**LEXICAL SEMANTICS AND COMPOSITIONALITY**

Barbara H. Partee

Department of Linguistics
University of Massachusetts
Amherst, MA   01003-7130

partee@cs.umass.edu

December 1994

# LEXICAL SEMANTICS AND COMPOSITIONALITY

### Barbara H. Partee

## 0. INTRODUCTION.

Semantics is an inherently interdisciplinary subject, and one which benefits from the intrinsically interdisciplinary perspective of cognitive science. "Semantics" has meant different things in different disciplines: that situation is not just accidental but neither does it necessarily reflect "turf battles" or disagreements; mainly it reflects the many different ways that different disciplines are concerned with meaning. And even within a single discipline, "semantics" often means different things within different schools of thought; and there it does often reflect serious disagreement about the nature of the "best theory", disagreement about which kinds of data are most important, even disagreements about such foundational issues as whether semantics is best viewed as a "branch of mathematics" or as a "branch of psychology" (see Partee 1979.) In this latter kind of case, the arguments are between whole theories, not just between competing definitions of key terms (the arguments are not "merely semantic", to use an idiomatic expression that semanticists do not appreciate!). Everyone does agree that semantics is the study of meaning. So the big question is: What is meaning?

It is not easy to tackle a question like that head-on; and while it's an important question to keep wrestling with, a total answer is not required in advance of doing fruitful work on semantics, any more than biologists wait for the answer to the still-difficult question "what is life?" before getting down to work. A scientific community just needs _some_ clear examples to get started, and then empirical and theoretical advances proceed together, along with further sharpening of key concepts.

Semantics has roots in linguistics, psychology, anthropology, logic and philosophy of language, artificial intelligence, and more. Traditional differences in approaches to semantics in these fields reflect at least two factors. For one thing, the central questions concerning meaning may come out quite differently if one focusses on language and thought, on language and communication, on language and culture, on language and truth, on the design of natural language man-machine interfaces, or on language "structure" per se. A more accidental but no less profound source of differences is the research methodology prevalent in the field within which one approaches questions of semantics. (This means, incidentally, that semantics is an area in which the student can get a good handle on historical differences among the fields that make up cognitive science, by reading articles on similar topics by scholars from the different disciplines.) To oversimplify a bit, one might say that when investigating meanings, some linguists have tended to look for feature structures (influenced by phonology and morphology) and other linguists expect to find a level of tree structures, something like a more abstract syntax-like representation; logicians tend to think in terms of formal systems and model structures; psychologists may be interested in studying concept discrimination, concept acquisition, and principles for scaling semantic fields; artificial intelligence researchers may approach meaning representations in terms of data bases and symbol manipulation; philosophers

ask whether there are such things as meanings at all and if so what sorts of things they might be.

The perspective of the author of this chapter, like that of the author of the preceding chapter, is that of "formal semantics" with roots in logic, philosophy of language, and linguistics, developed in an environment in which linguistic and cognitive science questions have been at the forefront, with logic and the philosophy of language providing important tools and foundations.

## 1. HOW MIGHT ONE APPROACH THE QUESTION "WHAT ARE MEANINGS?"

Let's imagine that we are starting from scratch with this question. The plan of this chapter is to begin with a little philosophical reflection on broad methodological problems and foundational concerns, and then to plunge into a series of case studies all dealing with various aspects of the semantics of adjectives in English. The generalities in this section will be brief, since some of the same points have been covered in the previous chapter. The case studies we will deal with in the following sections have been chosen to involve readily graspable empirical issues that have interesting theoretical ramifications, and to compress into a chapter several decades of advances in semantic theorizing, and last but not least, to help the novice to appreciate that there are always new questions opening up for new generations of researchers to work on, as well as old questions that need to be reexamined from new perspectives.

1.1. **Compositionality.** One of the starting points for thinking about what a semantic theory should be like is very similar to the main starting point of syntax: we need to account for a language user's ability to understand novel sentences, of which there are a potential infinity. Even before we have any handle on what sorts of things we should analyze meanings to be, this fundamental aspect of semantic competence provides an argument that they must be governed by some version of the Principle of Compositionality, or Frege's Principle:

> **Principle of Compositionality, First Version:** The meaning of a whole is a function of the meanings of the parts.

What are "parts"? Since one can form different, non-synonymous sentences with the same smallest parts (words and morphemes), we can conclude that the Principle of Compositionality requires a notion of part-whole structure that is based on syntactic structure.

> **Principle of Compositionality:** The meaning of a whole is a function of the meanings of the parts and of the way they are syntactically combined.

We will keep the Principle of Compositionality in the foreground as we work through our case studies, to see the powerful role it can play in helping to choose among alternative hypotheses about the meanings of words and phrases.

There are several key words in the Principle of Compositionality which on closer examination can be seen to stand for theory-dependent concepts. Sharpening the Principle of Compositionality requires a theory of syntax, to specify the nature of the relevant part-whole structure, and a theory of what meanings are and by what kinds of functions they are combined. (Here we might compare semantics with chemistry, looking at the history of theories of molecules and atoms. Clearly it was both possible and necessary to investigate chemical structure and chemical processes without knowing the nature of the smallest parts or the fundamental forces - a clear example of "bootstrap" progress.) Let's assume we have some syntax, building on Chapter One and some of the syntactic structures introduced in Chapter Two, while remembering that every hypothesis in syntax may have repercussions for semantics, and that the Principle of Compositionality may help choose among syntactic as well as semantic analyses, as we will see at several points below. Let us turn then to the key semantic concepts of meanings and functions that combine them. Given some commonsense ideas about meanings that all of us share pretheoretically, how might we identify a notion of meaning that will support fruitful theory-building?

## 1.2. <u>Two Useful Strategies.</u>

Not surprisingly, it is philosophers who have provided two particularly useful strategies for thinking productively about the question of what meanings are. The first comes from David Lewis (1970).

> <u>Lewis's Advice</u>: "In order to say what a meaning is, we may first ask what a meaning does, and then find something that does that." (p.22)

So let's think about what meanings do, besides combine in some way to make more meanings. For this, Max Cresswell (1982) has shown how a great deal of mileage can be gotten from a very minimal assumption. Cresswell notes that we don't have any good a priori conception of what meanings are, but we do know at least one thing about them, which he dubs his "Most Certain Principle".

> <u>Cresswell's "Most Certain Principle"</u>: "For two sentences $\alpha$ and $\beta$, if [in some possible situation -BHP] $\alpha$ is true and $\beta$ is false, $\alpha$ and $\beta$ must have different meanings." (p.69)

If we follow these two strategic pieces of advice, they lead rather inevitably to the idea, already spelled out in the previous chapter, that truth-conditions are at least one fundamental part of what should go into the notion of the "meaning" of a sentence (not necessarily all, by any means). And while truth-conditions may at first look much too austere to make up a very large part of what meanings should be, it turns out to be surprisingly non-trivial to assign meanings to the lexical items and principles for combining meanings of syntactically structured parts so as to eventually arrive at relatively correct truth-conditions for sentences.

Let's just look informally at an example of the force of these strategic suggestions. It's normally accepted that "half full" and "half empty" are synonymous, just two different ways of describing the same property. But
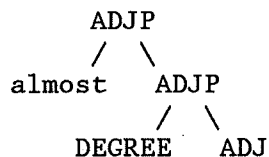
"almost half full" and "almost half empty" are clearly <u>not</u> synonymous, by
Cresswell's principle. Then by Lewis's Advice, if one of the things meanings
are supposed to do is combine to produce truth conditions of sentences, and if
the expression "almost half full" has as its main syntactic parts "almost" and
"half full", and similarly for "almost half empty", then we can argue that
"half full" and "half empty" must have different meanings after all, for how
else could one and the same meaning (the meaning of "almost") combine with the
meanings of those two expressions to give clearly different meanings as
result?  (The second of the "if"'s in the preceding sentence is a very big
"if" involving syntax, and in fact I think the hypothesis introduced about the
structure in question may well be false; we'll come back to it below.)

Is this result really counterintuitive? Well, first we may note that
with any other fraction than "half", there's not even apparent synonymy: "two-
thirds full" and "two-thirds empty" have clearly different meanings. That may
lead us to be more open to the possibility that "half full" and "half empty"
are expressions that don't really have the same meaning, but rather have
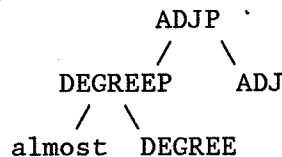different meanings that happen to be applicable to the same state of affairs.

The same example can be used to show the importance of thinking about
syntax and semantics together when evaluating proposals for compositional
semantic analyses. Suppose that the correct syntactic bracketing is rather
"almost half" plus "full" and "almost half" plus "empty". Then the argument
above suddenly evaporates. All we need then is that "full" and "empty" have
different meanings, which is uncontroversial, and that the relation between
them is such that when they are modified by any fraction other than "half",
the resulting meanings are different; but there is no longer any argument
against saying that "half full" and "half empty" are synonymous.

