .5$, and $c_{apple \cap not\ apple}(a) = .5$. But *apple that is not an apple* is a self-contradictory predi-

[7] Techniques can be introduced to restrict the domains of predicates to appropriate sorts of objects, for example, "prime" to numbers, "red" to physical objects, so as not to be committed to the truth of the "number three either is red or is not red". But the issue of vagueness is not directly affected by these modifications, since vagueness typically involves cases where the object is of the right sort for the predicate to be true or false of it (although we could also expect to find vagueness as to whether a given object is even of the right sort for a given predicate).

[8] One aspect of this proposal that some might find uncongenial is that the characteristic functions for concepts assign real numbers to objects. One might reasonably wonder what justification there might be for claiming that $c_{apple}(a)$ equal, say, .628 rather than .629 or .627. This is a difficulty that attends nearly all applications of the theory of probability to the analysis of problems in cognition (including the notion of subjective probability itself.) Some of the things we will put forward in sections 4.2 and 5.5 are afflicted by it too. We recognize this difficulty as one that requires a careful response. To respond adequately, however, would take us too far away from the central issues of this paper.

cate and thus should have a characteristic function that always gives the value 0. Also, *fruit that either is or is not an apple* should have the same characteristic function as the simple *fruit*.[9]

Suppose, however, that a is clearly a fruit and thus that $c_{fruit}(a) = 1$; then intuitively a is also a clear instance of the predicate *fruit that is or is not an apple*. So the c-function for that predicate ought to assign a the value 1; but as we noted above, the c-value for a according to fuzzy-set theory in this case is .5.

These counterintuitive results are not a necessary outcome of any theory of vagueness but result from Zadeh's proposals (i) and (ii) for computing the characteristic functions of conjunctions and disjunctions from the characteristic functions of their parts. The fundamental error in those rules is that they take no account of the possibility that the predicates A and B might not be independent of each other. This applies equally to the alternative of Goguen (1969), who takes $c_{A \cap B}(x)$ to be $c_A(x) \cdot c_B(x)$. On this formulation *apple that is not an apple* assigns a c-value of .25 to a if $c_{\text{apple}}(\text{a}) = .5$. This is less than .5 but still not 0.[10]

Note that by either Zadeh's "minimum" rule or Goguen's multiplication rule, if $c_{apple}(a) = .5$, then the characteristic functions for *apple that is not an apple* and *apple that is an apple* will have identical values for the object a (both .5 for Zadeh, both .25 for Goguen). It should be emphasized that these counterintuitive results have nothing to do with prototype theory but represent a basic shortcoming of multivalued logics in general, of which fuzzy logic is a special case. The trouble comes from looking at the "fuzzy characteristic function" as assigning degrees of truth to predications of the sort "a is P" and thinking of these degrees of truth as forming a linear scale, while still attempting to define the truth values of conjunctions and disjunctions from the truth values of their components.

It is clear that there is no way that this conception can meet even such minimal requirements as the simultaneous satisfaction of the following:

(i) $\varphi \& \neg\varphi$ should always have truth value 0.
(ii) $\varphi \& \varphi$ should always have the same truth value as φ.

For in any case in which φ and $\neg\varphi$ receive the *same* truth value (.5 in the

[9] Maybe "this is either P or not P" should come out false when P is in the unclear range; but note that on the fuzzy set theory treatment .5 is the lowest value this statement can ever attain. See sections 4.4 and 6.2 for more discussion of cases with *or*.

[10] Goguen's rule is closer to standard probability theory than Zadeh's. But note that in probability theory, $p(A\&B) = p(A) \cdot p(B)$ holds only when A and B are independent; the general rule for $p(A\&B)$ is $p(A\&B) = p(A) \cdot p(B|A)$ where $p(B|A)$ is the probability of B given the occurrence of A. According to this rule, $p(A\&\neg A)$ is always 0, for whenever $p(A) \neq 0$, $P(\neg A|A) = 0$.

cases cited), any truth-functional definition of conjunction must treat $\varphi\&\varphi$ and $\varphi\&\neg\varphi$ just alike.[11]

4.2. Supervaluation theory

An alternative approach to the logic of vague predicates is the method of supervaluations.[12] Central to this approach is the notion of a two-valued partial model. A two-valued partial model $\mathcal{M}$ for a language containing simple predicates like *apple* will consist of a universe of discourse $U_{\mathcal{M}}$ (a non-empty set) and for each predicate P a partial function $P_{\mathcal{M}}$ which assigns the value 1 to some objects (those in the "positive extension" of P according to $\mathcal{M}$), 0 to some others (those in the "negative extension" of P according to $\mathcal{M}$), and is undefined for the remaining objects (those in P's "truth value gap" or "range of indeterminacy" in $\mathcal{M}$). Thus $apple_{\mathcal{M}}$ will assign 1 to all objects in $\mathcal{M}$ that are definitely apples, 0 to the objects that are definitely not apples, and be undefined for those objects (if any) which qualify neither clearly as apples nor clearly as non-apples.

The rules for determining the truth values of logically compound sentences formed with *and*, *or*, and *not* from the truth values of their constituents are the classical ones, with the proviso that if a constituent lacks a truth value, and the value of this constituent could make a difference to the truth value of the whole, then the whole lacks a truth value also (see Kleene, 1952). What motivates these rules is that the indeterminate cases are regarded as lacking a truth value, not as having some third, "intermediate", truth value: this is precisely the difference between *two*-valued *partial* logic and *three-valued* logic.

Nothing in what we have said so far makes any distinction among the sentences $\varphi\&\neg\varphi$, $\varphi \vee \neg\varphi$, and $\varphi\&\varphi$ in those cases where φ lacks a truth value: any compound sentence containing φ as a part will also lack a truth

[11] It was pointed out to us by Filem Novák (p.c. to B.H.P.) that there are some versions of fuzzy-set theory (such as that presented in Zadeh, 1978) on which there is not just a single formula for conjunction but potentially infinitely many, and one must take into account the content of the conjuncts in deciding which operation to use in a given case. The choices are bounded by "bold product" at the lower extreme, where $a * b = \max(0, a + b - 1)$, and minimum at the upper end. If one chooses bold product for the "and" of "apple and not an apple", the desired value of 0 can be obtained.

We should make it clear that when we refer to fuzzy-set theory, we, like O&S, are referring principally to its classic version in Zadeh (1965), the basic principles of which we have heard Zadeh himself continue to articulate well into the 1980s. The alternative just described should clearly be an unsatisfactory account by anybody's criteria, since it suggests that the meaning of "and" is itself vague and context-dependent, whereas on the supervaluation theory exactly the given range of possible outcomes is predicted, and furthermore the choice among them explained, using exactly the classical and unambiguous logical "and".

[12] The method is due to Van Fraassen (1969) and was first applied to the logic of vague predicates by Fine (1975) and Kamp (1975).

value. The supervaluation technique, about to be described, was motivated in large part to reflect the intuition that there is a difference among such cases: if φ lacks a truth value, so should $\varphi \& \varphi$, but $\varphi \& \neg\varphi$ should nevertheless be counted as definitely false and $\varphi \vee \neg\varphi$ as definitely true.[13] The way this is achieved is as follows. Call the method of assigning truth values to sentences in a partial model that we have developed so far an ordinary (partial) *valuation*. To extend this to a *supervaluation*, consider the set of possible *completions* of the given partial model, that is, the set of valuations which eliminate all truth value gaps by extending the positive and negative extension of each predicate so that they jointly exhaust the domain.

Starting from a partial model with $U_{\mathcal{M}}$ as described above, for instance, each completion would include a total function f on $U_{\mathcal{M}}$ which agrees with $apple_{\mathcal{M}}$ where $apple_{\mathcal{M}}$ is defined, and which assigns a 1 or a 0 to each of the objects for which $apple_{\mathcal{M}}$ is undefined.

Different completions $\mathcal{M}'$ will divide up the original range of indeterminacy of $\mathcal{M}$ in different ways. For some applications it is appropriate to consider all logically possible completions of a given partial model, but for use in conjunction with prototype theory it is more appropriate to consider as possible completions only those which divide up the range of interdeterminacy of $\mathcal{M}$ in a manner consistent with the prototypicality rankings associated with given predicates.

Suppose, for example, that a concept, say *apple*, is given as a quadruple as above, including a function c_{apple} which assigns ranging from 0 to 1 to objects in its domain; and suppose we start with a partial model $\mathcal{M}$ in which $\|apple\|^{+}_{\mathcal{M}}$, $\|apple\|^{-}_{\mathcal{M}}$ are based on c_{apple} as follows:

(12) **Definition:** Let

$$
\begin{array}{ll}
x \in \|apple\|^{+}_{\mathcal{M}} & \text{if } .95 \leq c_{apple}(x) \leq 1 \\
x \in \|apple\|^{-}_{\mathcal{M}} & \text{if } 0 \leq c_{apple}(x) \leq .05 \\
x \in U_{\mathcal{M}} - (\|apple\|^{+}_{\mathcal{M}} \cup \|apple\|^{-}_{\mathcal{M}}) & \text{if } .05 \leq c_{apple}(x) \leq .95
\end{array}
$$

The possible completions of $\mathcal{M}$ *consistent with* c will then be all those completions $\mathcal{M}' = \langle U_{\mathcal{M}'}, \| \;\|_{\mathcal{M}'} \rangle$ of $\mathcal{M}$ which satisfy the constraint

[13] The intuitions that $\varphi \& \neg\varphi$ is always false and that $\varphi \vee \neg\varphi$ is always true have not gone unchallenged, particularly in the literature on prototype theory and fuzzy logic. We will return to this issue in section 4.4. The advantages of the supervaluation technique over fuzzy logic do not depend on the full strength of these assumptions, and can be maintained even on the weaker and less controversial assumptions that there *can* be a *difference* in truth values among φ, $\varphi \& \psi$, and $\varphi \vee \psi$ in cases where both φ and ψ are right "in the middle" of the range of indeterminacy (".5 true" in fuzzy logic terms).

(13) If $a \in \|apple\|^{-}_{\mathcal{M}'}$ and $b \in \|apple\|^{+}_{\mathcal{M}'}$, then $c_{apple}(a) < c_{apple}(b)$.

In words, if a given completion puts one object into the positive extension of *apple* and a second object into the negative extension, then the first must be a "better apple" than the second. With this constraint, different completions would differ from each other simply in where they drew the cutoff line on the original ranking, and would correspond intuitively to all the legitimately possible ways of making an originally graded concept into a sharp all-or-none one.

Given the set of possible completions of a given partial model, we are finally ready to define the supervaluation itself. Let us use the metavariable $\mathcal{M}$ for partial models, $\mathcal{M}'$ to range over possible completions of $\mathcal{M}$, and $\mathcal{M}^*$ for what we call a *supermodel*, consisting of $\mathcal{M}$ and all its possible completions. A *supervaluation* based on $\mathcal{M}$ is then defined as follows:

(14) **Definition:** The truth value of a sentence φ with respect to $\mathcal{M}^*$ is:

- 1 if its truth value is 1 in all completions $\mathcal{M}'$ of $\mathcal{M}$;
- 0 if its truth value is 0 in all completions $\mathcal{M}'$;
- undefined otherwise (i.e., if 0 in some and 1 in others).

The effect of this definition can be illustrated with our original case of c_{apple}. *Suppose a* is some object such that $c_{apple}(a) = .5$, and let "joe" be the name of *a* in our object language. Then in the model $\mathcal{M}$ which reflects this situation the atomic sentence "apple(joe)" receives no truth value, nor do any of the compound sentences "apple(joe) & apple(joe)," "apple(joe) & ¬apple(joe)" "apple(joe) ∨ ¬apple(joe)." In each completion $\mathcal{M}'$ of $\mathcal{M}$, "apple(joe)" will receive a definite truth value, 0 in some and 1 in others; whichever truth value it receives, its negation will receive the opposite. Thus in *every* completion the sentence "apple(joe) & ¬apple(joe)" will come out false and "apple(joe) ∨ ¬apple(joe)" will come out true. We thus obtain the result that in the supervaluation, "apple(joe)" is still without a truth value, as is "apple(joe) & apple(joe)", while "apple(joe) & ¬apple(joe)" receives the truth value 0 and "apple(joe) ∨ ¬apple(joe)" receives the truth value 1. This result would not be possible on any approach which, like fuzzy logic, defines the truth value of compound sentences directly on the numerical values of their constituents.

Supervaluation theory does not by itself provide all the answers to the problems of compositionality for prototype theory, as we will show in the next section. But it does provide a sound logical framework in which prototype theory and classical logic can peacefully coexist. We do not have to abandon any of the standard intuitions about the behavior of negation, conjunction and disjunction in order to make sense of a multivalued function c_{apple}; all we have to do is stop thinking of its values as "intermediate truth values" and think of them instead as providing constraints on the possible completions of a two-valued partial model.[14]

We can further illustrate the advantages of working with families of partial models and supervaluations by considering a somewhat richer example. Consider the concepts *boy*, *man*, *child*, *adult*, and *male*. The first four of these concepts are vague, but not independent of one another: the more clearly an individual is a man, the more clearly he is not a boy, and likewise for adult and child. Moreover, the vague line between boyhood and manhood is correlated with the vague line between childhood and adulthood; to make the example more general, we will not assume the gradients are identical, but allow that a boy might be an adult, although a man cannot be a child.

In keeping with our view that the prototypicality rankings given by c-functions act as constraints on partial models, we assume that someone who has learned these concepts is in a position to observe the correlations among them, and hence knows, explicitly or implicitly, the following constraints on partial models and on their possible completions:[15]

(i) If $x \in \|child\|^+_{\mathcal{M}}$ and $x \in \|male\|^+_{\mathcal{M}}$, then $x \in \|boy\|^+_{\mathcal{M}}$.

(ii) If $x \in \|man\|^+_{\mathcal{M}}$, then $x \in \|adult\|^+_{\mathcal{M}}$ and $x \in \|male\|^+_{\mathcal{M}}$.

(iii) If $x \in \|child\|^+_{\mathcal{M}}$, then $x \in \|adult\|^-_{\mathcal{M}}$.

(iv) If $x \in \|boy\|^+_{\mathcal{M}}$, then $x \in \|man\|^-_{\mathcal{M}}$.

Now consider some individual, Bob, who is definitely male but whose age is in the range of indeterminacy on both the boy–man and child–adult scales. Let us assume the following prototypicality values:

(15) $c_{boy}(Bob) = .5 \quad c_{man}(Bob) = .5 \quad c_{male}(Bob) = 1$
$c_{child}(Bob) = .45 \quad c_{adult}(Bob) = .55$

Now no matter what particular rules a fuzzy logician proposes for the various logical connectives, the fact that in fuzzy logic these prototypicality

[14] In our example we considered only a single partial model $\mathcal{M}$ and made a somewhat arbitrary stipulation about where its range of interdeterminacy lay. It is equally natural to consider *families* of partial models in the first place, whose membership is constrained by the function c_{apple} in much the same way that the possible completions are constrained. As we discuss later, different choices in practice as to where to draw the line between positive extension, range of indeterminacy, and negative extension are heavily affected by context of use. This suggests model structures in which contexts are associated with different partial models, each with its own contextually determined indeterminacy range. See Kamp (1975) for discussion.

[15] If we had made the assumption that the vague line between boyhood and manhood was identical to that between childhood and adulthood, then constraints (i) and (ii) would be expressed as biconditionals. We make no empirical claims about the correctness of the assumptions employed in the example, which is provided just to illustrate the workings of the theory.

values are treated as intermediate truth values to be directly operated on by the rules entails that any two sentences which differ just by a substitution of *man* for *boy* must come out with identical truth values. This is not so for the supervaluation approach, where the prototypicality values are part of a system of constraints on permissible partial models and their completions. The following pairs of sentences differ in just this way; their truth values on the supervaluation approach[16] are given in parentheses:

(16) (a) If Bob is a child, then Bob is a boy. (True)
(b) If Bob is a child, then Bob is a man. (False)

(17) (a) If Bob is a boy, then Bob is an adult. (Indeterminate)
(b) If Bob is a man, then Bob is an adult. (True)

(18) (a) If Bob is a boy, then Bob is a man. (False)
(b) If Bob is a man, then Bob is a man. (True)

Take the first pair above, for instance. On the fuzzy-set theory approach, (16a) and (16b) would both be represented as (16′):

(16′) If .45, then .5

Whatever rule one proposes for fuzzy *if–then*, the value must be the same in each case – clearly an incorrect result. On the supervaluation approach, however, we consider not those particular values but the constraints imposed on possible completions, starting from a situation in which the truth values of both antecedent and consequent are indeterminate. Thus for any completion $\mathcal{M}'$ which satisfies the constraints (i–iv) and for which, in accordance with our assumption that Bob is definitely male, Bob $\in \|male\|_{\mathcal{M}'}$, it will be the case that if Bob $\in \|child\|_{\mathcal{M}'}$ then Bob $\in \|male\|_{\mathcal{M}'}$. In contrast, constraint (iv) guarantees that for any such completion $\mathcal{M}'$, if Bob $\in \|boy\|_{\mathcal{M}'}$ then Bob $\notin \|man\|_{\mathcal{M}'}$. It then follows that (16b) is false. The other examples work similarly.

4.3. Deriving new c-functions with supervaluation theory

The supervaluation account we have just presented delivers a *partial* truth definition: a sentence that is true in some completions of $\mathcal{M}$ but false in others will receive neither of the two truth values. So a simple sentence like "Bob is an adult" will come out without a truth value in the sort of model discussed in the previous section, in which Bob is in the range of indeter-

[16] We assume the classical logical treatment of the material conditional (true unless the antecedent is true and the consequent false) for *if–then*; while the material conditional may be inadequate for a full account of the semantics of English *if–then* sentences, it suffices for our purposes here, and a more sophisticated account would not affect the point at issue.

minacy for the predicate *adult*. This renders the supervaluation approach and the fuzzy-set theory approach incomparable in certain respects, since the supervaluation theory contains nothing to match the *c*-functions which the theory we sketched in section 4.1 posits as characterizations of the degrees to which things fit concepts. However, the general form of the supervaluation account as given here invites a modification that delivers some such characterization as well.

The intuitive idea is this: suppose that in $\mathcal{M}$ Bob and Alma are both in the truth value gap of the concept *adult*, but Bob is "closer" to the positive extension of *adult* than Alma, perhaps because he is a little older or a little more grown up or both. It is reasonable to suppose that this comparison is reflected by the set of completions in $\mathcal{M}^*$: more of these will have Bob in the extension of *adult* than Alma. More precisely: the set of completions in which Alma belongs to the extension of *adult* will be a proper subset of those in which Bob belongs to the extension.

This suggests that we might take the set of completions in which Bob belongs to the extension of *adult* as a measure of the "degree" to which he is an adult. More generally, the set of completions in which an object a belongs to the extension of a concept A indicates the degree to which a falls under A. Note that along these lines we arrive at a notion of degreehood which applies to simple and compound concepts alike. Just as the degree to which a is an A is indicated by the set of completions in which a belongs to the extension of A, so the degree to which a is an instance of a conjunctive concept A *and* B is indicated by the set in which it is in the extensions of both A and B, and similarly for other compound concepts.

To make this measure comparable to a *c*-function we must associate with the sets of completions appropriate elements of the interval [0,1]. Such an association should preserve the relative ranking of the sets by size; the larger the set, the higher the associated number should be. In particular the entire set should get the number 1 and the empty set the number 0. Furthermore, the numbers associated to a given set of completions and to its complement should sum up to 1, and more generally, if A is a subset of a set B of completions then the numbers associated with A and $B - A$ should sum to the number associated with B.[17]

Together these constraints imply that the function μ which assigns numbers from [0,1] to sets of completions is a [0,1]-*measure*.[18]

Note, however, that the constraints do *not* determine the function μ completely. Indeed, it is far from clear on what sorts of criteria a particular μ could or should be selected.

But let us ignore this difficulty for the moment and assume that with each supermodel $\mathcal{M}^*$ comes a measure function that assigns numbers in [0,1] to a

[17] These constraints are argued more extensively in Kamp (1975), from which our presentation of the supervaluation account largely derives.

[18] See any text on measue theory (e.g., Halmos, 1950).

sufficiently rich subfield of the set of completions of $\mathcal{M}^*$. Formally, we will represent such supermodels as triples $\langle \mathcal{M}, \mathcal{S}, \mu \rangle$, where $\mathcal{M}$ is, as before, the partial model that forms the core of the supermodel, $\mathcal{S}$ is a set of completions, and μ is the measure function.[19] A supermodel $\langle \mathcal{M}, \mathcal{S}, \mu \rangle$ determines, for each concept A, a fuzzy characteristic function μ_A defined by:

(19) for any $a \in U_{\mathcal{M}}$, $\mu_A(a) = \mu(\{\mathcal{M}' \in \mathcal{S}$: a is in the extension of A in $\mathcal{M}'\})$

How do μ-values compare with the numbers assigned by c-functions? To answer this question we must return to the distinction we drew in section 2, between c^e and c^p. μ-Values are unambiguously concerned with a graded notion of instantiation of a predicate, and the values fall strictly between 0 and 1 only for elements in the range of indeterminacy of a given predicate. Thus, only when a is in the truth value gap of A in $\mathcal{M}$ can μ_A assign it a value between 0 and 1; if a is in the positive extension of A then $\mu_A(a)$ is always 1 and if it is in the negative extension then $\mu_A(a)$ is always 0. This implies that μ_A may be identified with c^e_A but not with c^p_A (except of course where c^e and c^p coincide).

Identifying c^e_A with μ_A gives us a procedure for obtaining the c^e-function of complex concepts from the semantic properties of their components. For it is the application conditions of the component concepts A and B that determine the values that $\mu_{A\&B}$ assigns to any object a. By way of illustration, let us return to an example of section 4.2, and let us now make the explicit assumption that the c-functions defined there for *boy*, *man*, *child*, etc., are indeed c^e functions. Let us say that a supermodel $\langle \mathcal{M}, \mathcal{S}, \mu \rangle$ *agrees with* c^e iff the conditions in (20) below hold:

(20) (1) $\mathcal{S}$ is *consistent with* c^e, in the sense that for any P in the domain of c^e and $\mathcal{M}' \in \mathcal{S}$,
- (i) if $a \in \|P\|_{\mathcal{M}'}$, then $c^e_P(a) > 0$
- (ii) if $a \notin \|P\|_{\mathcal{M}'}$, then $c^e_P(a) < 1$
- (iii) for any a, $b \in U_{\mathcal{M}}$, a, b and P satisfy the constraint (13) of section 4.2, that is: if $a \in \|P\|_{\mathcal{M}'}$ and $b \notin \|P\|_{\mathcal{M}'}$, then $c^e_P(a) > c^e_P(b)$.

(2) For any P in the domain of c^e, any $a \in U_{\mathcal{M}}$,
- (i) $\{\mathcal{M}' \in \mathcal{S}: a \in \|P\|_{\mathcal{M}'}\}$ belongs to the domain of μ; and
- (ii) $\mu(\{\mathcal{M}' \in \mathcal{S}: a \in \|P\|_{\mathcal{M}'}\}) = c^e_P(a)$

Suppose that $\langle \mathcal{M}, \mathcal{S}, \mu \rangle$ is a supermodel that agrees with c, in the sense

[19] In general, we cannot insist that μ be defined on the set of all subsets of $\mathcal{S}$. We will assume, however, that the domain of μ is large enough for our purposes and that, in particular, μ is defined for all the arguments which we discuss in the text.

of (20). Then, for instance, (21) can be interpreted as the c-value assigned to Bob for the conjunctive concept *boy and adult*:

(21) $\mu(\{\mathcal{M}' \in \mathcal{S} : \text{Bob} \in \|boy\|_{\mathcal{M}'} \,\&\; \text{Bob} \in \|adult\|_{\mathcal{M}'}\})$

This, by the way, also illustrates a point we made in section 2, that the c-function of the compound need not be determined by the c-functions of the components alone. The supervaluation account is not forced into claiming that the degree to which an object a satisfies a conjunctive concept $(A\&B)$ is fully determined by the degrees to which it satisfies A and B. Indeed, the degree to which a satisfies $(A\&B)$ in $\mathcal{M}^*$ is, according to our present proposal:

(22) $\mu(\{\mathcal{M}' \in \mathcal{S} : a \in \|A\|_{\mathcal{M}'} \,\&\; a \in \|B\|_{\mathcal{M}'}\})$

or, equivalently:

(23) $\mu(\{\mathcal{M}' \in \mathcal{S} : a \in \|A\|_{\mathcal{M}'}\} \cap \{\mathcal{M}' \in \mathcal{S} : a \in \|B\|_{\mathcal{M}'}\})$

This number is *not* determined by the degree to which a satisfies A and the degree to which it satisfies B, that is, by $\mu(\{\mathcal{M}' \in \mathcal{S}: a \in \|A\|_{\mathcal{M}'}\})$ and $\mu(\{\mathcal{M}' \in \mathcal{S}: a \in \|B\|_{\mathcal{M}'}\})$. The degree to which a satisfies $(A\&B)$ *is* determined by the two *sets* $\{\mathcal{M}' \in \mathcal{S}: a \in \|A\|_{\mathcal{M}'}\}$ and $\{\mathcal{M}' \in \mathcal{S}: a \in \|B\|_{\mathcal{M}'}\}$; but the way in which it is determined depends crucially on how these two sets intersect, and this is information that cannot as a rule be recovered from the values which μ assigns to these sets.

4.4. Contradictions and tautologies

To close this section, we wish to remark on claims which are sometimes made in the literature to the effect that the laws of classical logic are invalid for vague concepts. For instance, it has been denied, particularly by advocates of fuzzy logic and explicitly by some critics of O&S, that (24) below is a contradiction and that (25) is a tautology:

(24) Bob is a man and not a man.

(25) This apple is either red or not red.

We believe that these claims stem from a dubious analysis of the significance of the indisputable fact that such sentences can indeed be used to express contingent propositions.

Consider the case of (24) first. When such a sentence is used, we believe that something roughly like the following occurs. The sentence is first given a literal interpretation with *man* interpreted identically on both occurrences, since that is the analysis provided by straightforward application of the compositional semantic rules. When that interpretation is perceived to be contradictory, we backtrack and look for a non-contradictory interpretation, since otherwise we have to take our interlocutor to be violating fundamental

Gricean conversational maxims. The simplest way to reinterpret the sentence as non-contradictory is to interpret the two occurrences of *man* differently, or to interpret each as if it were modified by something like "in some respects". Note that if we block that possibility by adding the same "respects" to each, as in (26), it becomes much harder to construe the sentence as non-contradictory:

(26) Bob is a man with respect to age and gender and not a man with respect to age and gender.

Note further that in those cases where we impose a construal that makes the sentence true, as it might be in (27), we would want to count the sentence as definitely true, and not just as ".5 true," the highest value fuzzy logic could assign to (24):

(27) Bob is a man (in such-and-such respects) and not a man (in such-and-such other respects).

We thus conclude that a combination of classical logic, Gricean principles, and cooperative reconstrual gives a coherent and plausible account of how we (re)interpret apparent contradictions as non-contradictory, whereas fuzzy logic, despite its claims to gain support from such examples, does not.[20]

The case of a tautology such as (25) is somewhat different. Here the maxim apparently violated is not one of truthfulness but one of informativeness; we have to find a reconstrual of the sentence under which it could possibly be false or lack a truth value, since an outright tautology is uninformative. If the apple in question is a borderline case, the sentence would lack a truth value if we gave it a simple valuation rather than a supervaluation; this seems intuitively to be what we tend to do in such a case, and we hence take the sentence to be asserting that the apple does not fall in the range of indeterminacy for the predicate *red*. (The effect of such sentences can sometimes be to "decree" a closing of a truth-value gap, as in "those who are not with us are against us." "Either we're friends or we're not: I want you to make it clear where we stand.")

Note here too the undesirable results of the fuzzy logic rule that would always yield .5 for a disjunction whose disjuncts both had value .5. Suppose we are classifying things as animal, vegetable or mineral and we encounter an object we consider definitely not mineral, but ".5 animal" and ".5 vegetable" (perhaps a sea anemone). In this case we would want to assign sentence (28) the value 1, not .5:

(28) This is animal or vegetable.