What is the right syntax in this case? Let's leave that for a thought
question. It's even possible that both syntactic analyses are correct, that
the phrases are structurally ambiguous, synonymous when both have structure
(1) and non-synonymous when both have structure (2).

```
(1)                              (2)
          ADJP                              ADJP
         /  \                               /  \
    almost   ADJP                     DEGREEP    ADJ
            /  \                        /  \
       DEGREE   ADJ                 almost  DEGREE
```

We have still not said very much about what sorts of things word
meanings are; but we have begun to get a handle on how to do detective work to
figure out some aspects of what we have to take word meanings to be in order
to get them to combine compositionally to produce the truth-conditional
aspects of sentence meanings.

## 1.3. Montague's Legacy.

The truth-conditional perspective just discussed did not penetrate into linguistic work until the 1970's. One of the main forces was the work of the philosopher and logician Richard Montague, who startled the linguistic and philosophical communities with his famously titled paper, "English as a Formal Language" (Montague 1970), containing the famously provocative pronouncement: "I reject the contention that an important theoretical difference exists between formal and natural languages" (p.189).

Well, of course there are important differences; but taking a perspective from which one can analyze both with the same kinds of tools, many of them developed or deployed in novel ways by Montague himself, has proved to be immensely fruitful not only for providing good formal tools for the analysis of natural languages but also for elucidating the very differences between natural and formal languages and even for suggesting some fruitful innovations for logic, computer languages, and AI. The cooperative work between linguists, philosophers, and logicians taking off from Montague's seminal works was initially known as Montague Grammar; in more recent years, as that work has evolved and innovations have led to distinct theoretical frameworks variously related to Montague's original theory, "Montague Grammar" has gradually been replaced by the broader enterprise of "formal semantics". Montague is still recognized as having laid the foundations and set much of the agenda for the field, but the contributions of linguists and others have done much to enrich it into a productively interdisciplinary endeavor. The tensions between the antipsychologistic perspective on semantics of Frege, Montague, and most logicians on the one hand and the explicitly psychologistic perspective on all of linguistics of Chomsky and the bulk of the linguistic and cognitive science communities on the other hand have led to interesting foundational debates but interestingly have not prevented very fruitful interdisciplinary progress on substantive problems in semantics. (Compare progress in the development of the differential calculus invented by Newton and Leibniz and the accompanying stormy debates about the coherence or incoherence of the notion of infinitesimals.)

## 2. CASE STUDY I. ADJECTIVE MEANINGS.

### 2.1. Groundwork: Internal Structure of NPs: Syntax and Semantics.

The previous chapter included syntactic and semantic rules for forming and interpreting NPs consisting of a DET and an N. Before we immerse ourselves in issues of adjective meanings, let's stand back a little and consider some of the other kinds of parts than can go into NPs, including adjectives (ADJ), prepositional phrases (PP), and relative clauses (REL).

In order to bring out the full range of kinds of semantic roles these parts can play in contributing to NP meanings, we need to add one more DET to our arsenal, namely "the", which has a broader range of distribution than any other DET. Whole books have been written about "the", and its semantics is much more controversial than that of "some", "every", or "no". For starters, let's work with the intuition that "the teacher", like "John", simply denotes

an individual. (We could later convert our analysis to a "set of sets" analysis; see question 2.3 in the previous chapter.) But unlike a proper name like "John", a singular definite description like "the teacher" has meaningful parts, so the interpretation must be compositionally derived. What individual does "the teacher" denote? The individual that's a teacher. Well, that answer presupposes that there is such an individual, and furthermore that there's only one. If we postpone worrying about how to make it clear that we mean one and only one in the relevant context of our utterance, and not one and only one in the whole world, we have the core of the "individual-denoting" semantic analysis of NPs of the form "the N", and we can write a semantic rule as follows:

(3) $|| [_{NP}$ the N$] || =$ the individual $\underline{a}$ such that $\underline{a}$ is the one and only member of $||N||$, if $||N||$ has one and only one member; undefined otherwise.

This is by no means the last word about the semantics of <u>the</u>; but it is a reasonable first approximation and will do for now.

Armed with this much of the semantics of NPs consisting of DET + N and our general methodological principles, let's see what we can figure out about the semantic contributions of other parts of NPs. Consider examples (4 - 5).

(4)    (a) the teacher from France
       (b) the teacher of French
       (c) the French teacher

(5)    (a) the student who was curious
       (b) the student, who was curious
       (c) the curious student

The PPs, ADJs, and RELs in these examples can all be loosely described as modifiers, adding greater specificity to the meaning of the NP, but on closer examination we can identify at least three semantic roles for these added parts: "arguments" (in 4b), "restrictive modifiers" (in 4a, 5a), and "non-restrictive modifiers (in 5b)." Let's start with (4a) vs. (4b) and examine the differences in inference patterns involving them, since that is one of the good ways to investigate systematic differences in truth conditions.

(6)    (a) If Chris is the teacher from France, then
            (i) Chris is a teacher?  YES, Valid.
            (ii) Chris is from France?  YES, Valid.

(6)    (b) If Chris is the teacher of French, then
            (i) Chris is a teacher?  YES, Valid.
            (ii) Chris is of French?  ??: Invalid; arguably not
                                          even well-formed.

The pattern in (6a) is one diagnostic for the classification of the PP "from France" as a <u>modifier</u> of the noun "teacher"; the PP adds an additional property of the individual denoted by the NP. The PP "of French", by contrast, doesn't name a further property of the individual denoted by the NP -- as
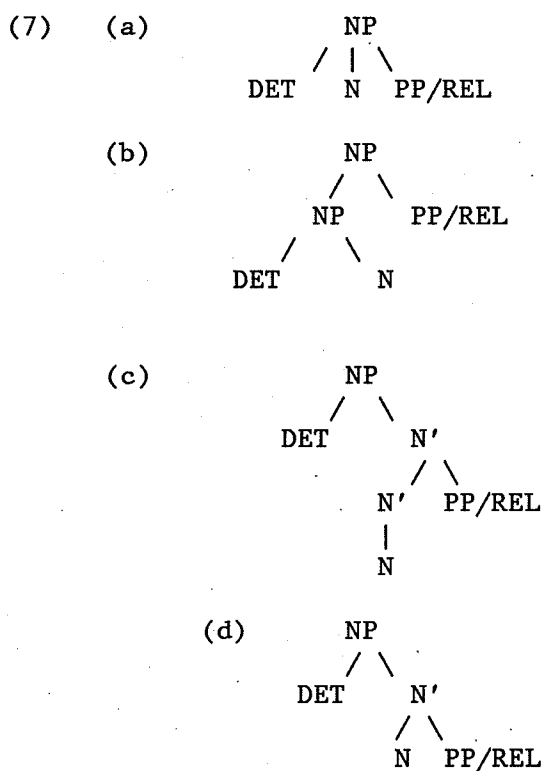
illustrated in (6b), it doesn't even sound coherent to ask whether Chris is "of French". The relation of this PP to the noun "teacher" is more like that of the direct object to the verb in "teaches French". "Teacher", like "king", "summit", "destruction" "author" or "price", can be used as a relational noun, the analog in the noun domain of a transitive verb. When an NP or a PP fills in a slot in such a relation, it is called an argument of the head; so "of French" is an argument of "teacher" in (4b) (and the subject and object(s) of a verb are likewise called arguments of the verb.) The line between modifiers and arguments is not always sharp, but there is general agreement about clear cases.

Within English NPs, "of"-PPs are mostly arguments rather than modifiers. PPs with "from" and other prepositions which express spatial or temporal or spatiotemporal relations are usually modifiers, since being "under the table" or "from France" express ordinary properties of individuals. This is one illustration of the distinction between "contentful" or "lexical" and "non-contentful" or "grammatical" prepositions. Prepositions in English are on the borderline between "open-class" and "closed-class" vocabulary, and this has interesting repercussions of several sorts. (Just to mention two: there are reports of differential responses of different kinds of aphasics to the two sorts of prepositions (see Friederici 1985); and there are different optimal strategies for dealing with the notorious difficulty of translating PPs across languages: in particular, when translating a PP with a grammatical preposition, don't try to translate the preposition at all, but just see what preposition is "demanded" by the translation of rest of the construction.)

Example (4c) is ambiguous[1]: "French" here can be interpreted either as a modifier, analogous to (4a), or as an argument, analogous to (4b). This is not a typical situation, however, as the reader can substantiate by asking whether "French" is an adjective on both readings. (If you know some other languages, try translating two readings of (4c); in many languages, it will be clear that the argument reading involves the noun "French", while the modifier reading is expressed with an adjective. And within English, you can test by substituting clear examples of adjectives and clear examples of nouns (e.g. "musical" and "music, "humorous" and "humor", and seeing what readings result.)