According to supervaluation theory, (28) would get the value 1 on the assumption that in this case the indeterminacy consists in just where the line

[20] For a similar view see Osherson and Smith (1982), Section 4.3.

between animals and vegetables is to be drawn, and that wherever we draw it, the object in question will fall in the positive extension of one or the other. For further discussion of disjunction, including other kinds of cases, see section 6.2.)

4.5. Conclusion

In this section we have considered how the basic logic of *and*, *or* and *not* should be modified or extended to deal with concepts that are vague rather than sharply all-or-none. While fuzzy logic is the most widely known attempt to deal with this important question, it has fundamental defects that have been appreciated by formal semanticists for some time. Supervaluation theory, while by no means the complete story for dealing with vague concepts,[21] is an established extension of classical logic which encompasses vague as well as sharp concepts, and which is proof against the logical defects of fuzzy logic we have discussed in this section. We have indicated how it can be combined with prototype theory so as to provide a compositional account of the meanings of compound concepts, while avoiding some of the problems O&S noticed. In other words, in relation to the data considered in this section, this version of prototype theory appears viable when combined with supervaluation theory in the way we have sketched, provided we construe the *c*-functions as specifying degrees of membership. It is important to stress, however, that in the compositional account of compound concepts we have given here, the *c*-functions play a more indirect role than O&S assume for them; they act not to provide intermediate truth values but to provide constraints on permissible completions of partial models.

We thus see that at least some of the problems raised by O&S can be traced directly to the defects of fuzzy-set theory and are not problems in principle for finding an account of conceptual combination for concepts characterized by prototypes.

5. Striped apples and context dependence

5.1. Supervaluation theory and the striped apple case

The supervaluation theory explains why the *c*-function for *apple that is not an apple* may assign 0 to an object a even though $c_{apple}(a) = .5$. But it does not get us out of the difficulties presented by another case which O&S discuss. They observe that the prototype of the concept *striped apple* is presumably quite different from both the prototype for *striped* and that for

[21] See, for instance, Dummett (1975), Wright (1975), Sanford (1976), Kamp (1981), and Pinkal (1984, 1985, 1989).

apple. Now let a be an object which is very close to the prototype of *striped apple*. Then the c-function for *striped apple* may assign it a value quite close to 1, and one that is higher than the value assigned to it by c_{apple}. O&S offer these data without distinguishing between degrees of membership and typicality. However, in the present case this is not crucial; whether c is taken to be c^e or c^p it seems plausible that there could be objects a which get a higher c-value for the compound than for either of its components. We will defer the problems having to do with c^p (where it differs from c^e) until later, and continue to focus on c^e.

The possibility O&S mention – that $c^e_{striped\ apple}(a) > max(c^e_{striped}(a), c^e_{apple}(a))$ – is one that neither fuzzy logic nor supervaluation theory allows for. For neither theory permits the value of a conjunction ever to be greater than those of its conjuncts. Should we conclude from this that prototype theory is indeed incompatible with the facts of conceptual combination? Not necessarily; we believe that O&S have indeed identified an important problem, or family of problems, well illustrated by this case, but it is our opinion that there are other factors at work here affecting the data – factors that are not well understood in any theoretical perspective that we are familiar with. We also believe that the issues raised by the *striped apple* case are considerably more general than the specific problems such cases pose for prototype theory. In this section we discuss those general issues of vagueness and context dependence which we believe play a crucial role in such examples. In subsection 5.2 we give evidence for the conclusion that the *striped apple* case does not reflect directly on the logic of conceptual combination, but depends at least as much on issues of context dependence in the compositional semantics of English. In subsection 5.3 we present some hypotheses about the relevant dynamics of context effects in the interpretation of combinations of vague predicates. In the final subsections we return to the implications of these considerations for prototype theory and mechanisms of conceptual combination.

5.2. Diagnosing the problem

As we noted in section 3.3, it is not always straightforward to distinguish problems about the semantics of English from problems in the characterization of concepts and conceptual combination. The *striped apple* case seems to us a good illustration of this difficulty. It is natural to assume, as O&S did, and as we did too prior to this work, that *striped apple* should be an instance of a conjunctive concept, since *striped* appears to be an ordinary intersective adjective, so that the combination should be semantically interpretable by the simple intersection (i.e., conjunction) rule given in section 3.2.2. On the assumption that the relevant semantics is simply conjunction, it is thus natural to look to the theory of the logic of vague concepts for an account of the facts of such combinations. And as we have seen, the logic of conjunction for vague concepts provided by supervaluation theory fares just as poorly as that of fuzzy-set theory in this case: neither allows for a conjunction ever to receive a higher c-value than either of its constituents.

But let us back up and re-examine the assumption that *striped apple* is a simple case of predicate conjunction. As noted briefly in section 3.2.6, the influence of context dependence on interpretation can produce effects which interact with the effects of a given semantic combination rule in such a way as to make it sometimes quite difficult to determine what semantic rule is at work in a given case. Here we will show that the effects may be even more drastic than was there suggested, and that the nature of the context-dependent effects is at least in part dependent on the mode of syntactic combination used to form the complex expression and not just on the nature of the constituent expressions. From this we can argue that the problem illustrated by the *striped apple* case is a problem for the semantics (or semantics-*cum*-pragmatics) of modifier-head constructions in English, and probably in natural languages quite generally, rather than a problem for the "logic of conjunction".

Before looking more closely into the *striped apple* case itself we will consider some other examples which clearly illustrate effects of context dependence on the meanings of modifier–head constructions. First consider examples (29a–c):

(29) (a) Sam is a giant and a midget.
(b) Sam is a giant midget.
(c) Sam is a midget giant.

It is our judgment and the judgment of others we have informally surveyed over the past several years that these three sentences are most naturally understood as conveying propositions with mutually distinct truth conditions, despite the fact that all three would appear to predicate of Sam a compound concept with the same pair of constituent parts. What makes this a particularly nice test case is that both predicates are capable of being used either as head nouns or as modifiers with quite symmetrical interpretations, so that the effects of the syntax on the interpretation can be neatly factored out from contributions of the individual lexical items.

The relevant observations about the interpretation of (29a–c) are as follows. In the case of (29a), with overt conjunction of the two predicates, the sentence is generally interpreted as contradictory, unless one can find grounds for imposing an interpretation that implicitly adds different "respects" to the two, for example, a mental giant and a physical midget, much as in the case of overt contradictions discussed in section 4.3. Note that both *midget* and *giant* are vague and context-dependent terms; one who counts as a midget on a college basketball team will probably be larger than one who counts as a giant on a basketball team of 10-year-olds. When the terms are directly conjoined as in (29a), it appears that the default case is for them to be interpreted relative to the same context, and it follows from their semantic content that whatever counts as a giant relative to a given context *ipso facto* does *not* count as a midget relative to that same context.

In (29b) and (29c), on the other hand, one predicate serves as head noun and the other as modifier, and the difference in interpretation is striking. Our informants agree that a giant midget must be an unusually large midget, and a midget giant an unusually small giant. That is, it is the predicate serving as head noun that is interpreted relative to the given external context (a boys' basketball team, a family of circus midgets, the fairy tale of Jack the giant-killer, or whatever), and the predicate serving as modifier appears to be "recalibrated" in such a way as to be able to make distinctions *within* the class of possible referents for the head noun. So whereas in (29a), *giant* and *midget* are normally construed as mutually exclusive categories, in both (29b) and (29c) the modifier–head construction seems to virtually force one to construe them as compatible if at all possible, apparently by adjusting the interpretation of the modifier in the light of the local context created by the head noun (see Keenan, 1974).

When such an adjustment is made to the interpretation of the modifier, there is then no obstacle to interpreting the semantics of the modifier-head construction as conjunction;[22] but without such an adjustment, a conjunction interpretation of the construction would lead to the false prediction that (29a–c) should all have the same interpretation. Combinations such as *tall tree* illustrate the same point; *tall* is a vague term whose interpretation is affected by both the linguistic and non-linguistic context as illustrated in section 3.2.6; but once the interpretation is specified (as something roughly like "at least d tall" for some degree of height d) the combination can be treated as simple conjunction. The moral we would draw from these observations is that one must exercise great caution in drawing inferences from the interpretation of adjective–noun combinations to conclusions about the logic of conceptual combination, since the interpretation of the constituents, insofar as it is partly context-dependent, may not stay constant when the constituents are combined.

Returning to the *striped apple* case, we can now suggest that it is at least possible that the evidence it appeared to provide against both fuzzy-set theory and supervaluation theory is rather to be explained as evidence for contextual readjustment of the adjective *striped* in the context of the head noun *apple*. We will present an extension of the supervaluation approach incorporating such contextual readjustment in section 5.5 below and discuss

[22] One should of course question whether the semantics of the combination is indeed conjunction in this case, or whether *midget* and *giant* as modifiers have a non-intersective reading that explicitly builds in relativity to the noun they modify. The *snowman* test of section 3.2.6 applies here as well and suggests that these modifiers, like *tall* and *short*, are indeed context-dependent intersective ones. For practical purposes the difference between the two analyses is in any case slight.

an alternative approach involving a derived prototype in section 5.8. First we will develop further, in sections 5.3 and 5.4, some of the implications of the kinds of context-dependent readjustment we have been suggesting. In the later subsections and in the two appendices we explore the possibilities for integrating these ideas with supervaluation theory in hopes of approaching a theory which might do justice to the full range of data.

5.3. The dynamics of context effects

We will extend the discussion of the previous subsection slightly to try to articulate further some of the kinds of principles that may govern the dynamics of context effects with various linguistic constructions. We do this partly to alert experimental researchers to the potential importance of such effects as a source of contamination of the data, and partly to draw attention to an inviting field of study which to our knowledge has received very little systematic exploration, but which we believe holds promise for casting light on both linguistic and non-linguistic cognitive principles involved in the effects of context on the ways in which vague language is understood and vague concepts are applied. Our own proposals in this section are to be taken as rather exploratory and tentative.

The two principles suggested by the examples of the preceding subsection are the following:

(PSE) Parallel structure effect: In a conjoined structure, each conjunct is interpreted in a parallel way relative to their common context.

(HPP) Head primacy principle: In a modifier–head structure, the head is interpreted relative to the context of the whole constituent, and the modifier is interpreted relative to the local context created from the former context by the interpretation of the head.[23]

Both of those principles involve sensitivity to the choice of linguistic structure; but there are other principles which seem to be quite general, possibly universal, and not specifically linguistic; these may either cooperate or compete with principles like the parallel structure effect and the head primacy principle. For example:

(NVP) Non-vacuity principle: In any given context, try to interpret any predicate so that both its positive and negative extension are non-empty.

In the *midget giant* example, for instance, the head primacy principle and the non-vacuity principle cooperate to produce the observed results: we first

[23] In the simplest cases, the effect of the interpretation of a head noun on a given context will be to restrict the local domain to the positive extension of the head in the given context.

interpret the head *giant* in the given context in such a way as to give *giant* both a positive and a negative extension in the domain of the context; then we interpret *midget* in such a way that it has both a positive and a negative extension within the positive extension of *giant*. This of course requires a very different "calibration" of c^e for *midget* than would be obtained in the global context, since *midget* and *giant* are incompatible relative to one and the same context.

In the *giant and midget* example, (29a), on the other hand, we find a conflict between the parallel structure effect, which will make the two predicates incompatible and their conjunction contradictory, and the non-vacuity principle, which bids us try to interpret the conjoined predicate as non-contradictory, perhaps leading us to search for "different respects", though this might in turn run counter to the parallel structure effect again.

In the *giant midget* and *midget giant* cases, it is the positive extension of the modifier that has to be stretched to satisfy the non-vacuity principle. In other cases, such as (30) below, the same principle leads us to shrink the positive extension and expand the negative extension, since otherwise the truth of (31) would make the modifier in (30) redundant:[24]

(30) This is a sharp knife

(31) Knives are sharp.

In still other cases, the non-vacuity principle seems to override the head primacy principle. Consider the phrase *stone lion*: is a stone lion a lion? With respect to the normal interpretation of this predicate the answer would seem to have to be "no": stone lions fail both scientific and everyday tests for lionhood, would never be counted in a census of the worldwide lion population, etc. Yet if we have to stick to the normal interpretation, there cannot be any stone lions; lions aren't made of stone, nor is there any way to stretch the predicate *stone* to apply truly to lion-flesh. But *stone lion* is also not just an idiom: any name of a material can be substituted for *stone*, familiar or novel (glass, chocolate, velveteen, . . .), and just about any concrete sortal noun can be substituted for *lion*. In this case it seems that the non-vacuity principle forces us to override the head primacy principle and reconstrue the head noun so that the modifier *will* have a positive extension within the positive extension of the head noun. This is not just a case of vagueness resolution, since it seems to force us to put into the positive extension of *lion* things which are clear negative cases in normal contexts. (It may also be the non-vacuity principle or some generalization of it that makes us so strongly inclined to reinterpret tautologies and contradictions as contingent statements, as discussed in section 4.4).

The kinds of considerations we have raised here about effects of context

[24] An exact formalization of the principle that the extension of the modifier is adapted so as to yield a meaningful (i.e., non vacuous) partition of the head can be found in Klein (1980).

on positive and negative extension – and we have barely scratched the surface – undoubtedly apply as well to judgments of typicality. We suspect, from informal sampling, that judgments about what is a typical fish will differ when the context is varied among restaurant dining, home aquariums, going to the beach, or sport fishing. And we are sure that even if robins are the prototypical bird, the prototypical white bird will not be an albino robin but something like a dove or a seagull: here too we recalibrate the head noun, going down the typicality scale for birds in order to come up with a large enough sample of white birds for there to be both more and less typical examples within it. In terms of c^P-values, one might recast the non-vacuity principle as saying that we want to be able to interpret any predicate, simple or complex, in such a way as to have the full range of c^P-values from 0 to 1 instantiable.