The added parts in (5a-c) are all modifiers; (5a) has a restrictive modifier, (5b) a non-restrictive modifier, and (5c) is ambiguous between a restrictive and a non-restrictive reading. The clearest test for this distinction, easiest to apply when the determiner is the or another definite determiner, is to ask whether the modifier does or does not play a role in determining the reference of the NP; if it does, then the modifier is restrictive; it is helping to answer the question "which N?" The role of a non-restrictive modifier is to add some further information about an independently established referent. So (5a) would typically be used in a situation in which there was more than one student, but only one who was curious, and the NP is picking out that one. But (5b) would be used in a situation in which there was only one student, or one student who was already salient in the discourse, and the non-restrictive relative clause adds an additional statement about that student. This time the ambiguity in (5c) is fully systematic and general; any adjective in that construction will produce the same ambiguity, an ambiguity which is usually readily resolved by context.

How might the distinctions among these different semantic roles, which so far we have just described informally, be captured in a compositional syntactic and semantic analysis?  Let's ask what the relevant part-whole structure should be in each case. When the modifier or argument comes after the noun, we have to decide among various possible tree structures such as the following; these do not exhaust the possibilities but represent four serious contenders.[2]  In these trees we have used PP/REL to mark the position where a PP or relative clause might go, and we have used the category N', as in the previous chapter, to stand for "common noun phrase" (the "N-bar" in the X-bar theory of Jackendoff 1977, the CN or CNP of Montague grammar) under the hypothesis that there is a syntactic division between DET and the "rest of the NP".

(7)    (a)                NP
                        /  |  \
                   DET    N   PP/REL


       (b)                NP
                        /    \
                     NP       PP/REL
                   /    \
               DET        N


       (c)                NP
                        /    \
                   DET        N'
                            /   \
                         N'    PP/REL
                          |
                          N


       (d)         NP
                 /   \
            DET       N'
                    /   \
                   N   PP/REL

Let's begin by considering the semantics of restrictive modifiers as in (4a), (5a). The principal alternatives that have been advocated in the syntactic literature have been trees (7b) and (7c), and it is worthwhile to consider how the requirement of compositionality can help choose between them. First let's eliminate trees (7a) and (7d). The simplest argument against trees (7a) and (7d) for restrictive modifiers is the fact that restrictive modifiers can be added <u>recursively</u>, that is, that we can always add more of them. Trees (7b) and (7c) allow for such recursion, because they involve <u>adjunction</u> structures of the form (8),

(8)           A
            /   \
           A     B
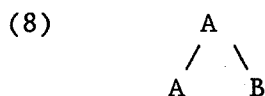
and such structures can be iterated to allow for multiple adjuncts, as in (9).

```
(9)          A
           /  \
          A    B
        /  \
       A    B
     /  \
    A    B
```
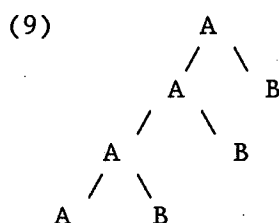
Now which is the right kind of structure for restrictive modifiers in NPs, (7b) or (7c)? We see that each analysis makes two binary subdivisions: in (7b) the DET and the N combine to make an NP, which then combines with the modifier to make a bigger NP, while in (7c) the N (by itself a minimal N′) and the modifier combine to make a bigger N′, and then the DET combines with that to make the NP. What we already know about the semantics of the is that the interpretation of a phrase of the form the $\alpha$ includes a presupposition that the set denoted by $\alpha$ has one and only one member. The noun teacher denotes a set of individuals (see the previous chapter), and so does the PP from France; similarly for the noun student and the relative clause who was curious. Under analysis (7c), the two set-denoting expressions are first combined to form a complex set-denoting expression, which can be interpreted as denoting the intersection of the two sets; combining the with the result leads to the correct presupposition that is that class that has one and only one member. On analysis (7b), on the other hand, the is first combined with the noun alone, which would lead to the presupposition that there is one and only one teacher in (4a), or one and only one student in (5a), and that inner NP would denote that individual; if the modifier combines with that already formed NP, there is no natural way for it to play a restrictive role. The structure (7b) is thus a good basis for the non-restrictive interpretation of the modifier in (5b), but not for the restrictive interpretation of (4a) or (5a) (this is argued in Partee 1973 and Rodman 1976, but see the alternative presented by Bach and Cooper 1978). Put another way, the problem with trying to use structure (7b) for the semantic analysis of restrictive modifiers is that the meaning of the phrase the student is not a part of the meaning of the phrase the student who was curious. Only by making the major subdivision between the and student who was curious can a uniform semantic treatment of the be maintained.

We have argued that (7c) is the right structure for restrictive modifiers, (7b) for non-restrictive modifiers. (7d) can be argued to be the right structure for arguments; we will not go through detailed arguments here, but note that (i) the non-recursivity of (7d) is an advantage for arguments, since unlike modifiers, arguments cannot be added ad infinitum; and (ii) the close association ("sisterhood") of the noun and the argument in (7d) make a good structural basis for the fact that the choice of noun governs whether and which arguments can occur. Baker (1978) also provides a nice argument to the effect that the "pro-form" one(s) that occurs in expressions like "a large green pencil and two small ones" acts as a "pro-common noun phrase", substituting for any well-formed N′ expression, but not for a noun without its arguments. So the distribution and interpretation of that use of one can also be used to discriminate between modifiers and arguments, and its behavior is

consistent with the correlation between structures and interpretations that we have argued for here (see Problem 2).

What about structure (7a)? It does not seem to be optimal for any of the constructions examined, and this, together with many other such cases, has led a number of researchers to propose that a great deal of syntactic and semantic structure is binary-branching. There is some suggestive evidence, but not unchallenged, that young children may have flatter structures like (7a), and a corresponding lack of sensitivity to some of the semantic distinctions discussed in Section 2.2. below which can be made only with nested binary-branching structures; see Matthei 1979, but also see Hamburger and Crain 1984, 1987.

The example shows that the requirement that semantic interpretation rules correspond to syntactic structure can put very strong constraints on syntactic analyses. But caution is needed with this claim, since without independent constraints on permissible syntactic analyses and permissible means of semantic combination, the compositionality constraint by itself would have no teeth: it would be possible to construct a syntactic analysis to support virtually any desired semantic analysis. That is why many formal semanticists view compositionality as a methodological principle or a working hypothesis rather than as a testable empirical hypothesis; see Janssen, 1983, Gamut 1991. Compositional semantic analysis is typically a matter of working backwards from intuitions about sentences' truth-conditions (the most concrete data we have, according to Cresswell's "Most Certain Principle"), reasoning our way among alternative hypotheses concerning (a) lexical meanings, (b) syntactic structure, and (c) modes of semantic composition. Choices of any one of those constrain choices among the others; some choices lead to dead ends or at least make things much harder, others survive. "Solutions" are rarely unique and almost never final, since in any argument we are examining some particular set of alternative hypotheses with a great many assumptions explicitly or implicitly held constant, and a new idea about any part of the syntax or semantics can affect the choices among existing alternatives or open up new alternatives for consideration.

## 2.2. The Semantics of Adjectives.

For the rest of this case study, let us focus on the restrictive modifier uses of adjectives and inquire further about their semantics. In the preceding section, we took the semantics of restrictive modification to be just set intersection, and in fact that analysis works perfectly well for all restrictive modifier uses of PPs and all restrictive relative clauses. It is also a simple and appealing hypothesis about the semantics of restrictive modifier uses of adjectives, or adjective-noun semantics for short, but we can argue, based largely on work by Parsons (1972), Kamp (1975), and Siegel (1976), that while it is adequate for many examples, it is not an adequate analysis for adjective-noun modification in general.

## 2.2.1. The Intersection Hypothesis.

In most of what follows, our attention will be mainly on phrases containing one adjective and one noun; a simplification in terminology may therefore be in order. The recursive nature of the modifier structure requires that the explicit rules combine ADJ with N' to form a new N'; the "bottom" occurrence of N' will consist of an N alone, as in the following tree.

```
        N
       / \
     ADJ  N'
         / \
       ADJ  N'
             |
             N
```
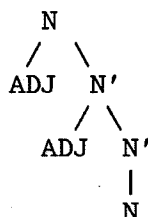
For simplicity, however, we may often speak of ADJ as combining with N, ignoring the intervening N' node. The explicit rule below is careful about using N', but subsequent informal references to it will often allude simply to intersecting the ADJ set with the N set.

Let us recast the intersection hypothesis in explicit rule form.

(10)  Given the syntactic configuration  $[_{N'}$ ADJ N'], the semantic interpretation of the whole is $||ADJ|| \cap ||N'||$

We illustrate with the example <u>carnivorous mammal</u>, from Kamp and Partee (forthcoming):

(11)  $||$carnivorous$||$        $= \{x \mid x$ is carnivorous$\}$
      $||$mammal$||$             $= \{x \mid x$ is a mammal$\}$
      $||$carnivorous mammal$||$ $= ||$carnivorous$|| \cap ||$mammal$||$
                                 $= \{x \mid x$ is carnivorous and $x$ is a mammal$\}$

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

## 2.2.2. Nonintersective Adjectives.

An adjective like <u>carnivorous</u> is intersective, in that (12) holds for any N.

(12)  $||$carnivorous N$|| = ||$carnivorous$|| \cap ||$N$||$

But what about an adjective like <u>skillful</u>? At first it seems all right to think of it as intersective; after all, a skillful carpenter is skillful and is a carpenter. But as Parsons (1968) pointed out, if the principle (12) were true with <u>skillful</u> substituted for <u>carnivorous</u>, then the inference shown in (13) should be valid, which it clearly is not.

(13)  Francis is a skillful surgeon.
      Francis is a violinist.
      Therefore Francis is a skillful violinist.

So skillful isn't intersective; it does not simply pick out a set of individuals who are skillful "period", but rather carves out a subset of the set corresponding to the noun it combines with. We will say more about how this might be achieved in section 2.2.4 below.

Since skillful does obey the principle exemplified in (14), which says that the set of skillful surgeons is a subset of the set of surgeons, it is called a "subsective adjective". (Note that the intersective adjectives are all subsective; subsectivity is a weaker property than intersectivity.)

(14)  Subsectivity: $||$skillful N$|| \subseteq ||$N$||$

Note that (14) doesn't give the meaning of skillful nor of the combination skillful + N; it merely specifies one property of the meaning of skillful (and of many other adjectives), a property which might be considered important enough to mark as a "semantic feature" in the lexicon. The principle in (14) can then be thought of as the semantic content, in model-theoretic terms, of such a semantic feature. Such principles or constraints are sometimes called meaning postulates in the philosophical and formal semantic literature. We will say more about meaning postulates in Section 2.2.4.

## 2.2.3. Nonsubsective Adjectives.

But are all adjectives at least subsective? The reader might start trying to think of some that are not. These should be adjectives that fail both the condition of intersectivity and the weaker condition of subsectivity; it should be possible for something to be an "ADJ N" without being an "N".

Here are some: the adjectives former, alleged, counterfeit are neither intersective nor subsective.

(15)      (a)  $||$former senator$|| \neq ||$former$|| \cap ||$senator$||$
          (b)  $||$former senator$|| \not\subseteq ||$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 (15b) asserts, it isn't even true that the set of former senators is a subset of the set of senators. Among the nonsubsective adjectives we might further distinguish a subclass of "privative" adjectives, those for which an instance of the ADJ + N combination is never an instance of the N alone. (See Problem 4.) Counterfeit is of this sort, while alleged is not, since an alleged murderer, for instance, may or may not be a murderer. For some nonsubsective adjectives it is not completely clear whether they are privative or not; and the answer may be one that is dependent on the context and the domain of discourse. Is a fake gun a gun? Can it be? If a fake gun is necessarily not a gun, how can one make sense of a question like "Is that gun real or fake?" (See Thought Question 3.) Readers who have thought of some additional examples of possibly non-subsective adjectives might compare notes

and see if they agree about which of them indeed are non-subsective, and about which of them are furthermore privative. (See Problem 3.)

The semantic properties of adjectives that we have examined so far have been chosen because they provide increasingly strong counterexamples to the intersection hypothesis. They are by no means the only interesting semantic properties of adjectives, and we will consider some others in later sections. First let's think about what kind of semantic interpretation of adjectives can do justice to the range of ADJ + N combinations we have just seen.

### 2.2.4. Adjectives as Functions.

Parsons (1968), Montague (1970), and others argued that the simplest rule for the interpretation of ADJ + N 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 they combine with onto the value of the ADJ + N combination. And from the evidence of the nonintersective and nonsubsective cases they argued that the relevant semantic values must be _properties_ rather than _sets_, i.e. must be _intensions_ rather than _extensions._ Let's look at some of the arguments for this analysis, introducing the key concepts as we need them.

First of all, we already have in (13) the germ of the argument that we can't do justice to the semantics of ADJ + N if the semantic value of the N is just the set that it denotes (its "extension" in the given state of affairs.) Let's just extend (13) a bit by considering a possible state of affairs in which every surgeon is a violinist and every violinist is a surgeon: i.e., in which the set of surgeons IS the set of violinists. Then no matter what function we were to take the adjective _skillful_ to denote, if it has to apply to the set denoted by the noun, then there is no way that we could get the semantic value of _skillful surgeon_ to come out different from the semantic value of _skillful violinist_. Why? Because we would have the same function applying to the same argument in both cases, necessarily giving us the same value. So we need to find semantic values for _surgeon_ and _violinist_ that can be different even when the sets denoted by those N's are the same.

The idea, which traces back to Frege (1892) and was further developed through the work of such philosophers as Carnap (1956), Hintikka (1969), Kripke (1963), and Montague (1970), is that every noun expresses a _property_, which we will call its _intension_; that property, together with the facts in a given state of affairs, determines what set the noun happens to denote (as its _extension_) in the given state of affairs. The _intension_ comes much closer than the _extension_ to what we ordinarily think of as the meaning of the noun; the intension is more like a characterization of what something would have to be like to count as instance of that noun. The nouns _unicorn_ and _centaur_ both have (presumably) the same extension in the actual world, namely the empty set: there aren't any of either. But they don't have the same meaning, and that correlates with the fact that there are fictional or mythical states of affairs where the two nouns have different extensions.

Intensions and extensions can be modelled using the notion of _possible world_ (possible situation or state of affairs, possible way things might be),

a notion which may be approached from various angles (see the collection Allén (ed.) 1989). Linguists working on formal semantics tend to view possible worlds as a formal tool for illuminating a certain kind of semantic structure, without necessarily taking a stand on the many deep philosophical issues that can be raised about them. It is worth noting, however, that some such notion is probably essential for an understanding of some very basic aspects of human cognition. Evidence for conceptualization of "other possible worlds" can be seen even at a pre-linguistic level in any child or animal that can show surprise, since surprise signals mismatch between a perceived state of affairs and an expected state of affairs. The notion of alternative possible worlds should therefore be understood not as a matter of science fiction but as a fundamental part of the ability to think about past, future, and ways things might be or might have been.

To a first approximation, we can take the extension of the predicate surgeon at a time $t$ in a possible world $w$ to be the set of things that have the property of being a surgeon 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 extensions; but if they differ in meaning they should differ in intension. Or more accurately, if they differ in truth-conditional aspects of meaning, they should differ in intension. Frege (1892) notes that there are things like "tone" or "emotional affect" that might also be ingredients of meaning in the broadest sense that have no effect on determining extension and are therefore not reflected in intension. Two terms differing only in "tone" or "connotation" or the like might therefore have the same intension but not be considered to have quite the same meaning.

Limiting our attention to truth-conditional aspects of meaning, the reasoning we have gone through suggests that we want to assign properties as the semantic values of nouns and other simple predicates. And it is commonly (though not universally) accepted that the property a given predicate stands for is completely determined by the "spectrum" of actual and possible extensions it has in different possible worlds. In other words, the property is completely identified by the function which assigns to each possible world $w$ the extension of the predicate in $w$. We therefore take such functions as our formal analysis of properties, and assign them in the lexicon as the intensions of simple predicate expressions such as common nouns.

If properties are identified with predicate intensions, then we can see how adjective meanings can be understood as functions from intensions to intensions. The distinctions among the various subtypes we have looked at in the preceding subsections can be characterized in terms of restrictions on the kinds of functions that are expressed by the different classes of adjectives. Formally, these restrictions can be expressed as meaning postulates; informally, we can think of classifications like "subsective" and "intersective" as semantic features on adjectives like skillful and

<u>carnivorous</u> respectively, cashed out as restrictions on the corresponding functions requiring them to obey restrictions analogous to the respective conditions (14) and (12) above.

The notion of meaning postulates is worth more discussion than we can give to it here. The central issue to which they are relevant is a long-standing debate concerning whether all lexical meanings can be fully analyzed via "lexical decomposition" into some sort of presumably universal semantic "atoms", representing basic or primitive concepts. That view, while appealing and recurrent, may well be too strong, and the notion of meaning postulates is offered as a technique for capturing significant generalizations about extractable regularities within lexical meanings without presupposing total decomposability. On both views, there are important regularities concerning semantic properties of lexical items that need to be captured; on the anti-decomposition view, some lexical items have idiosyncratic "residue" parts of their meanings which cannot reasonably be analyzed further into compositions of simpler parts. Meaning postulates can then express whatever regularities there are to be found without entailing that what can be said about the meaning of a given item with meaning postulates should be supposed to exhaust its meaning. Another, perhaps related, problem in the field is whether and how we can draw a line between information that belongs in the lexicon  and information that is part of our "world knowledge" or beliefs about the referents of our terms. This is an area where there are some clear cases and many unclear ones. Meaning postulates might be a helpful tool here as well, since they make the <u>form</u> of some kinds of lexical information no different in kind from the form of some kinds of general knowledge. That would make it possible to hypothesize that the very same "fact", e.g. that whales are mammals, could be in either of two "places", a storehouse of lexical knowledge or a storehouse of empirical knowledge, and whether it's part of the meaning of "whale" or not need not be fixed once and for all but could vary across different individuals or subpopulations or different historical times.