5.4. *Compositionality and stone lions*

The observations made above about the radical readjustment to the positive extension of *lion* needed for the *stone lion* case can be seen to create problems for any attempt to construct a systematic account of compositionality, not just for a semantics based on prototype theory. It would seem that in order to interpret *stone lion*, we do not just apply some semantic operator to the meanings of the parts but rather actually change the meaning of *lion* first.[25]

How does a language-user know how to do this? It would seem that part of knowing the meaning of a word should have to involve knowing how the basic meaning(s) could be stretched, shrunk, or otherwise revised in various ways when necessary; since the possible revisions are probably not finitely specifiable, such a conception of meaning would take us well beyond the normal conception of the lexicon as a finite list of finite specifications of idiosyncratic information about the particular lexical items.

We mention this not because we have a suitable theory of the lexicon to offer,[26] but because it suggests that some of O&S's data may point to problems for compositionality in general, and not just for prototype theory. We return to the *stone lion* case in section 5.6 after offering an explicit mechanism for handling the more central *striped apple* case in section 5.5 below.

5.5. *Context dynamics and supervaluation*

In section 5.1 we observed that the supervaluation account sketched in section 4 cannot deal with compounds like *striped apple*. Our subsequent

[25] Maybe not literally "first"; maybe there is backtracking here.

[26] But see Chien (1985) for a conception of public or shared meaning as a set of constraints and possible speaker's meanings on possible occasions of use.

discussion of adjective–noun combinations like *giant midget* has shown the reason for this: the supervaluation theory of section 4 does not allow for the "recalibrations" which, as we suggested in sections 5.2 and 5.3, are involved in determining the extensions of such compounds. If the interpretation of adjective–noun combinations does involve such "recalibrations", then evidently a compositional theory of predicates will have to allow for them. As the supervaluation theory of section 4 does not allow for recalibration it is inadequate to the challenge which such adjective–noun combinations present.

But this does not mean that the supervaluation account is to be simply rejected. Rather we should first see how it might be combined with an account of context-sensitive recalibrations of predicates. In the present section we give a sketch for adjectival recalibrations of the *striped apple* kind which we discussed in section 5.2. In 5.6 we briefly consider this same problem in connection with the interpretation mechanisms for combinations of the *stone lion* variety.

In order to bring about a satisfactory merger between the supervaluation approach and the recalibration mechanism of section 5.2 it will be helpful to recast the supervaluation framework into a slightly different form. In section 4 we introduced supermodels as triples $\langle \mathcal{M}, \mathcal{S}, \mu \rangle$ where $\mathcal{M}$ is a partial model, $\mathcal{S}$ a set of completions of $\mathcal{M}$ and μ a measure over some field of subsets of $\mathcal{S}$. Of the three components which constitute these models the measure μ is the one for which it seems most difficult to give an exhaustive justification. In fact, we argued for only two kinds of constraints on μ: first, that it have the properties of a measure (so that for $T \subseteq \mathcal{S}$, $\mu(\mathcal{S} \backslash T) = 1 - \mu(T)$, etc.) and second that μ must be consistent with some antecedently given degree function c (in the sense that for any predicate P and $a \in U_{\mathcal{M}}$, $c_P(a) = \mu(\{\mathcal{M}' \in \mathcal{S}: a \in \|P\|_{\mathcal{M}'}\})$). Perhaps there are other constraints on μ; but we are unaware of them, and for the remainder of this paper we assume that there are not.

This implies that given a core model $\mathcal{M}$, a degree function c and a set of completions $\mathcal{S}$, any μ that obeys the mentioned constraints will be as good as any other. Thus $\mathcal{M}$, c, and $\mathcal{S}$ determine a set of supermodels $\langle \mathcal{M}, \mathcal{S}, \mu \rangle$ where each μ is consistent with $\mathcal{M}$, c and $\mathcal{S}$ in the sense just explained.

In fact, given $\mathcal{M}$ and c there is a natural set of completions of $\mathcal{M}$, consisting of all those which are consistent with c in the sense of (20.1). If $\mathcal{S}$ is chosen in this way, each pair $\langle \mathcal{M}, c \rangle$ will determine a set of supermodels $\langle \mathcal{M}, \mathcal{S}, \mu \rangle$. We will refer to pairs $\langle \mathcal{M}, c \rangle$ as "presupermodels." More precisely, a *presupermodel for* a given language L is a pair $\langle \mathcal{M}, c \rangle$ where (i) $\mathcal{M}$ is a partial model for L and (ii) c is a function which assigns to each primitive predicate of L a "graded extension", that is, a function which assigns to each object in $U_{\mathcal{M}}$ a number in the interval [0,1].

Let us assume that the adjective A and the noun N are among the primitive predicates in L and that $\langle \mathcal{M}, c \rangle$ is a presupermodel for L. Our problem is to articulate how $\langle \mathcal{M}, c \rangle$ determines a semantic value for the complex predicate AN and to show how this semantic value contributes in its turn to the semantic values of the larger expressions in which AN occurs as a constituent (in particular to the truth conditions of sentences of L containing AN).

The determination of the semantic value of AN should be seen as a two-stage process, consisting of a recalibration of A in the context provided by N. What we have said about the first stage of this process suggests a formal reconstruction, but does not determine any formal procedure in all detail. So the procedure we are about to sketch does not follow from our informal discussion of section 5.2, but it seems to be about the simplest way of implementing that discussion. For the sake of simplicity let us assume that the predicate N is sharp, that is, that for all $a \in U_{\mathcal{M}}$ $\|c_N(a)$ is either 1 or 0. So N has a well-defined extension $\|N\|_{\mathcal{M}}$. For sharp N, the recalibration procedure we are proposing comes to this: we treat the best cases of A within $\|N\|_{\mathcal{M}}$ as definitely within the positive extension of A/N ("A relative to N") and, similarly, the worst cases of A within $\|N\|_{\mathcal{M}}$ are treated as definitely in the negative extension of the recalibrated A/N. The intermediate cases are adjusted proportionally. Thus we arrive for the recalibrated predicate A/N at a c-function defined by

$$c_{A/N}(a) = \frac{c_A(a) - c^-_{A/N}}{c^+_{A/N} - c^-_{A/N}}$$

where

$$c^+_{A/N} = \sup\{c_A(a) \colon a \in \|N\|_{\mathcal{M}}\}$$

and

$$c^-_{A/N} = \inf\{c_A(a) \colon a \in \|N\|_{\mathcal{M}}\}$$

We now make use of supervaluation theory to define the degree to which any object a satisfies the conjunction AN. However, the supermodel to be used in this definition is *not* the one based on the given presupermodel $\langle \mathcal{M}, c\rangle$ but on a presupermodel $\langle \mathcal{M}, c'\rangle$ where c' incorporates the recalibration of A in the context of N. That is, c'_A equals $c_{a/N}$, whereas c'_N equals c_N. Such a presupermodel will determine a set $\mathcal{S}'$ consisting of all completions of $\mathcal{M}$ that are consistent with c'. Moreover $\langle \mathcal{M}, c'\rangle$ determines a family of supermodels $\langle \mathcal{M}, \mathcal{S}', \mu'\rangle$ where μ' is a measure over $\mathcal{S}'$ consistent with c'. Although μ' need not, as we saw, be fully determined by c', we can nevertheless ascertain for any $a \in U_{\mathcal{M}}$ that

(i) if $a \not\in \|N\|_{\mathcal{M}}$, then $\mu'(\{\mathcal{M}' \in \mathcal{S}' \colon a \in \|A\|_{\mathcal{M}'}$ and $a \in \|N\|_{\mathcal{M}'}\}) = 0$
(ii) if $a \in \|N\|_{\mathcal{M}}$, then $\mu'(\{\mathcal{M}' \in \mathcal{S}' \colon a \in \|A\|_{\mathcal{M}'}$ and $a \in \|N\|_{\mathcal{M}'}\}) = c_{A/N}(a)$.

For if $a \not\in \|N\|_{\mathcal{M}}$ then $c'_N(a) = 0$ and so there is no $\mathcal{M}'$ in $\mathcal{S}'$ such that

$a \in \|N\|_{\mathcal{M}'}$. So there is a fortiori no $\mathcal{M}'$ such that $a \in \|A\|_{\mathcal{M}'}$ and $a \in \|N\|_{\mathcal{M}'}$ and consequently $\mu(\{\mathcal{M}' \in \mathcal{S}': a \in \|A\|_{\mathcal{M}'} \text{ and } a \in \|N\|_{\mathcal{M}'}\}) = 0$. If, on the other hand, $a \in \|N\|_{\mathcal{M}}$, then $c'_N(a) = 1$, and so $a \in \|N\|_{\mathcal{M}'}$ for all $\mathcal{M}' \in \mathcal{S}'$. So

$$\{\mathcal{M}' \in \mathcal{S}': a \in \|A\|_{\mathcal{M}'} \text{ and } a \in \|N\|_{\mathcal{M}'}\} = \{\mathcal{M}' \in \mathcal{S}': a \in \|A\|_{\mathcal{M}'}\}$$

and so

$$\begin{aligned}&\mu'(\{\mathcal{M}' \in \mathcal{S}': a \in \|A\|_{\mathcal{M}'} \text{ and } a \in \|N\|_{\mathcal{M}'}\})\\ &\quad = \mu'(\{\mathcal{M}' \in \mathcal{S}': a \in \|A\|_{\mathcal{M}'}\}) = c'_A(a) = c_{A/N}(a)\end{aligned}$$

Note that these identities are independent not only of the particular choice of μ' but also of the complete specification of c'. (Recall that we only said of c' that c'_A was to coincide with $c_{A/N}$ and c'_N with c_N; we left open how c' should relate to c for the other predicates of L.) However, if we want to use supermodels associated with $\langle \mathcal{M}, c' \rangle$ for the evaluation of sentences in which AN co-occurs with such other predicates, this question is no longer immaterial. What model will be appropriate for the evaluation of the sentence will depend on whether these other predicates are to be understood as recalibrated or not. Evidently this is a question that cannot be answered in abstraction – it will depend on how these other predicates occur within the sentence – for instance, some of them may occur in prenominal position and thus may be, like the predicate A, in need of recalibration, but quite possibly with respect to different contexts. In such a case the supermodel suitable for evaluation of the sentence may stand in a quite complicated relation to the original presupermodel $\langle \mathcal{M}, c \rangle$. This indicates that in general the supermodels involved in the valuation of a given sentence s must be those which are determined by a presupermodel $\langle \mathcal{M}, c' \rangle$ where c' differs from the original c in such a way as to account for all the recalibrations which concept words in s require in view of the particular ways in which they occur in s. To develop a semantical theory in which this observation is implemented systematically is no trivial matter, and it is not one we claim to have solved in detail. However, this is a task which transcends the purpose of the present paper, and we will not pursue it any further. We hope, however, to have indicated how the process of recalibration can be reconstructed in formal terms compatible with the supervaluation method, and how within such an extended formal setting the problem presented by combinations such as *striped apple* finds its resolution.

We have only considered the case where N is sharp, which is reasonable enough for *striped apple*, but not for combinations like *midget giant*. In fact, we do not have very clear intuitions about what recalibration should be like in situations in which N fails to be sharp. It is possible to formulate certain analogues of the procedure described above – for example, by carrying the procedure through for each of the possible sharp extensions of N that are consistent with c_N, and then integrating these results into a single supermodel. But we will not undertake such an analysis here, given the tentative nature of the particulars of this proposal. Some of the problems facing our proposals for cases like *midget giant* or for those involving open-ended scalar adjectives like *tall* are discussed in Appendix I.

5.6. Supervaluations and stone lions

In the previous section we have seen how one could combine context-dependent "recalibration" mechanisms with the supervaluation theory to give a plausible account of adjective–noun combinations such as *striped apple*. As we earlier suggested in sections 5.3 and 5.4, a combination like *stone lion* requires a more drastic form of reinterpretation, involving mechanisms that may well fall outside the purview of any kind of compositional semantics. Here we have a combination AB in which we are not just recalibrating the extension of the modifier A in the context of the head B but shifting the interpretation of the head noun B substantially (roughly from animals to artifacts, in this case). What seems do drive the reinterpretation in a case like *stone lion* is the fact that on the original or primary meaning of the head noun *lion* there is no way to construe the modifier as having a positive extension that overlaps the positive extension of the head noun at all. So literal interpretation, even enriched with the contextual and supervaluation mechanisms we have introduced, results in a conflict with the non-vacuity principle – a situation which can presumably trigger the search for a reconstrual that will permit a non-vacuous interpretation. Whether the reinterpretation of *lion* as "lion representation" is a metaphorical one or a secondary literal meaning is a question we do not know how to settle and do not need to; in either case what we have is the activation of a secondary meaning in a context where the primary meaning fails to give a coherent non-vacuous interpretation.

What we cannot expect the regular compositional semantic mechanisms to tell us is what the most salient available secondary meaning will be if the primary meaning fails. This will clearly be different for different concepts and in different contexts, and will frequently involve all sorts of real-world knowledge in addition to linguistic knowledge. A complete theory of understanding including theories of ambiguity resolution, metaphor interpretation, and lexical shifts, etc., has to involve such issues, and we recognize that there are debates about how far one can separate such issues from what we are taking to be the domain of compositional semantics. But we are inclined to believe that, whether part of semantics or not, the strategies or principles involved in seeking or constructing a secondary meaning to use if the primary meaning "fails" or leads to anomaly are different in kind from the sorts of principles and mechanisms articulated in section 4 and in section 5.5, and we therefore make no attempt here to formalize this sort of reinterpretation.

Once the meaning of the head noun has been shifted, however, the semantics of the combination is straightforward predicate conjunction. *Stone* as a modifier means "made of stone", and the shifted meaning of *lion* we take to be something like "lion representation" or "object in the shape of a lion." For these predicates, simple conjunction, or set intersection, is unproblematic, and we do not need to invoke any of the further mechanisms of supervaluation theory or contextual readjustment described in section 5.5.

5.7. *Interlude: no compositionality for prototypes*

The account of conceptual combination we have developed in this and the previous section has not made any direct mention of prototypes. In particular, it explains how the membership function of a complex concept is determined by its constituents without ever referring to the complex concept's prototype. In many cases this is just as well. For there are numerous complex concepts such that:

(a) there is no systematic connection between membership and prototypicality, so that an account of how the prototype is determined would be of no help in accounting for membership, and/or
(b) the factors which determine the prototypes appear to lack the degree of systematicity which a compositional account of prototype choice presupposes.