Now a further issue concerns generality. We seem to have two choices. We can, as suggested above, treat <u>all</u> adjectives as functions from intensions to intensions, with the subclasses like subsective and intersective defined by meaning postulates that are satisfied by the corresponding functions in those cases. Alternatively, could we perhaps say that different adjectives have different sorts of semantic values: the intersective ones are simply predicates whose extensions are sets and whose intensions are properties, and only the non-intersective ones have to be interpreted as functions from intensions to intensions. Compositionality in its strictest form requires a single semantic type for each syntactic category: therefore this second strategy would require that adjectives also be divided into distinct syntactic subcategories corresponding to these different semantic types.

If we do want a single semantic type for all adjectives, that can only be achieved if all adjectives are treated as functions from intensions to intensions: it is possible, for instance, to treat an intersective adjective like <u>carnivorous</u> 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 adjectives like

<u>carnivorous</u> as the general case, treat adjectives in general as simple predicates and take the interpretation of the ADJ + N rule to be set intersection, because, as we argued earlier, there is no comparable way to treat <u>former</u> or <u>skillful</u> as a restricted subcase of that.

In Section 4 we will introduce the concept of "type-shifting" and explore the possibility of assigning to each category a family of types rather than a single type, with a slight weakening in the principle of compositionality. From that perspective we will be able to say that the interpretation of adjectives as functions from functions is the one type of interpretation that all adjectives can have, while some adjectives also have simpler meanings, such as interpretations as simple predicates expressing properties and denoting sets. We will appeal to that possibility at various points in what follows, but will defer a fuller discussion to Section 4.

The reader may be concerned that we have said nothing about the fact that adjectives also occur in predicate positions, typically after <u>be</u> or <u>become</u>; it turns out on examination that it is normally only the extensional, intersective adjectives that occur in predicate position, where the simple interpretation of adjectives as denoting sets and expressing properties suffices. Siegel (1976a,b) discusses both this generalization and its apparent exceptions, and argues for a corresponding syntactic distinction in subtypes of adjectives, noting that some languages, such as Russian, even have different morphological forms for predicate and function interpretations of adjectives. Of course there is much more to say about the different uses of adjectives and the connections among them, but we will continue to restrict our attention principally to adjectives occurring as modifiers of nouns.

Let's sum up what we've concluded from this case study. We wanted to figure out what sorts of things adjective meanings might be. Following Lewis's Advice, we asked what we wanted adjective meanings to "do", what sorts of properties they needed to have. To get at that question, we focussed on the principle of compositionality and on the truth-conditions of sentences containing adjective-noun combinations. Assuming that a common noun (phrase) (N or N') denotes a set (i.e. has a set as its extension), we saw that the simplest case of restrictive modification could be analyzed as set intersection. But for the most general case we needed something more than sets and set intersection, and that motivated the notion of intensions: N's express properties (have properties as their intensions), restrictive modifiers are most generally interpreted as functions from properties to properties, and restrictive modification is most generally understood as function-argument application. Within that general analysis, the intersective adjectives can still be singled out as a natural simplest subclass, characterizable as such by meaning postulates (or on a type-shifting approach, as will be discussed in Section 4, treated as having predicative interpretations as their simplest interpretations and as having derivative interpretations as functions via general type-shifting principles.)

## 3. <u>CASE STUDY 2: Vagueness, Context-Dependence, and Point-of-View.</u>

### 3.1. <u>Vagueness and Context-Dependence.</u>

In section 2.2.2. above we indicated that the inference pattern (13) was a test of whether an adjective was intersective. By this test, it looks like vague adjectives like <u>tall</u> and <u>young</u> are non-intersective:

(16)   a. Win is a tall 14-year-old.
        b. <u>Win is a basketball player.</u>
        c. ?? Therefore Win is a tall basketball player.

The inference in (16) seems to be invalid, just like that in (13), so by that test it would seem that <u>tall</u> is non-intersective. But <u>tall</u> is a vague adjective, and maybe it's interpreted differently in the premise (16a) and the conclusion (16c); and if so, maybe this test isn't conclusive. When an adjective is vague, like <u>tall</u> or <u>young</u> or <u>blue</u>, the lexical semantics isn't simply a matter of classifying entities into those that are tall and those that are not. Intuitively, it seems that for vague adjectives, there may be some clear positive cases and/or some clear negative cases, and there are unclear cases, or a "range of indeterminacy", in between. For some vague adjectives, like <u>blue</u> and <u>round</u>, there are some "absolute" clear cases, which seems to correlate with the fact that these adjectives can sensibly occur with modifiers like <u>perfectly</u> and <u>absolutely</u>, whereas for others, like <u>big</u> and <u>old</u>, there may be no unequivocally clear cases, and only relative to a given context can we identify any cases as clearly positive or negative. But for either kind of vague adjective, context clearly plays a major role in how we reduce or eliminate the range of indeterminacy. Context will lead us to draw the lines in particular ways for what will count as definite positive instances (the "positive extension") and what will count as definite negative cases (the "negative extension"); some contexts will eliminate indeterminacy, others may leave a (smaller) range of indeterminacy.

So let's think about how vagueness and the influence of context might affect what we should conclude from our judgements about (16). Perhaps <u>tall</u> appears to fail the test of (16) simply by virtue of the influence of the noun on the context. As Kamp (1975) suggested and Siegel (1976b) argued at length, "relative" adjectives like <u>tall</u>, <u>heavy</u>, and <u>old</u> 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 ADJ + N construction. Why is this evidence against treating <u>tall</u> as a nonintersective adjective like <u>skillful</u>, interpreted as a function applying to the noun's intension? Well, here's an argument. In (17) below, we see another example of the influence of contextual cues on the (partial) resolution of the vagueness of <u>tall</u>, only this time we have the same noun <u>snowman</u> in each case.

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

Further evidence that there is a difference between truly nonintersective subsective adjectives like <u>skillful</u> and intersective but vague and context-

dependent adjectives like _tall_ was noted by Siegel (1976b): the former occur with _as_-phrases, as in _skillful as a surgeon_, whereas the latter take _for_-phrases to indicate comparison class: _tall for an East coast mountain_. (An adjective can be nonintersective and **also** vague, and then one can use both an _as_-phrase and a _for_-phrase: _very good as a diagnostician for someone with so little experience_.)

It is both difficult and important to sort out the effects of context-dependence on the interpretation of different sorts of adjectives and nouns occurring alone or in combination. 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 nonintersective (almost certainly including _former_ and probably including _skillful_.) There may be many disputable or borderline cases and there may be cases which involve homonymous or polysemous doublets (as suggested by Siegel 1976b for _clever_, which she analyzes as having one vague but intersective reading meaning having to do with a general evaluation of persons and another non-intersective (also vague) reading that is a function applying to role-properties like _experimenter_ or _magician_ or _con artist_.)

Vagueness and context-dependence are in principle independent properties, although they often co-occur. There are adjectives like _left_ and _right_ which are context-dependent but not (very) vague, and we will say more about context-dependence of that kind in section 3.2; and there are nouns like _vegetable_ and _bush_ which are vague but not (very) context-dependent. Even the line between vague and non-vague predicates is vague; a concept may count as sharp for most purposes but vague relative to the demands of scientific or legal or philosophical argument. Probably almost every predicate is both vague and context-dependent to some degree.

Let's go back to our examples of vague adjectives and think about how to analyze their meanings. The existence of vague intersective adjectives like _tall_ poses a puzzle for the semantic account we gave in section 2. We argued that for the intersective adjectives, the semantic interpretation of ADJ + N could be analyzed as set intersection, and we noted that there were various ways to maintain that claim even when we generalized ADJ + N semantics to function-argument application, either via meaning postulates or via the type-shifting analysis to be discussed in Section 4. So one way or another, we can still regard a simple intersective adjective as denoting a set which gets intersected with the noun set. But sets are sharply defined: any given element either is or is not an element of the set. How can we model the vagueness we've just been talking about? How can the intersective adjective _tall_ denote a set?

We probably need to complicate our theory a little. Let's go back to Lewis's Advice and think about what the meaning of _tall_ does, and find something that does that. We've already seen some aspects of the behavior of _tall_, in thinking about tall 14-year-olds vs. tall basketball players, and in the examples about the tall snowmen. And we've noted that such adjectives generally have in any given context a positive extension, a negative extension, and often still some range of indeterminacy, and these vary from

context to context somehow. Let's see if we can find out some more about how this variation works.

Here's an example in which we have a pair of words that can both be used either as nouns or as adjectives, with virtually the same meaning in either use, and we find a striking effect depending on which is which.

(18)  a. Bobo is a giant and a midget.
      b. Bobo is a midget giant.
      c. Bobo is a giant midget.

It seems 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 Bobo a compound concept with the same pair of constituent parts. So there seems to be some interesting effect of syntax here that affects how the resolution of the vagueness of these predicates is (partially) resolved.

In the case of (18a), with overt conjunction, the sentence is generally interpreted as contradictory, unless one can find grounds for imposing an interpretation that implicitly adds different "respects" to the two, e.g. a mental giant and a physical midget. 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 (18a), 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 (and in a given respect) _ipso facto_ does _not_ count as a midget relative to that same context.

In (18b) and (18c), on the other hand, one predicate serves as head noun and the other as modifier, and the difference in interpretation is striking. Everyone seems to agree that a giant midget must be an unusually large midget, and a midget giant an unusually small giant. This tells us that the predicate serving as head noun 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 the modifier appears to be "recalibrated" in such a way as to make distinctions _within_ the class of possible referents for the head noun. So whereas in (18a), _giant_ and _midget_ are normally construed as mutually exclusive categories, in both (18b) and (18c) the modifier-head construction virtually forces us 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.

If we set up our theory so that it accommodates such adjustments to the interpretation of the modifier, then there will be no obstacle to maintaining the interpretation of the modifier-head construction as predicate conjunction (set intersection)[3]; but without such an adjustment or "recalibration", a conjunction interpretation of the construction would lead to the false prediction that (18a-c) should all have the same interpretation. Combinations such as _tall tree_ should presumably be handled similarly: _tall_ is a vague term whose interpretation is affectd by both the linguistic and the non-linguistic

context as illustrated above, but once the interpretation is specified, as something roughly like "at least $\underline{d}$ tall" for some degree of height $\underline{d}$, the combination <u>tall tree</u> can be treated as simple set intersection.

Let's try to articulate some of the kinds of principles that may govern the dynamics of context effects with vague linguistic expressions. These tentative hypotheses, from Kamp and Partee (forthcoming), may be viewed as an invitation to more systematic exploration of a relatively undeveloped field of study which may cast 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.

Two principles suggested by the examples we have just considered are the following:

> (19) **Parallel Structure Effect**: In a (syntactically) conjoined structure each conjunct is interpreted in a parallel way relative to their common context.

> (20) **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.

In the simplest cases of the application of the Head Primacy Principle, the effect of the interpretation of the head noun will be to restrict the "local context" for the modifier to the positive extension of the head; so that one will for instance be interpeting the modifier "giant" in (18c) in a local context which consists only of midgets.

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:

> (21) **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 <u>midget giant</u> example, for instance, the Head Primacy Principle and the Non-Vacuity Principle cooperate to produce the observed results: we first interpret the head <u>giant</u> in the given context (e.g. the fairy-tale world of Jack and the Beanstalk) in such a way as to give <u>giant</u> both a positive and a negative extension in the domain of the context; then we interpret <u>midget</u> in such a way that it has both a positive and a negative extension within the positive extension of <u>giant</u>. This of course requires a very different "calibration" of <u>midget</u> than would be appropriate for the global context, since <u>midget</u> and <u>giant</u> are incompatible relative to one and the same context.

In the <u>giant and midget</u> example, (18a), 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 <u>giant midget</u> and <u>midget giant</u> cases, it is the positive extension of the modifier that has to be stretched to satisfy the <u>Non-Vacuity Principle</u>. In other cases the same principle may lead us to shrink the positive extension and expand the negative extension: see Problem 5.

In still other cases, the Non-Vacuity Principle seems to override the Head Primacy Principle. Consider the phrase <u>stone lion</u>: Is a stone lion a lion? With respect to the normal interpretation of the predicate <u>lion</u> the answer would seem to have to be "no": stone lions fail both scientific and everyday tests for lionhood, would never get counted in a census of the world lion population, etc. Yet if we have to stick to normal interpretation, there can't be any stone lions; lions are not made of stone, nor do we seem inclined to try to stretch the predicate <u>stone</u> to apply truly to lion-flesh. But <u>stone lion</u> is also not just an idiom: any name of a material can be substituted for <u>stone</u>, familiar or novel (glass, chocolate, velveteen, ...), and just about any concrete noun can be substituted for <u>lion</u>. In this case it seems that the Non-Vacuity Principle overrides the Head Primacy Principle and leads us to reconstrue the head noun (as "representation of a lion" or something like that) so that the modifier can have a positive extension within the positive extension of the head noun. This is more like a case of meaning-shift (see Section 4.1 below) than straight vagueness resolution, since we are moving things into the positive extension of <u>lion</u> which are ordinarily clear negative cases.

It may also be the Non-Vacuity Principle or some generalization of it that makes us so strongly inclined to reinterpret sentences which by their form and the Parallel Structure Effect should be tautologies and contradictions, like (22) and (23). It's likely that the hearer interprets (22) not as a tautology but as a denial of the existence of a range of indeterminacy for <u>new</u>, and interprets (23) not as a contradiction but with different "respects" implicitly filled in, much as one might do to make sense of the <u>midget and giant</u> case (see Kamp and Partee (forthcoming) for discussion.)

(22) Either it's new or it's not new.
(23) Well, he's smart and he's not smart.

An explicit formalization of a compositional semantics of vague adjectives and their combinations with nouns is beyond the bounds of this chapter; there are several proposals in the recent literature, and this is still very much an open research area. The main point of this subsection has been to illustrate the strategy for working towards such an analysis, and to show how there is constant interplay between hypotheses about the meanings of individual words and hypotheses about the compositional semantics of various syntactic constructions. A good analysis of vagueness may allow us to maintain

the simple compositional rule of set intersection for many more modifiers than
we could if we didn't disentangle the context-dependent aspects of vagueness
from the semantics of adjective-noun combination. Progress in semantics, as in
many other fields, often depends on finding ways to understand complex
phenomena as the result of interactions of several relatively simple
principles.

It is also worth noting as we close this discussion of vagueness that as
one studies how vagueness works in more detail, one quickly overcomes the
common prejudice that vagueness is always a bad thing, that it is some kind of
"defect" of natural language. Not only is it impossible for a natural language
to get along without vague expressions (the reader is invited to argue this
point in Question 4 for Further Thought), but natural languages provide the
means for reducing vagueness where needed, both by the implicit interactions
with context discussed above, and by explicit stipulation. And as we have
taken some initial steps towards showing, vagueness is no obstacle to formal
modelling; a theory of the semantics of vague terms can be a perfectly precise
theory. And as a practical goal, if we can understand vagueness better, we may
be able to develop richer and more versatile computer languages and more user-
friendly human-machine interfaces.

## 3.2.  Context-dependence and point-of-view.

In the previous section we saw one aspect of the context-dependence of
natural language semantics, in the effect of context on vagueness reduction
and vagueness resolution. In this section we will look at another dimension of
context-dependence, one which shows up in clearest form when we think about
the meanings of "demonstrative" or "deictic expressions like this, that, and
there, and "indexical" expressions like I, here, today, and now.[4] Learning the
meaning of a demonstrative or indexical is not like learning the reference of
a proper name, since the reference of that or I changes from occurrence to
occurrence; nor does it mean simply learning an associated property, as we do
for the nouns and intersective adjectives we have looked at so far. Anything
could be the referent of a use of the word that: a boiling teakettle, a
number, a color we have no name for, the expression on a baby's face; what
that refers to in a given utterance depends on the particular intentions of
the speaker on that particular occasion of utterance. Few demonstratives and
indexicals are quite as wide-ranging as this and that, but even more narrowly
delimited ones like yesterday and today depend for their reference on the
context of the given utterance, in this case on the time at which the
utterance occurs,

Following Lewis's Advice again, it is generally agreed that the meanings
of demonstratives and indexicals are best thought of as functions from
contexts to referents.[5] So meanings have become more complex in another
dimension. In general, whenever we discover that the reference of some
expression depends on X, Y, and Z, we model meanings as functions from Xs, Ys,
and Zs to referents. (And then when working on case studies in which Y and Z
play no role, we may simplify and treat meanings as just functions from X's.)

Now an interesting phenomenon that linguists have studied is that there
are quite a number of open-class lexical items like nouns, verbs, and

adjectives that show mixed properties: their meanings involve a combination of ordinary properties with indexical or demonstrative-like behavior. Fillmore (1971, 1975) opened up this area with his classic studies of come and go, take and bring, and in front of, which involve interesting interplays of factors such as the position and orientation of the speaker and of the hearer, direction of motion, if any, and for the case of in front of, the question of whether the object in question has an intrinsic front or not: "in front of the car" is more ambiguous than "in front of the tree" (though even trees may derivatively acquire fronts and backs if closely associated with things like houses that have an intrinsic front). (See Problem 6 and Questions 5 and 6 For Further Thought.) The difference between "Please come out to play!" and "Please go out to play!" is not a difference in the action being requested, but a difference in the point of view of the utterer - we know which of these is said by a neighbor kid on the doorstep and which by a parent in the house.