An example which illustrates both these points clearly is the concept *male nurse*. Both *male* and *nurse* come fairly close to being sharp concepts, and we can assume that they are sharp without distorting the point that concerns us. Moreover, *male* seems to be a pretty good example of an intersective adjective.

Thus the extension of *male nurse* is determined in the simplest way possible, namely as the standard logical conjunction of the constituent concepts *male* and *nurse*. (The function $c^{e}_{male\ nurse}$ is the function which only takes the values 0 and 1 and which maps an object a onto 1 iff it is both male and a nurse, that is, if $c^{e}_{male}(a) = c^{e}_{nurse}(a) = 1$, and maps it onto 0 otherwise.)

This establishes that no prototype for *male nurse* is needed to fix its extension. But it is not just that no prototype is required for this; it is hard to see how the prototype for *male nurse* could play any significant role in determining membership. This does not mean that speakers do not have prototypes for such a concept. However, different speakers' prototypes for *male nurse* will undoubtedly vary considerably according to their individual experiences (i.e., what examples of male nurses they have encountered in real life, or seen on TV, etc.) and will probably bear no systematic relation to their prototypes for *male* and for *nurse*; and in any case, whatever prototypes speakers may have for *male nurse* will have little bearing on the question of what is to count as a male nurse.

Since, as we saw, the extension of *male nurse* is fixed by principles that do not involve its prototype, the apparently unsystematic ways in which speakers get their prototypes for it and the fact that those prototypes may vary widely from one speaker to the next pose no real threat to the principle of compositionality. However, with concepts for which there is a logical connection between membership and resemblance to prototype the situation is different. This situation is discussed in the next section, in which we return once more to the concept *striped apple*.

5.8. Membership and prototype resemblance

As early as section 2, we distinguished between two c-functions: c^e, which specifies the degree of membership, and c^p, which specifies degree of resemblance to the prototype. We asserted then, without giving arguments, that we believe that these two functions sometimes do and sometimes do not coincide. In the last subsection we described the case of *male nurse* – a clear example of a concept for which the two functions do not coincide. In this case c^e is the characteristic function of the set of people who are both nurses and male – a function which, at least in first approximation, may be assumed to take only the values 0 and 1; while c^p is a function which we saw no way to define, which, unlike c^e, may vary from speaker to speaker, and which, as resemblance to the prototypical male nurse (whatever that may be) is surely a matter of degree, assigns many individuals values other than 0 or 1.

We believe, however, that there are also concepts for which two functions do coincide – concepts for which degree of membership *is* a matter of resemblance to the prototype. We think this is so, for instance, for color concepts, such as *red*, for many artifact concepts, such as, say, *chair*, and for such character or personality concepts as, for example, *bully*: for all these concepts there is a genuine sense in which falling under the concept is a matter of resemblance to the cases which speakers conceive as prototypical or central.

We think, moreover, that there is such a logical connection between prototype and membership not only for certain simple concepts, but also for a range of complex ones. For complex concepts for which there is such a connection, indeterminacy of the prototype – such as found in the case of *male nurse* – would seem to threaten the possibility of giving a compositional account of their extensions. By contraposition, if our grasp of complex concepts whose extensions are determined by their prototypes *is* compositional, we should expect their prototypes to depend on their constituent concepts in a systematic way. To see whether this is indeed so let us look once more at *striped apple* – a concept of which we think it plausible that its membership is connected with prototypicality.[27]

[27] This may well be controversial. But even if it is, the general point in this section stands as long as there are *some* adjective–noun combinations which on the one hand are interpreted via recalibration of the adjective, and on the other hand have extensions that are determined by prototype resemblance.

Inasmuch as it is true that speakers associate a prototype with *striped apple* it is likely, as in the case of *male nurse*, that these prototypes will vary – depending, say, on what kinds of apples are most common in the speaker's environment.[28] But if, as we have assumed,

(i) the prototype of *striped apple* is connected with its extension, and
(ii) the extension of *striped apple* is determined by recalibration,

then we should expect that the choice of prototype is at the same time constrained in a way that, say, the prototype for *male nurse* is not. To see that this should be so we argue as follows. First, (i) entails the following principle:

(32) The prototype for C must be as good an instance of C as there can be.

For concepts whose prototypes determine membership via resemblance (32) may be considered analytic. For nothing will resemble the prototype more than it resembles itself. So if degree of membership is degree of resemblance to the prototype it follows that the prototype must have as high a degree of membership in the concept as any object. In particular, the prototypical striped apple, P_{SA}, will have as high a degree of membership as any other individual in the concept *striped apple*. And if we assume, as we did in the reconstruction of recalibration we offered in section 5.5, that *apple* is sharp, it follows that P_{SA} is as good a case of *striped* – that is, it gets as high a value under $c_{striped}$ – as any apple. In other words, for any a in the relevant domain we have:

(33) $c_{SA}(a) \leq c_{SA}(P_{SA})$

where c_{SA} is defined as in section 5.5. In view of the definition this is equivalent to the requirement that

(34) $c_{SA}(P_{SA}) = 1$

The upshot of this is that for a complex concept which satisfies conditions (i) and (ii) above an account of its application conditions could be given in two distinct ways. The first account would follow the proposal sketched in

[28] In fact the example of a striped apple shown in the illustration in O&S is an example of a striped apple *picture* with straight black and white stripes that look more like stripes on a shirt or on a zebra than like anything that could occur in nature on a real apple, where stripes or streaks virtually always follow longitudinal lines, are most noticeable near the stem end, and are irregular and shaded. The context of psychology experiments is indeed often one of line drawings or cardboard cutouts, so that the mechanisms that govern speakers' (apparently effortless) reinterpretation of *lion* or *apple* as *lion representation* or *apple representation* might profitably be explicitly investigated and their interaction with the factors addressed here studied. Such investigations might also help to shed some light on the apparent context dependence of judgments about whether a fake gun is a gun.

section 5.5, namely contextual recalibration of the adjective followed by the normal supervaluation treatment of conjunction. The second would spell out how, and under what constraints, the prototype is selected, and then define c_{SA} in terms of conceptual distance from that prototype. Although we are missing details from the second story it seems nonetheless safe to conclude, first, that the assumptions (i) and (ii) we have made about *striped apple* are in principle compatible; and second, that under those assumptions the membership function *can* be construed as prototype resemblance, but can also be obtained from the constituent concepts by the first procedure, in which prototype selection plays no active part.[29]

Although the assumptions (i) and (ii) we have made about *striped apple* strike us as very plausible, we are not convinced of them beyond all possible doubt. If we were to be wrong about (i) this would have no serious implications that we can see. It would simply mean that membership in the concept *striped apple* can be accounted for along the lines of section 5.5 without any need to worry about links to a possibly erratic and variable prototype. In contrast, if we were to be wrong about (ii), while at the same time (i) was correct, the implications would be serious. For in this eventuality there would be the threat of a breakdown of compositionality. Specifically, there would be in this situation a genuine possibility of two speakers X and Y, while in complete agreement about the component concepts *striped* and *apple*, nevertheless diverging not only in their prototypes for *striped apple* but also, via their differing prototypes, in their membership functions, $c^X_{striped\ apple}$ and $c^Y_{striped\ apple}$, for the complex concept. In such a situation $c^X_{striped\ apple}$ and $c^Y_{striped\ apple}$ might be equally legitimate, with the implication that the membership function for *striped apple* is not fixed by the component concepts.

As we said, we are inclined to believe that *striped apple* does satisfy assumption (ii) as well as (i), so that the threat to compositionality is in this instance academic. However, compositionality will be in jeopardy as soon as there is any complex concept for which (i) holds but (ii) fails. We do not know whether such concepts exist, but leave this as a question for further investigation.

6. Varieties of concepts and conceptual combination

6.1. Varieties of concepts

In the preceding sections our focus has been on concept composition. One of our main points was that there are a variety of strategies for interpreting

[29] For a discussion of the question under what conditions the two methods would produce the same result see Appendix II.

complex concept expressions, in particular those combining adjectives and nouns. In the course of that investigation we discovered differences, not just among ways of combining concepts, but also among the concepts themselves. We found it helpful to distinguish (i) between vague and sharp concepts, (ii) between concepts that do and those that do not come with a prototype, and (iii) within the class of concepts that do come with a prototype, between those where degree of membership is a matter of prototype resemblance and those where it is not. These three different features enable us to give a rough typology of concepts as in Table 1, where we abbreviate the three contrasts we have been considering as follows:

+V/−V: vague/sharp
+P/−P: does/does not have a prototype
+PE/−PE: prototype does/does not determine extension

Evidently the typology which results from looking at just these features is a very rough one. There are many other important aspects of concepts such as their context sensitivity, how their vagueness relates to their context sensitivity, whether they come with a simple prototype or a collection of prototypes, the universality of their prototypes, etc. All these aspects are important to how concepts work, and a respectable and useful conceptual typology will have to take them into account as well. However, to develop a refined typology of concepts that accounts for all the systematic logical and semantical differences between them is a large and difficult task. We wish to emphasize how important we think further work in this direction will be. Here we must limit ourselves to a few brief comments on the types that are represented in Table 1, and on the properties by which they are defined.

We must begin by saying a few words about the distinction between vague and sharp. In the preceding sections we have assumed of a number of concepts that they are sharp; there also appear some concepts under the head −V in Table 1. Strictly speaking, however, absolutely sharp concepts are quite rare; possibly they are found only in abstract domains such as those of pure mathematics. For a concept that applies to concrete things it is

Table 1

	−P	+P	
		−PE	+PE
+V	tall wide heavy not red	adolescent tall tree	red chair bully shy
−V	inanimate odd (number) not a bird	bird fish grandmother	(none)

almost always possible to imagine objects for which the concept is indeterminate. Moreover, often such objects are known to exist in reality. But even then the concept may be *de facto* sharp relative to large domains which include all that most of us ever deal with or talk about, and may consequently be construed as sharp. It is in this sense – of being *de facto* sharp and construed as sharp in most normal contexts of use – that the term "sharp" should be understood. Vague concepts, as we use the term "vague" here, are those which are not sharp even in this weak sense.[30]

The second feature we must consider is prototypicality. The distinction between concepts with and concepts without prototypes is, like that between vague and sharp concepts, not easy to make precise. A large part of the problem is that we are uncertain just how one should understand many of the extant claims to the effect that a given concept has a prototype, or has such-and-such for its prototype. As we have noted above, we have the impression that the term "prototype" has been used in a number of overlapping yet non-identical ways. We do not think that our failure to offer such an analysis here affects any point of substance we have made in the preceding sections, but it does imply that our classifications of concepts into those that have and those that do not have prototypes are necessarily somewhat tentative.

Notwithstanding this element of uncertainty there is at least one prototype-related distinction between concepts of whose importance we are firmly convinced. This is the difference between concepts whose prototypes determine membership and concepts whose prototypes do not. On this point we differ from O&S, who suggest that there never exists any logical connection between prototype and membership.

As examples of concepts for which there exists such a connection Table 1 lists *red*, *chair*, *bully*, and *shy*. We have also classified these concepts as vague. Indeed, one would expect that a concept for which membership is a matter of prototype resemblance must be vague, as resemblance is a matter of degree. For color concepts, such as *red*, and for character concepts, like *bully* and *shy*, this seems plausible enough; for *chair* the matter is perhaps more dubious. But insofar as one may have doubts on this point, that only goes to underscore how tenuous the line is which separates the vague from the sharp. In fact, for each of the four concepts there will be contexts in which it is *de facto* sharp because everything in the context either resembles its prototype to a very high or else to a very low degree. Nonetheless we

[30] It would be a useful task, but one we have not undertaken here, to look for linguistic tests for whether a given predicate is construed as sharp or vague in a given usage – tests such as co-occurrence with various predicate modifiers such as *completely*, *perfectly*, *very*, *almost*, etc. (See also the "shouting test" proposed in Unger, 1975.) The distinction is not *overtly* grammaticized in English, but it may have systematic consequences in the semantics of predicate modification and predicate combination. One and the same predicate term may well have both sharp and vague senses, much as nouns often have both mass and count variants, etc. – a fact which must be taken into consideration in devising classifactory diagnostics.

have classified each of the concepts as vague because we believe that there are many "normal" contexts that involve objects for which the concept is indeterminate.

For concepts with prototypes that determine extension, then, we have good reason to expect those concepts to be vague. But among concepts where there is no connection between prototype and membership, we may expect to find some which are vague but others which are sharp. And indeed there appear to be many concepts which have been claimed to come with prototypes but which nonetheless are sharp with respect to extension – or at any rate very close to that. Among them are some of the concepts that have been most thoroughly investigated in the work on prototypicality, such as *bird* and *fish*. *Bird* and *fish* belong to a large class of concepts known in the philosophical literature as *natural kind concepts*. Natural kind concepts have been the subject of intensive scrutiny since the provocative studies, in the late 1960s and early 1970s, by Kripke and Putnam, who were the first to recognize that such concepts enjoy a special status. Putnam's work in this domain is especially relevant to the theory of prototypes because of its emphasis on the thesis that for a natural kind concept there is *no* logical connection between extension and prototype. The argument, though well known, is important enough to merit a brief summary here.

Putnam and Kripke (see in particular Kripke, 1980; Putnam, 1975) argue that the concepts denoted in natural language by such words as *bird*, *robin*, *apple* or *water* have a reality that transcends what particular knowledge about them may be available to individual users of those terms or even to the entire language community. The application criteria of such concepts are determined by the essential properties of those objects that served and serve as exemplars for the concept: a thing falls under the concept if and only if it shares the essential characteristics of those exemplars. What is an essential property has to do with the ultimate structure of the natural world, and thus the essential properties of the things falling under the concept may well be quite different from those which users employ to determine whether something falls under the concept or not. The extension of such a concept, Putnam argues, may therefore be largely independent from the *stereotype* – the bundle of criteria which the ordinary user associates with it.