Let's continue looking at adjectives and see what sorts of demonstrative and indexical-like behavior we can find among them. Sample candidates include right, left, nearby, far, ahead, behind, close, closest, local, foreign, but the reader can surely find or think up many more (like vagueness, context-dependence of this kind seems to be almost everywhere once you look for it.) Let's look more closely at right and left, since they illustrate two different points about this family of words. (And more besides; we won't try to go into all the interestingly different properties these words show in their different uses such as "on the right", "to the right", "to my right", "turn right", "turn toward the right", etc.; we will just look at one tip of this particular iceberg.) In the first place, there is one kind of context-dependence that just seems to be elliptical for an of-phrase: the context-dependent (24a) may be viewed as elliptical for the (more nearly) context-independent (24b).

(24a) Sew the team insignia to the right.
(24b) Sew the team insignia to the right of the "A".

But as the author realized when the instruction (24b) was received in a packet with her son's Little League baseball cap (the "A" was on the front, in the middle), (24b) still contains a crucial point-of-view ambiguity: to the right of the "A" as you look at it, or to the right of the "A" from the point of view of the person wearing the cap? That is, there is another context-dependent aspect of the meaning of right and left, much like that of come and go, which has to do with point of view and has no simple non-elliptical variant (although you may be able express the point of view with some added phrase like "from the point of view of so-and-so").

As words like foreign show, context-dependence and point of view are not limited to spatio-temporal aspects of the context. Who's foreign depends not simply on where you are but on political boundaries and things like citizenship; when I'm in your country, it's me that's foreign (although a parochial tourist may not realize that and may remark on the experience of being surrounded by foreigners). And also worth noting is that the relevant context isn't only the situation in which the utterance occurs: utterance-internal expressions can also affect the local context for the interpretation of successive expressions, much as we found the interpretation of vague adjectives to be affected by the local context established by their head

nouns. How many interpretations can you think of for (25)?

(25) Most foreigners speak a foreign language.

On one reading of (25), as spoken, say, by an American, (25) says that most
people from other countries than America speak a language other than their own
native language; this might be said by someone arguing for more teaching of
foreign languages in the American schools. So on this reading, which is just
one of several, the first occurence of foreigner is "anchored" to the
utterance context ("foreigner relative to this country"), while the later
occurrence of foreign is anchored to that earlier occurrence of foreigner: "a
language foreign to that foreigner".[6] (See Problem 7.)

So what are contexts?  What does a theory have to look like to model the
context-dependence of meaning?  Some of the simplest examples might suggest
that a context can be represented as an n-tuple of a speaker, a place of
utterance, and a time of utterance.  But we have already seen examples that
show that there is no such simple enumeration specifiable once and for all of
aspects of contexts that may play a role in context-dependence (see Cresswell
1973 for a classic discussion of this point).  Trying to specify all the
ingredients of "contexts" is as fruitless as trying to specify the ingredients
of possible worlds or possible situations that go into the analysis of
properties and propositions and intensions in general.  In any particular
discourse situation, or in any sample model or particular computational
application, it may well be feasible and desirable to reduce possible worlds
and possible contexts to "state-descriptions" in terms of values on some
specified set of parameters.  But for the workings of a natural language in an
uncircumscribed context, all we can expect to do is specify that an analysis
must be in terms of a set of possible worlds or situations W and a set of
possible contexts C, leaving the elements of W and C as primitives in the most
general case.

So the extension of an adjective like foreign relative to a given
context (providing a point of view) and a given state of affairs (a
specification of "the facts") will be a set of individuals (persons and
objects and languages and more - we must take the notion of "individual" very
broadly). The intension of foreign relative to a given context is a function
that maps each possible state of affairs onto the corresponding extension --
two different states of affairs that differ in Vladya's citizenship may result
in two different answers to whether Vladya is in the extension of foreign. And
what we might optimistically call the "meaning" of foreign (until Lewis's
Advice tells us to go searching for further ingredients) is a function from
contexts to the relevant intension: the context has to tell us whose point of
view establishes what is to count as foreign, and then the intension encodes
how the extension depends on the facts.

Context-dependence of this sort, like vagueness, is one of the things
that makes natural languages so versatile. User-friendly computer programs
have learned to exploit context-dependence: the same short command may have
different effects depending on the local context in which it is used. "Exit"
or "Quit" are typically context-dependent (in fact "indexical") commands,
meaning exit from "this", from whatever process you are currently in. By the

same token, context-dependence can be a source of misunderstandings (both in human-human and human-computer communication) that may be hard to pin down because of the usually implicit nature of the relevant context-dependent aspects of meaning. (In case of the Little League sewing assignment mentioned earlier, I was curious to see what other parents (well, back then, other mothers) would do. It turned out 50-50; but I suspect that relatively few people noticed that there was an ambiguity. You could do experiments to explore this kind of phenomenon.)

Context-dependence also provides us with an opportunity to think about the role of simplifying assumptions in ordinary uses of language as well as in science. We noted above that _right_ and _left_ are indeterminate without a point of view. In contrast, _up_ and _down_ are context-independent. Ah, but that's only because of gravity; what if you're in a gravity-free environment? Then _up_ and _down_ are just as much in need of a designated point of view as _right_ and _left_ are. Since gravity is a universal of terrestrial life it is not surprising that languages evolved with gravity presupposed in the interpretation of some "vertical" prepositions and adverbs; this is a good example of a language universal that should probably not be attributed to the language faculty. (As a relevant side note, it reportedly turns out that astronauts prefer to have a "designated floor" and "designated ceiling" to prevent disorientation.)

### 3.3  Compounds vs modifiers and the limits of compositionality.

The remarks made so far about ADJ + NOUN combinations are intended to apply to all cases of modifier-head constructions, including cases where a noun is converted into an adjective and used to modify another noun, as in _stone lion_, _oak table_, _cardboard box_.  But they are not intended to apply to compounds, either of the noun-noun or adjective-noun variety. Compounds in English can generally be recognized by their heavier stress on the first word; see (26), in which heavier stress is indicated by capitalization.

> (26)  Modifier-head (compositional)
>        (a) black BOARD
>        (b) brick FACTORY
>        (c) toy STORE
>
>        Compound (idiomatic)
>        (d) BLACK board
>        (e) BRICK factory
>        (f) TOY store

The contrast in the case of adjective-noun combinations, as in (26a) vs. (26d), 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 nonsubsective, although this is debatable (see the discussions of _fake gun_ in Section 2.2.3 and _stone lion_ in Section 3.1.) In compounds, on the other hand, there is no general rule for predicting the interpretation of the combination, intersective or otherwise.[7] A TOY store (in typical contexts) is a store that sells toys, a TOY box is a box that holds

toys, etc. Semanticists in general do not expect a semantic theory to provide a compositional semantics for compounds but do expect a compositional semantics for modifier-head construction. The reasoning is that a native speaker cannot generally interpret a novel compound on first hearing on the basis of knowledge of the language alone, but can do so for a novel modifier-noun construction.

It is an interesting question just what a semantic theory should say about compounds, but one that goes beyond the scope of this discussion. Presumably the semantics for English must at least tell us that the syntactic head is also the semantic head in the sense that a BRICK factory is a kind of factory, not a kind of brick. One of the challenging parts of the problem is how to articulate the interface between linguistic and non-linguistic contributions to interpretation so that the semantics specifies that this a place where non-linguistic knowledge has to fill in some "relevant" property that "saliently" involves the first element of the compound to form a modifier or specifier of the second element. The semantic constraints are extremely weak and general; there seems to be no limit in principle on how that inferred property is related to that first word. This very absence of semantic constraints, however, results in a fairly strong pragmatic constraint: for a novel compound to be understood on first hearing, there must be a unique most salient or plausible interpretation, so very far-fetched possibilities will only be usable in very "rich" contexts. (Silent exercise: construct three different possible interpretations for the compound "computer puzzle" and imagine scenarios in which each would be readily understood as the intended interpretation. See also Problem 8.) The existence of compounds, like the existence of vagueness and context-dependence, appears to exploit the cognitive capacities of language users in ways that allow natural languages to be much more flexible than we can allow computer languages or other formal languages to be. In the past, it was commonplace for logicians and others to criticize natural languages for their "sloppiness" or "messiness" and to take pains to avoid properties like vagueness in the design of formal languages; but as these aspects of natural languages come to be better understood and appreciated, we are beginning to understand that that may be a little like criticizing living organisms for being messier than machines. Formal languages, like machines, can be extremely useful tools, but natural languages, like people, don't have to apologize for not being more like them (respectively.)