We hesitate to say that the notion of a stereotype as it figures in Putnam's writings and that of a prototype as it has emerged from the psychological literature coincide. But at the very least there are close connections between them. Usually – perhaps invariably – a thing will fit the stereotype of a prototypical concept if and only if it closely resembles the prototype. The concept *lemon* provides a good illustration: its stereotype includes a characteristic shape, color, texture, taste, etc. – just the salient properties of the ordinary person's prototypical lemon. Putnam points out that although we expect most lemons to have these stereotypical properties and though perhaps every lemon will have at least some of them, fitting the stereotype is not the *criterion* for lemonhood (see Rey, 1983). It is part of our using *lemon* as a natural kind concept that what makes something a lemon is a certain internal (in this case presumably genetic) structure. The stereotypical properties are common manifestations of that internal structure, but it is compatible with the lay person's understanding of such matters, and (to our knowledge) even with the present state of science, that something may have the internal structure without displaying any of the stereotypical symptoms. Putnam concludes that the properties of being a lemon and of fitting the stereotype of lemon are not identical; they could, and for all we know they actually do, have different extensions. Resembling the prototype of *lemon* must therefore be distinguished from being a lemon: the extension of the concept is not determined by its prototype.

The doctrine that to belong to a natural kind is to have the essential properties which characterize that kind might tempt one to think that natural kind concepts must be sharp: for either something has the essential properties and so falls under the kind concept, or it fails to have at least one of those properties, and thus does not fall under the concept. This presupposes, however, that the essential properties are themselves sharp and there is no general argument to show that this presupposition is warranted. Nevertheless, the vast majority of natural kind terms are sharp in the strict sense of being determinately true or false of everything that is found in the real world. For instance, to belong to a particular biological species an individual must have the DNA of that species; and almost without exception this is a property which an individual organism either definitely has or else definitely lacks. Being sharp in this strict sense, the vast majority of natural kind concepts are *a fortiori* sharp in the weaker sense in which the term is being used in our scheme of classification.

Natural kind concepts are not the only sharp concepts with come with prototypes. Others are, for instance, *grandmother* and the complex concept *male nurse* which we discussed in section 5.7. Again, to classify these concepts as sharp is problematic, for they too admit of borderline cases. But the class of contexts in which they are *de facto* sharp seems large enough to make such a classification seem reasonable.

Our decision to classify *male nurse* and *grandmother* as having prototypes reflects our belief that speakers typically associate representative examples with these concepts; as we already noted in relation to *male nurse* in section 5.7, what the prototype is may vary from speaker to speaker, but since the prototype does not affect questions of membership for these concepts this is of little import. (We suspect that there is less variation in the prototype for *grandmother* than there is for *male nurse*; but this too is a matter of merely anecdotal interest.)

The examples of concepts with the features +P, −PE we have just discussed were, we claimed, for practical purposes sharp. However, there are also concepts with these features that it is reasonable to classify as vague. Our table gives two instances of this category: *tall tree* and *adolescent*. We already mentioned the case of *tall tree* in section 5.2, but a brief remark on *adolescent* may be in order. We have identified *adolescent* as +P and −PE because, as for *male nurse* and *grandmother*, it seems reasonable to assume that normal speakers have an idea of what the prototypical adolescent is like, while realizing that what counts for membership in the concept is just age and/or biological maturity – if you are over 30, for instance, you simply cannot qualify as an adolescent no matter how much you resemble the prototype in dress, behavior, etc.; and if you are 15, then you presumably qualify even though you do not resemble it very much at all. Yet, unlike *male nurse* and *grandmother*, *adolescent* has been classified as vague. For there does not seem to be a well-defined point in chronological or biological age when you first become an adolescent, nor one where adolescence ends.[31]

As regards membership, we should expect there to be little difference between concepts of the +P, −PE variety and concepts which do not have a prototype at all. So one might expect that just as there are both vague and sharp concepts of the first kind there are also vague and sharp concepts of the second. This seems to be so. As an example of a prototype-free vague concept we have given *tall*. Over *tall's* vagueness there can hardly be any argument. We also think it is quite clear that *tall* has no prototype. This has to do with the fact that it can be applied to an indefinite variety of things and with the circumstance that there is in general no natural upper bound to how tall things can be. Other unbounded scalar concepts, such as *heavy*, *big*, *wide*, etc. also belong to this type.

Examples of concepts that are without prototype and sharp are of various kinds. Some are abstract concepts, such as arithmetic notions like *odd* and *even* (for discussion see Armstrong et al., 1983). Some are complex concepts, which we discuss further in section 6.2 below. A concept formed via negation, for example, may lack a prototype even when its unnegated counterpart clearly qualifies as prototypical (by the loose and informal criteria we are applying). For instance, many speakers have been shown to associate a prototype with *bird*; but it is hard to see what the prototype could be for *not a bird*. We will return to this point in the next section.

Only one of the types for which our table provides is without instance altogether. This is the type of concept for which membership is determined by prototype resemblance but which is sharp nevertheless. There is no *a priori* reason why there could be no concepts of this sort. For even though resemblance is in principle a matter of degree, it might well be that, as things are, everything either resembles the prototype of some given concept to a maximal or else to a minimal degree. However, we have not come upon

[31] In fact, *adolescent* is one of those concepts that give rise to the so-called "bald man paradox", according to which a man with a full head of hair must be accepted as an instance of the predicate "bald" as soon as we accept that one hair more or less cannot make a difference to the question whether one is bald or not. See the references to Pinkal, as well as Kamp (1981), for discussion of the phenomenon.

any convincing instances of this type. (It may well be that even if there were such concepts they would be hard to recognize as sharp and yet as +PE, so that learnability considerations might strongly disfavor them.)

In this subsection we have described a very rough typology of concepts in terms of features that the earlier sections revealed to be important. We have briefly described the different types of this typology, given some illustration of each of the types that we believe to be instantiated and tried to motivate our choices. In the next subsection we briefly consider what happens when concepts of given types are combined to form complex ones.

6.2. *Concept types and concept combination*

Once concepts have been classified into different types the question arises how the type of a complex concept is related to the types of its constituents. This appears to be a fairly complicated question – even when we restrict ourselves to only the types contemplated in section 6.1 and to those modes of conceptual combination that were discussed in the preceding sections – and a serious attempt to come up with a detailed answer would require another paper. Here we will restrict ourselves to a few observations relating to complex concepts C formed by negation, conjunction and disjunction. Our main concern is to show that the properties illustrated in Table 1 are in general not preserved under such combinations. That is, the classification of a complex concept need not be the same as that of its constituent or constituents. We will mainly focus on the property of having a prototype but we will also have something to say about the property of being a concept for which prototype is relevant to membership.

In particular, the negation of concept C does not as a rule appear to have a prototype even if C itself does. This is true both for concepts whose prototypes determine membership and for those whose prototypes do not. As an example of the second sort, take *apple*. Presumably most speakers associate some kind of prototype with *apple*. But what might be the prototype of the concept (*is*) *not an apple*? A pear? A banana? An artichoke? Kareem Abdul Jabbar? The number 27? All of these are perfectly good cases of non-apples; but they are very different from each other, which is probably why none of them is plausible as a *prototypical* non-apple.[32]

Since the extension of *apple* is not determined by prototype and the

[32] An anonymous referee has pointed out to us that he or she seems to have a coherent prototype for "non-athlete"; this prototype excludes artichokes, for instance, and "just includes the really clumsy types with particular body builds". This accords with our own intuitions as well, and makes it clear that we cannot strengthen the claim above to a claim that the negation of a concept that has a prototype itself *never* has a prototype; we have no immediate suggestions as to which negations of concepts will have prototypes, or whether this depends more on intrinsic semantic properties or more on properties of the history of the use of the term in actual contexts.

extension of the concept (*is*) *not an apple* is the set-theoretic complement of the extension of *apple*, the lack of a prototypical non-apple has no consequences for the question of membership. For a concept whose prototype does determine membership the situation is necessarily different. Consider *red*. The concept (*is*) *not red* does not appear to have a prototype; for how might one resolve the choice among *white*, *green*, *black*, *yellow* and all the other colors that *red* excludes? Nevertheless the degree of membership in the concept *not red* is a matter of prototypicality. Only, the relevant prototype is not some prototype for *not red* but the prototype for *red*, and the degree to which something is not red is a matter of how *little*, rather than how much, it resembles that prototype.

Similar remarks apply to complex concepts that are obtained by conjoining one concept with the negation of a second, as in *chair that is not red*. We doubt that it makes sense to speak of prototypical non-red chairs any more than it is possible to make sense of a prototype for *not red*. Yet whether or to what degree something is a chair that is not red does, we think, involve an element of prototypicality, but the prototype involved is not some prototypical non-red chair, nor for that matter some prototypically non-red color patch. Rather, our understanding of the concept rests, in some way, on what we have by way of prototypes for *chair* and *red*.

Note that a supervaluation account (whether formulated as in section 4.3 or as in 5.5) will handle the (graded) extension of these complex concepts correctly. In particular, a presupermodel in which c_{red} (and therewith the spectrum of extensions of *red* in the completion set $\mathcal{S}$) is based on resemblance to the prototypical red will assign the concepts *not red* and *chair that is not red* membership functions which agree with the intuitions just expressed.

Disjunction appears to be a mode of combination which does not seem to preserve prototypicality either. Robins are reportedly the prototypical bird; and trout, we may assume for the sake of argument, are prototypical fish. But neither could be *the* single prototype for the concept "is either a bird or a fish". Ducks, penguins, or flying fish hold no better title.

Disjunctive concepts raise an issue which we have not touched upon so far, but which a discussion of prototypicality cannot ignore. Can a concept be prototypical in the sense of being associated with several prototypes rather than with just one? We mention this as yet another topic for further analysis. It is worth noting in this connection that when the constituent concepts of a disjunctive concept not only have prototypes but have prototypes that determine membership, then these prototypes may both be essential to the extension of their disjunction. For instance, the degree to which something satisfies the concept *red or green* is presumably a matter of how much it resembles either the prototype for *red* or the prototype for *green*.

However, there are other disjunctive concepts for which this does not appear to be the correct analysis. Consider the instance the disjunctive concept *red or pink*. A color shade which lies on the borderline between *red* and *pink* – and, it may therefore be assumed, is "equidistant" between the prototype for *red* and that for *pink* – is neither a clear case of red nor a clear case of pink, for it does not have enough resemblance to either prototype. Nevertheless many speakers are prepared to regard such a shade as a perfect case of *red or pink*. The explanation that has been offered for this, and which in part motivated the supervaluation approach, is that no matter where the line between red and pink should be drawn, the shade would fall squarely within the union of the extensions *red* and *pink* (see Fine, 1975).

Once more it appears that the supervaluation theory delivers the desired results. It will assign to *red or green* a membership function which is connected with the prototypes for *red* and *green* in the way described; and it will make a shade in the borderline of red and pink a good case of *red or pink*.[33]

The few observations we have made in this section reinforce the picture which has been gradually emerging in the course of this paper. According to that picture, concepts come in a variety of different types; there are concepts without and concepts with prototypes, and among the latter there are ones for which prototype determines membership and others for which it does not. The present section has shown how concepts of one type can lead, via negation or disjunctive combination, to concepts of another.

One of the conclusions to which all this points is that a theory of conceptual combination must satisfy two requirements. On the one hand it should be able to do justice to the logical connection between prototype and membership for those cases in which such a connection exists. On the other hand it must provide a general account of concept membership which applies to arbitrary concepts and treats the concepts that show a connection between prototype and membership as one type among others. The framework we have sketched in sections 4.3 and 4.5 is no more than a first

[33] Some speakers are inclined to deny that a shade on the borderline of red and pink is a clear case of red-or-pink. It is possible to do justice to this intuition within the supervaluation framework if we are prepared to accept that the interpretation of *red-or-pink* on which these speakers fasten is paraphrasable in some way as "either clearly red or clearly pink". Fine (1975) outlined a way of accounting for the relevant sense of "clearly" in terms of an operator "definite", but his proposal is technically too complicated to spell out here. It is also possible that the expression *red or pink* is genuinely ambiguous with one of its readings amounting to a disjunctive description of a single color (in which case a borderline case would satisfy the description) whereas the other reading presents a disjunction of two colors (in which case a borderline color would be a poor example insofar as it is a poor example of either one). But we do not know of any evidence that would support such a claim of ambiguity nor of any analysis that would predict it. The following examples give some weak evidence for an ambiguity, but not conclusive evidence.

(i) If it's red or pink, put it in basket 1 or basket 2 respectively; if it's any other color, put it in basket 3.
(ii) If it's red or pink, put it in basket 1; if it's any other color, put it in basket 2.

pass at such a theory; but we hope that it provides some idea of how a more comprehensive semantic theory might be constructed.

6.3. Compositionality again

A compositional theory in any domain is a theory which offers an explanation of some property P or family of properties $P_1, \ldots, P_n$ of some "complex" A in terms of the corresponding properties of some simpler or constituent As; the idea is to show how the value of P (or the Ps) for the more complex A is a function of the values of P (or the Ps) for the simpler As and of the particular relation of the simpler As to the complex one. In the case of semantics, P is meaning, As are expressions. In the case of probability theory, P is a distribution in a probability space, As are event types. For prototype theory, O&S originally articulated the goal in terms of P being the c function, As being concepts; we have argued for revising the goal so that P is a fuller set of semantic properties of a concept, on the grounds that the demand for a compositional determination of c-functions from c-functions alone reflects an excessively narrow view of compositionality in general, besides being unworkable for the kinds of reasons O&S put forward.[34]

Indeed, one of the points of this paper has been to show that part of the problem of finding a compositional account of the concept combinations O&S consider is of this very sort. Even if one's primary concern is that of accounting for the c-functions of complex concepts, there is no *a priori* reason why it should be possible, as O&S seem to implicitly assume, to do this just by looking at the c-functions of the parts. While agreeing with O&S's arguments to the effect that *such* an account could not work, we hope that the positive suggestions we have made in this paper lend plausibility to the notion that a compositional account of graded extension is possible nevertheless when the c-functions are seen as part of a larger complex of semantic and logical attributes.

In the account we have proposed here these attributes include the particular way in which the possible sharp extensions of a concept are distributed over the members of the completion set $\mathcal{S}$. (We have seen how this distribution is constrained by the c-function but not fully determined by

[34] The search for a compositional theory in any domain may be valuable precisely because it may lead us to discover a rich structure of interconnected properties in the given domain; it may lead us to such discoveries as we find it to be impossible to achieve compositionality starting from a too impoverished set of properties. For instance, in semantics the move from extensions to intensions was motivated and argued for by the failure of compositionality at the level of extensions alone; similarly in probability theory one needs a Boolean-structured probability space, not just a linear scale of probabilities, since as a rule the numerical probability values fail to provide crucial information about the independence or non-independence of different event types that is necessary to compute the probability of a complex event type in terms of the probabilities of the relevant simpler event types.

it.) But this constraint is not enough, we saw, to deal with the *striped apple* problem. To deal with that problem we had to assume that the process of conceptual combination proper is preceded by one of contextual calibration. Thus we came to postulate a two-step procedure to get us from the meanings of *striped* and *apple* to that of *striped apple*. This introduces a new kind of complexity, that of *modularity*: certain semantic processes involve the operation of two (or more) distinct "modules". Each module operates according to its own principles, presupposes its own conceptual distinctions and involves its own input and output structures. The different modules interact either in that, as in the case under discussion, one operates on the output of another, or in that different modules compute different components of the same, joint, output.[35]

The "modularity" of the account we offer in section 5 for the compositionality of combinations such as *striped apple* raises an important question which we have not yet touched upon: should we see the recalibrations described in section 5 as recalibrations of concepts or as recalibrations of expressions? Where in the preceding sections our wordings seemed to carry any commitment at all on this point, it always implied that we saw the process as one of *concept* calibration. This should not be interpreted, however, as expressing a firm conviction that this is what calibration comes to always. In fact, though we have not explored the question in detail, we are inclined to think that the answer may well have to be different in different cases. Thus we believe on the one hand that *tall* can best be seen as expressing a single concept, and that it is this *concept* that gets calibrated in the various contexts which different head nouns (often in cooperation with other factors) determine. On the other hand, the "recalibration" of *lion* in the context of *stone* looks to us more like a process where the word *lion* is reinterpreted, so that it comes to express a different concept – the concept of being a "lion representation" rather than a real lion. For many other concept words we have at present no clear idea what the answer should be. Thus it is not at all clear to us what to think about, say, *striped* – do we, in recalibrating *striped* in the context of *apple*, exploit a single, unitary concept to optimize its interaction with the concept *apple*, or do we switch to a new concept of "apple stripes"?

[35] Until fairly recently, modularity played only a marginal role in formal semantics, although it has been central to the conception of more comprehensive descriptions of language – for example, those which deal with pragmatic as well as semantic phenomena; or, alternatively, theories which deal with both semantics and syntax. There exist certain precedents for modularity within specifically semantic contexts, such as the proposals for type shifting (cf. Partee & Rooth, 1983, Rooth & Partee, 1982; Partee, 1987). Also, there are a number of recent approaches to the description of natural language which are based on the conviction that the traditional boundaries between semantics and pragmatics are theoretically untenable and in which "semantics" and pragmatics" are intimately intertwined. The architecture of these theories tends to incorporate assumptions of modularity somewhat in the same way as one finds in earlier theories which present themselves as combinations of semantics and pragmatics.

The question when contextual calibration is true recalibration of a concept and when it is a matter of reinterpreting a word is only one of several which complicate the assessment of what has been for us, as it was for O&S, the central issue: are prototypicality and compositionality compatible, and, if the answer is yes, *how* are prototypical concepts combined into compound concepts? While our conclusions cannot be said to offer any particular support for prototype theory, we hope to have shown that O&S's negative conclusions on this point were derived more from defects of fuzzy-set theory than from demonstrated defects of prototype theory, and we have argued that the phenomena involved are interestingly complex in ways that are partly understood and partly in need of further investigation. Thus, for instance, we found it necessary to distinguish two different c-functions, c^e and c^p; and with that distinction came another, that between concepts whose prototypes are instrumental in determining their extensions, and those for which this is not so, a distinction whose importance has also been emphasized by Armstrong, Gleitman and Gleitman (1983). This, we saw in section 6.1, is only one of a number of different ways in which concepts must be distinguished before we can hope to come up with a coherent theory of concept combination. The matter is further complicated, we noted in the last section, by the fact that the component concepts that enter into a conceptual combination may have distinct characteristics; moreover, there does not appear to be any simple correspondence between the characteristics of the component concepts and those of the compound. Consequently the number of different types of concept combination is substantial even if we take no other concept characteristics into account than those we discussed in section 6.1.

For these and other reasons conceptual combination is a much more diversified problem than traditional accounts of compositionality give them credit for. This means that a theory of compositional semantics will have to be considerably more complex than those currently in use. But of what such a theory should look like in detail we have only been able to give the merest hint. In connection with most problems that such a theory should address we have done no more than point out that the problem exists. Almost all the real hard work on the topic will still have to be done.

In the course of discussing apparent counterexamples to the possibility of rendering prototype theory and compositionality compatible and identifying the factors that seem to us to be involved in an explanation of the phenomena, we have emphasized that many of the issues involved arise in the analysis of vague predicates generally, both those with which we associate prototypes and those with which we do not. In particular, the principles we suggested in connection with the dynamics of context dependence and the readjustment of the interpretation of parts and wholes to satisfy principles like the non-vacuity principle are all principles which extend to arbitrary vague concepts, and do not just apply to those that have prototypes.

What we seem to have ended up with is a view of the phenomena which is compatible with prototype theory but one in which the notion of a prototype is not particularly central. From this perspective prototype theory appears a partial theory of vague concepts – one that may well be consistent with and embeddable in the kind of richer and more comprehensive theory that will be necessary to account for the full spectrum of facts concerning concepts and conceptual combination, but which on its own can deal only with a limited and arguably non-natural subclass of those phenomena.

7. Conclusions

The central question posed by Osherson and Smith (1981), and reiterated even more sharply in Smith and Osherson (1984), is whether prototype theory is or is not compatible with the phenomena of conceptual combination. In the later paper, their skeptical conclusion is "that no extant version of the gradient thesis can serve as a framework for a theory of conceptual combination", where by *the gradient thesis* they mean "the view that conceptual application is essentially graded" (as opposed to *the binary thesis* – the view that conceptual application is essentially dichotomous.)

Our arguments in this paper have supported O&S's narrower claims about the unsuitability of prototype theory in combination with fuzzy-set theory as a theory of conceptual combination, but we have not endorsed their suspicion "that the latter phenomena [of conceptual gradation] are insulated from the logic of complex ideas, and hence, that the theoretical machinery needed to explain typicality phenomena is different in kind from that needed to explain compositionality". It is probably true that the kinds of theoretical machinery that have been developed in the two domains thus far have been rather different in kind; we expect that it will take a considerable amount of work to develop more comprehensive theories that do justice to the union of the central concerns of theories of compositionality and of theories of prototypicality.

We do believe, however, that there are cases where resemblance to a prototype does play a central role in determining extension, and where the phenomena of vagueness are therefore inextricably linked up with the "logic" of concepts and of conceptual combination. In these cases there must be a way to specify compositional principles for interpreting conceptual combinations, or the basic principle of compositionality is threatened and the predictability of speaker–hearer's interpretations of novel combinations becomes a mystery.

However, it has become clear that the factors and mechanisms involved in any possible compositional semantics for prototypical concepts interact in complex ways that makes their investigation difficult. In light of the simplicity and elegance of the logic of combination for sharp concepts – at least for basic combinations such as conjunction – it is natural and reasonable to start by examining the possibility of rather simple modifications to truth values and truth functions, as in fuzzy logic. But the arguments of O&S and others make it clear that no such simple operations on "fuzzy characteristic functions" can correctly determine the fuzzy characteristic function of a complex concept from those of its parts.

We have argued that there are a number of factors that must be taken into account in trying to articulate a compositional theory adequate to the challenge presented by vague concepts. One of the important factors we identified was context dependence. Wherever there is vagueness we may expect to find context dependence, since the context is what helps us resolve or narrow down the vagueness in most cases; and the effects of context, both linguistic and non-linguistic, are clearly quite diverse. It is often very hard to draw lines between linguistic, conceptual, and real-world knowledge; vagueness clearly adds to the difficulty of discriminating between attributes intrinsic to a given concept and those that may be accidentally yet very frequently associated with it. Likewise, as we pointed out in section 6.3, it may be hard to tell whether a given word is uniformly interpreted as a single intrinsically vague concept or whether there is instead or in addition some shifting of the meaning(s) of the given word in different contexts.

When the vagueness is connected with resemblance to a prototype, as in many of the examples dealt with here, there are added difficulties created by the problem of distinguishing the roles of the two functions c^e (degree of membership in the extension of a given concept), and c^p (degree of resemblance to the prototype) – difficulties which we discussed in sections 5.8 and 6.1. There are hard issues here even for simple concepts, and the problems become harder still when such concepts are combined, as we illustrated at various points in the paper, particularly in sections 5.2, 5.4, and 6.2.[36]

Our positive suggestions in response to the problems we have pointed out are largely *samples* of *kinds* of methods that we believe might be fruitfully exploited in constructing a comprehensive compositional theory of conceptual combination. We have conducted no experiments and we make no pretense of exhaustiveness of coverage. The first goal of the paper and the main subject of sections 2 through 4 was to give a clear demonstration of the advantages of supervaluation theory over fuzzy-set theory as a theory of the

[36] O&S stimulated a substantial number of reactions and further studies, in which a good many of these difficult problems are taken up. Some of these papers came to attention while we were already quite far along with our work on the present paper. To have devoted to those papers the careful discussion which they deserve would have required significant changes to its structure as well as a considerable increase in length, neither of which seemed advisable to us. Among the papers which ideally we should have discussed we wish to mention in particular Hampton (1987, 1988a,b), Malt and Smith (1984), Medin & Smith (1984), Murphy (1988, 1990), Osherson and Smith (1982), Smith and Osherson (1984), Smith, Osherson, Rips and Keane (1988), Zadeh (1982).

logic of vague concepts and their combination. But the failure of supervaluation theory alone to provide an account of reported data for the *striped apple* case required that we suggest at least a possible account of the effects of context dependence and their interaction with supervaluation theory. This we did in section 5. The aim in that section and the next has been primarily to identify the *kinds* of issues that need to be resolved, to show that several factors seem to interact at once, and to try to make plausible that one can discern robust compositional principles of combination underneath the tangle of idiosyncratic and variable detail.

To the skeptic it might seem that compositionality has little force if the basic principles are complicated and almost obscured by such a variety of contextual factors, lexical differences, etc. But it must be reiterated that compositionality in some form appears to be a crucial part of any account of semantic competence. And it is also possible to look at the demand for compositionality, together with other quite general principles like those suggested in section 5.3, as central driving principles which induce much of the recalibration and reinterpretation that we seem to find.

Our main positive points are probably best stated in terms of the broad working hypotheses of compositionality and modularity. Modularity drives the search for relatively autonomous subsystems, and could motivate the attempt to try, for instance, to derive the c-function of a complex whole from the c-functions of its constituents. Compositionality requires us to identify systems of interacting properties such that the properties of the whole can be derived from the properties of the parts. We have shown that excessive isolation of properties into modules is incompatible with compositionality; thus, as we argued in sections 4.1–4.3, one cannot derive the c^e of the whole from the c^e of the parts. Similarly, in section 6.2 we emphasized that one cannot determine the prototype of the whole from the prototype(s) of the parts; the whole may not *have* a prototype even when the parts do, even in cases where resemblance to the prototype does determine extension. But these facts, we have argued, do not support O&S's global pessimism about the possibility of a compositional theory of combination for gradient concepts, because the having of a prototype and the relation of prototype to extension are just *parts* of a full semantic description of a concept, and do not form an autonomous "module". Our very speculative suggestions involving presupermodels, measure functions, etc., are designed to show that it may be possible to do a much better job of accounting for conceptual combination with a richer cluster of properties in the relevant semantic "module", or system of interacting modules. And we have suggested, though not proved, that when a suitably rich compositional theory adequate for conceptual combination of vague concepts is developed, prototypes will be seen not as incompatible with the logic of vague concepts but as one property among many which only when taken together can support a compositional theory of combination.

Acknowledgements

In the long and sometimes tortuous course of the writing of this paper, the authors have accumulated a considerable debt of gratitude to many who have contributed in a variety of ways to its form, its content, and to its very existence and eventual completion. For early encouragement (1982) and valuable discussions (1984) we are especially indebted to Dan Osherson and Ed Smith; and the reader will readily see that it was their 1981 paper that provided the inspiration and impetus for the present one. We received extremely useful comments on our 1986 draft from Fred Landman in 1988 after he had used it in a Cognitive Science seminar at Cornell – comments which led to substantial improvements. Other valuable comments and suggestions came at various stages from Lyn Frazier, Vilem Novák, Laszlo Kálmán, and members of audiences to whom one or both of the authors presented the material, including an MIT Cognitive Science seminar in 1982, a University of Massachusetts Cognitive Science seminar in 1984, the Research Institute for Linguistics, Hungarian Academy of Sciences in 1989, and a seminar on grammar and semantics of the Czechoslovakian Scientific and Technical Society in Prague in 1989.

In the course of preparing and revising numerous drafts of this paper and formatting it in LaTeX so that our formulas as well as our text would be readable, we had invaluable technical assistance from Emma Pearse, Dikran Karageuzian, Jonni Kanerva, Purna Mishra, John Brolio, Louis Conover, and Bernd Langner. Kathy Adamczyk typed many drafts and did a remarkable job of deciphering combinations of marked-up old drafts, hard-to-read handwritten additions, and cassette tapes.

We are most grateful to an anonymous referee for this journal for a number of specific helpful suggestions and, equally importantly, for pointing out to our surprise and dismay that our penultimate draft in places sounded condescending towards psychologists in general and Osherson and Smith in particular; we hope we have corrected that unintended and unfortunate tone in the final version, for that was quite the opposite of our intent.

We are also indebted to many sources for assistance in getting together to work on the paper; the difficulty of finding times and places to work together was one of the main impediments to progress, but our own schedules were the major culprit, since we generally had generous assistance with transportation costs. We acknowledge with thanks support from all of the following: the Cognitive Science Program of the University of Texas at Austin; the Center for the Study of Language and Information at Stanford University; System Development Foundation Grant #650 to the University of Massachusetts for Research on the Formal Foundations of Semantics; the Centrale Interfaculteit of the University of Amsterdam; and the organizers and sponsors of the First Symposium on Theoretical Aspects of Knowledge Representation at Asilomar, California.

Last but not least, we are most grateful to our respective spouses, Blanca

Bartoš and Emmon Bach, for their supportive accommodation of the disruptive schedules that marked our periodic bursts of progress on this project.

It is inevitable that with a paper this long in the works we have forgotten to acknowledge some to whom acknowledgement is surely due; to them too our thanks as well as apologies.

Appendix I (to section 5.5)

The proposal presented in section 5.5 should be seen as something like a first approximation, which moreover is applicable only in some of the cases where, intuitively, recalibration occurs. First, in cases where the proposal might be considered to have some plausibility (as, in our own view, that of *striped apple*), one would want evidence that it is this, and not some different process, with equivalent or nearly equivalent output, that is involved in recalibration.

More seriously there are many cases for which our proposal does not seem to work at all and among those appear to be in particular those on which we concentrated in our discussion of adjective–noun combinations in section 5.2, namely, *giant midget* or *midget giant*. In fact, the proposal seems unworkable for a very wide, and much discussed, class of cases – those where the adjective is a "scalar" such as *tall*, *heavy*, or *loud*. In order to apply our recalibration proposal to a combination like *tall tree*, for instance, we would have to suppose that the c-function for *tall* with which we start takes a value within [0,1] for every tree, and that these values correctly reflect comparative height (such that if a is a taller tree than b, then $c_{tall}(a) > c_{tall}(b)$. As long as we confine our perspective to objects of sizes roughly comparable to trees, this assumption may not seem unreasonable. However, it does seem reasonable to demand that the function c_{tall} should be able to serve as a basis not only for recalibration in the context of a concept such as *tree*, but also in those of, say, *mushroom* and *mountain*. To serve as a suitable basis for all these recalibrations c_{tall} must assign numbers in [0,1] to objects ranging from the smallest mushrooms to the tallest mountains. Of course, from a purely formal point of view, it is clear that such functions must exist since the set of all positive numbers can be isomorphically mapped onto the open interval [0,1]. However, it is hard to see such a function as having any direct bearing on *membership* in the concept *tall*. This reflection suggests that short of a calibrating context, membership in the concept *tall* is not really defined at all. Rather, inasmuch as *tall* has a conceptual content, it is one that manifests itself indirectly, for example, through the various concepts that emerge when *tall* is calibrated in the context of some particular set or noun.

Indeed, one of the existing semantic analyses of scalar adjectives (or "degree adjectives" as they are also called) is precisely as functions from contexts to (partially defined) concepts (see Klein, 1980). Other analyses treat scalars as concealed relations (i.e., as having a meaning that becomes explicit in comparative constructions such as "*a* is taller than *b*") and construe "*a* is a tall *x*" as something like "*a* is taller than most *x*s" (see Platts, 1979). Yet others have argued for a relational analysis of scalars, according to which they are relations not between individuals to both of which the concept is applicable but rather between such an individual and a degree (the degree to which the concept applies to the individual; cf. Cresswell, 1978). There is no question of surveying the fairly extensive semantical literature on this topic here; but the little we have said should suffice as an indication not only that more should be said about recalibration than we have done, but also that much more could be said by pursuing the leads of already existing proposals.

The combinations *giant midget* and *midget giant* present an additional problem. As nouns *giant* and *midget* are, we think, understood as having probably vague but nevertheless clearly disjoint boundaries. That is, even though c_{midget} and c_{giant} admit of intermediate values, we have for all a that if $c_{giant}(a) \neq 0$ then $c_{midget}(a) = 0$. This means that when our procedure is applied to yield the recalibration of *midget* in the context of *giant* and uses the function c_{midget} as input, it clearly will not give us what we want. For, on the apparently inescapable assumption that the set of giants relative to which we recalibrate consists of individuals a for all of which $c_{giant}(a) > 0$, it follows that for every a in that set $c_{midget}(a) = 0$, and so the formula for $c_{midget/giant}(a)$ is undefined (it has 0 in both numerator and denominator). The resolution of this problem must, we think, be sought in the following direction. The prenominal use of *midget* induces, for a start, the reinterpretation according to which it comes to mean something like "midget-like" – a concept akin to *tiny* with possibly some overtones of "dwarf-like". This adjectival concept is then to be calibrated in the context of *giant*, just as *tiny* would be calibrated in that context.[37]

Appendix II (to section 5.8)

In section 5.8 we mentioned two methods for arriving at the function c_{AN}, where AN is a complex concept for which membership is a matter of resemblance to prototype. This raises the question whether the two methods in fact yield the same result.

Whether they do cannot be established in the absence of a more detailed treatment of the second method than we have given. But let us at least

[37] Note in this connection that *midget* can be used not only in combination with nouns that, unmodified, are instantiated by individuals that are substantially larger than midgets, but also, though perhaps somewhat less naturally, with concepts whose instances are smaller, as in *midget mushroom* or *midget squirrel*.

briefly explore on what the question depends. According to the first method, spelled out in some detail in section 5.5,

$$c_{AN}(a) = c_{A/N}(a) = \frac{c_A(a) - c^-_{A/N}}{c^+_{A/N} - c^-_{A/N}}$$

for $a \in \|N\|$ and

$$c_{AN}(a) = 0 \text{ otherwise.}$$

As regards the second method, we have shown that a prototype p for the concept AN would have to have the maximal value that a reaches within $\|N\|$, thus $c^+_{A/N}$. The second method should, according to the suggestions made by O&S, yield $c_{AN}(a)$ as a monotone decreasing functional F that applies to some suitable function D which measures distance from p; F and D should be such that in particular p itself receives the value 1, i.e.

(a) $F(D(p)) = 1$

If $D(a)$ measures the distance of a from p, D presumably has (as O&S assume) the form $d(a, p)$, where d is a binary distance function in the usual sense (see section 2). Thus the requirement that the two methods produce the same results means that d and F must be such that

(b) for all $a \in \|N\|$ $F(d(a, p)) = \dfrac{c_A(a) - c^-_{A/N}}{c^+_{A/N} - c^-_{A/N}}$

and

(c) for all $a \in \|N\|$ $F(d(a, p)) = 0$

As d is a distance function, $d(p, p) = 0$; so (a) reduces to

(a′) $F(0) = 1$

As long as no more is given about F and d, the question whether (a′), (b) and (c) are satisfied is ill defined: clearly there are combinations F and d which will satisfy them, while on the other hand most combinations will not. If we settle on a particular F, for example, $F(D(a)) = max(1 - D(a), 0)$, the question becomes[38] that of whether

(b′) for $a \in \|N\|$ $max(1 - d(a, p), 0) = \dfrac{c_A(a) - c^-_{A/N}}{c^+_{A/N} - c^-_{A/N}}$

and

(c′) for $a \not\in \|N\|$ $max(1 - d(a, p), 0) = 0$

(c′) requires that $d(a, p) \geq 1$ for $a \not\in \|N\|$. (b') implies that for $a \in \|N\|$

[38] Evidently (a) and (a′) are satisfied in this case.

$$d(a, p) = 1 - \frac{c_A(a)}{c^+_{A/N} - c^-_{A/N}}$$

We do not know, however, of any independent motivation for precisely this distance function. Nor do we, for that matter, know any plausible justification of $max(1 - D(a), 0)$ as the choice for the function $F(D)$.

References

Armstrong, S.L., Gleitman, L.R., & Gleitman, H. (1983). What some concepts might not be. *Cognition*, *13*, 263–308.

Carnap, R. (1952). Meaning postulates. *Philosophical Studies*, *III*, 65–73.

Chien, A. (1985). *On vagueness*. PhD dissertation, University of Massachusetts at Amherst.

Chierchia, G., Partee, B., & Turner, R. (Eds.) (1989). *Properties, types and meaning, Vol. 1: Foundational issues*. Dordrecht: Kluwer.

Clark, R. (1970). Concerning the logic of predicate modifiers. *Nous*, *4*, 311–336.

Cohen, B., & Murphy, G.L. (1984). Models of concepts. *Cognitive Science*, *8*, 27–58.

Cresswell, M.J. (1978). Prepositions and point of view. *Linguistics and Philosophy*, *2*, 1–42.

Dowty, D.R., Peters, S., & Wall, R. (1981). *Introduction to Montague grammar*. Dordrecht: Reidel.

Dummett, M. (1975). Wang's paradox. *Synthese*, *30*, 301–324.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, *30*, 265–300

Goguen, J.A. (1969). The logic of inexact concepts. *Synthese*, *19*, 325–373.

Halmos, P.R. (1950). *Measure theory*. New York: Van Nostrand.

Hampton, J.A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory and Cognition*, *15*(1), 55–71.

Hampton, J.A. (1988a). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(1), 12–32.

Hampton, J.A. (1988b). Disjunction of natural concepts. *Memory and Cognition*, *16*(6), 579–591.

Kamp, H. (1975). Two theories about adjectives. In E.L. Keenan (Ed.), *Formal semantics for natural languages*, Cambridge, UK: Cambridge University Press. (pp. 123–155).

Kamp, H. (1981). The paradox of the heap. In U. Mönnich (Ed.), *Aspects of philosophical logic*. Dordrecht: Reidel.

Katz, J.J., & Fodor, J.A. (1963). The structure of a semantic theory. *Language*, *39*, 170–210.

Keenan, E.L. (1974). The functional principle: generalizing the notion of "subject of". *Proceedings of the Chicago Linguistic Society*, *10*, 298–309.

Kleene, S.C. (1952). *Introduction to metamathematics*. Amsterdam: North-Holland New York: Van Nostrand.

Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, *4*, 1–45.

Kripke, S. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.

Malt, B.C., & Smith, E.E. (1982). The role of familiarity in determining typicality. *Memory and Cognition*, *10*, 69–75.

Malt, B.C., & Smith, E.E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behaviour*, *23*, 250–269.

Medin, D.L., & Smith, E.E. (1984). Concepts and concept formation. *Annual Review of Psychology*, *35*, 113–138.

Murphy, G.L. (1988). Comprehending complex concepts. *Cognitive Science*, *12*, 529–562.

Murphy, G.L. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language*, *29*, 259–288.

Osherson, D.N., & Smith, E.E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, *11*, 237–262.

Osherson, D.N., & Smith, E.E. (1982). Gradedness and conceptual combination. *Cognition*, *12*, 299–318.

Parsons, T. (1968). *A semantics for English*. Unpublished manuscript.

Partee, B.H. (1987). Noun phrase interpretation and type-shifting principles. In J. Groenendijk, et al. (Eds.), *Studies in discourse representation theory and the theory of generalized quantifiers* (pp. 115–143). Dordrecht: Foris.

Partee, B.H., & Rooth, M. (1983). Generalized cognition and type ambiguity. In R. Bäuerle, U. Egli, & A. von Stechow (Eds.), *Meaning, use and the interpretation of language. Berlin: De Gruyter.*

Pinkal, M. (1984). consistency and context change In F. Landman & F. Veltman (Eds.), *Varieties of formal semantics. Dordrecht: Foris.*

Pinkal, M. (1985). *Logik und Lexikon. Die Semantik des Unbestimmten.* Berlin: De Gruyter.

Pinkal, M. (1989). Imprecise concepts and quantification. In R. Bartsch & J. van Benthem (Eds.), *Semantics and textual Interpretation. Dordrecht: Reidel.*

Platts, M. (1979). *Ways of meaning: An introduction to a philosophy of meaning.* Boston: Routledge & Paul.

Putnam, H. (1975). The meaning of meaning. In K. Gunderson (Ed.), *Minnesota studies in the philosophy of science* (Vol. VII). Minneapolis: University of Minnesota Press.

Rey, G. (1983). Concepts and stereotypes. *Cognition*, *15*, 237–262.

Rooth, M., & Partee, B.H. (1982). Conjunction, type ambiguity, and wide scope "or". In D. Flickinger et al. (Eds.), *Proceedings of the First West Coast Conference on formal Linguistics*, Stanford University, pp. 353–362.

Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T.E. Moore (Ed.), *Cognitive development and the acquisition of language*. New York: Academic Press.

Sanford, D. (1976). Competing semantics of vagueness: many values vs. super-truth. *Synthese*, *33*, 195–210.

Siegel, M. (1976a). *Capturing the adjective*. Unpublished PhD dissertation, University of Massachusetts at Amherst.

Siegel, M. (1976b). Capturing the Russian Adjective. In B.H. Partee (Ed.), *Montague grammar*. (pp. 293–309). New York: Academic Press.

Smith, E.E., Medin, D.L., & Rips, L.J. (1984). A psychological approach to concepts: comments on Rey's "concepts and stereotypes". *Cognition*, *17*, 265–274.

Smith, E.E. and Osherson, L.J. (1984). *Conceptual combination for prototype concepts*. Unpublished manuscript, MIT.

Smith, E.E., Osherson, D.N., Rips, L.J. and Keane, M. (1988) Combining prototypes: a selective modification model. *Cognitive Science*, *12*, 485–527.

Unger, P. (1975). *Ignorance*, Oxford: Oxford University Press.

Van Fraassen, B. (1969). Presuppositions, supervaluations and free logic. In K. Lambert (Ed.), *The logical way of doing things*. New Haven: Yale University Press.

Wright, C. (1975). On the coherence of vague predicates. *Synthese*, *30*, 325–365.

Zadeh, L.A. (1965). Fuzzy sets. *Information and Control*, *8*, 338–353.

Zadeh, L.A. (1978). PRUF – a meaning representation language for natural languages. *International Journal of Man–Machine Studies*, *10*, 395–460.

Zadeh, L.A. (1982). A note on prototype theory and fuzzy sets. *Cognition*, *12*, 291–297.