In this section we have seen several aspects of the interaction of meaning with context. Some word meanings, those we have characterized as being wholly or partly indexical or demonstrative, involve instructions for fixing aspects of their reference as a function of the context. The relevant context may be the situation in which the utterance occurs, or it may be a more abstractly constructed context, built up by the speaker and hearer together on the basis of the evolving discourse. We have seen examples where we have to take into account very "local" contexts internal to sentences: the interpretation of one expression in a sentence can give rise to context-dependent effects in the interpretation of other expressions in the same sentence. The resolution (complete or partial) of vagueness is a very widespread case of context-dependence of meaning, and we scratched the surface of what is still a wide-open domain of inquiry in posing some tentative

hypotheses concerning the various linguistic and cognitive principles that interact in influencing how we "recalibrate" vague adjectives depending on the context in which they occur. And in this last subsection, we have seen that compound constructions represent an extreme of context-dependence -- the rules of language provide only some very general constraints on the interpretation of a compound, with inferences from context and plausibility bearing the bulk of the load, and with concomitant benefits and costs in flexibility of interpretation and potential for failure of interpretation.

## 4. CASE STUDY 3: MEANING SHIFTS

We closed the previous section with some thoughts about constructions in which the semantics only constrains but does not completely determine the meaning and a great deal is left to plausible inference from context. Another domain in which we can observe some tension between the drive for uniformity of principles of interpretation and the drive for flexibility of interpretation is the domain of "meaning-shifting" principles. We will look at some examples of such principles at work in the domain of adjectives, and use these examples as a basis for some speculations about more broadly cognitive notions of "natural" functions and relations of which these meaning-shifting principles might be linguistic instantiations.

### 4.1. Type-Shifting and Meaning-Shifting Principles.

In the previous chapter, Exercise 2.3 asked you to work out how proper names could be treated as generalized quantifiers. In the first part of this chapter we analyzed definite descriptions as denoting individuals, but remarked that we could reanalyze them as generalized quantifiers by the same means. Compositionality seems to demand that for each syntactic category there be one and only one corresponding semantic type. And since the only type general enough to accommodate all kinds of NP meanings is the type of generalized quantifiers, uniform compositionality leads us to analyze all NPs as generalized quantifiers, including those which otherwise could have been analyzed in a simpler way, e.g. as denoting an individual.

In section 2 of this chapter, we saw something similar with adjective meanings. The simplest adjectives could be adequately analyzed as denoting sets (and expressing properties), with set intersection for the semantics of adjective-noun commbination. But we argued that that wouldn't do in the general case because of the nonintersective and nonsubsective adjectives. For uniformity of adjective semantics we have to "generalize to the hardest case" and treat all adjectives as functions from intensions of common noun phrases to intensions of common noun phrases, for nothing simpler will work for adjectives like _former_ and _alleged_. That means even the simple intersective adjectives also have to be analyzed as such functions, and we described how meaning postulates could be used to capture the fact that the corresonding functions in that case would be highly constrained ones which in effect would mimic set intersection.

There is another approach to this situation, one which slightly weakens the form of the compositionality constraint but allows us to do more justice to the intuition that the "simpler" cases, like the proper names and the intersective adjectives, are indeed simpler. On this approach, known as "type-shifting", each syntactic category is associated with a set of semantic types rather than a single uniform type. Each lexical item is entered in the lexicon in its simplest type, and there should be 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 (with a set as extension and a property as intension), 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 are required. (See Partee (1987), Partee and Rooth (1983) for details in the case of NPs and verbs; the discussion of adjectives here is modelled on those treatments but without developing formal details.)

In the case of adjectives, we can take this approach a step farther and let the syntactic rules themselves have a family of interpretation rules rather than a single uniform one, and following what is called "Type- Driven Translation" (Klein and Sag 1985), let the choice of combining rule itself depend on the semantic properties of the parts. Let's illustrate.

Consider again examples (11) and (15):

(11)   carnivorous mammal
(15)   former senator

We are still assuming that (ordinary) nouns are interpreted as one-place predicates, that their extension is a set of individuals, their intension a property. Now since _carnivorous_ is a simple intersective adjective, its simplest interpretation is also as a one-place predicate. And it turns out to be a general semantic principle that any modifier-head construction in which the modifier and head are both one-place predicates is interpreted as set intersection.[8] In that case, not only can we keep our first hypothesis about the semantics of the intersective adjectives, but we would not even need to formulate a type-shifting rule to turn them into functions when they occur prenominally; the simple intersection rule that we gave in section 2.2.1. would in fact be the applicable rule without having to be "stipulated" as such.

In the case of (15), we would have no such option. _Former_, as we remarked earlier, does not have any simple predicative interpretation: we can't classify entities in a given domain into those that are former and those that are not. This correlates with the fact that _former_ does not occur in predicate position; we can't normally say sentences like (27) (except in poetry, advertising, and jokes, where creativity beyond the existing bounds of grammar may be welcome.)

(27)   ##Some of my best friends are former.

So the only available meaning for *former* is as a function from intensions to intensions. Type-driven translation will predict in that case that the semantic principle by which the meaning of (15) is derived is application of the function denoted by *former* to the intension of the noun *senator*. Note that besides the different types of adjective meanings involved here, we are implicitly bringing in the assumption that sometimes a noun contributes its extension to the meaning of the whole, sometimes its intension. That is also all right as long as there are principles which predict which it is in any given case.

Going back to the case of (11), suppose instead that we wanted to say that modifier-head semantics was always function-argument application. (One reason for such a rejection of set-intersection interpretation of ADJ + N combinations, even in the case of simple intersective adjectives, could be the evidence given in the discussion of vagueness in Section 3.1. that there is an asymmetry in modifier-head constructions that is not present in conjoined constructions, and set intersection is basically a formalization of predicate conjunction.) Then what we might say about *carnivorous* could be as follows. Its simplest interpretation is still as a one-place predicate, and it gets that meaning when it occurs in predicate position, as in "is carnivorous". But when it occurs in prenominal position as in (11), it is automatically reinterpreted by means of the following general principle:

(28)   **Predicate-to-Prenominal Shift:**
       If an ADJ has an interpretation as denoting a set $S_{ADJ}$, then that
       ADJ also has a possible interpretation as a function applying to a
       set, namely as the function $F_{ADJ}$ such that $F_{ADJ}(S_N) = S_N \cap S_{ADJ}$.

Note that when we pack set intersection into the shifted "function-type" meaning of the adjective itself, as we did in (28), the semantic interpretation principle for combining the shifted adjective and a noun must then be function-argument application. By distinguishing function from argument but defining the function in terms of set intersection, one can capture the modifier-head asymmetry needed for the Head Primacy Principle while still asserting that the interpretation involves set intersection.

The two choices we have illustrated with type-driven translation of (11) on the one hand and the meaning-shifting principle in (28) on the other hand represent two alternative strategies which are in principle available in this case and which would have the same effects except possibly for the issue of modifier-head asymmetries. Further investigation of such possibilities in a broader context of principles of grammar and processing should help to determine which of these accounts is closer to correct.

There are other kinds of meaning shifts that are common and familiar, some involving shifts in type (e.g. from one-place predicate to two-place relation or function, from entity to generalized quantifier, etc.) and some not. Of course individual lexical items can acquire shifted meanings through all sorts of idiosyncratic routes, and not all individual instances of meaning shifts reflect any general principles. But many do; we noted earlier that any English concrete noun X can turn into a modifier with the meaning "made of X"

(this is not true in all languages; many languages require a prepositional phrase like "of N" to express material.) Another common shift is a shift of a concrete count noun to a corresponding mass noun denoting the stuff that the count noun is made of: a potato, some potato on my plate; an egg, some egg on my plate. While not all count nouns have mass noun counterparts in common use, David Lewis has argued that that is just because we don't normally have occasion to use them all; he invites us to consider a "Universal Grinder": put a chair in one end, turn the crank, and there will be chair all over the floor. There are also familiar shifts of noun meanings to verb meanings of various sorts illustrate by verbs like "to can", "to dust" (this one can mean either to put dust on or to take dust off; as in the case of compounds and in ambiguity resolution, plausibility in a given context plays a large role in selecting among available meanings), "to staple". The reader is invited to explore the extent to which there are predictable subclasses of noun-to-verb meaning shifts and the extent to which they are predictable from properties of the subclasses of nouns involved. (Marchand 1960 is a classic source of data.)

A given kind of meaning shift may be productive or semiproductive in one language and not productive at all in another.[9] One minor pattern which seems to be semiproductive in English but not in any other language known to the author is the shift of a certain class of nouns we may call "attribute nouns" that allows NPs they head to be used as predicates or modifiers without having to be put into PPs or changed into adjectival phrases. This pattern is illustrated by the uses of the noun "color" in examples (29a-c) and the noun "length" in (30 a-c), as discussed in Partee (1987).

(29)    (a) Blue is a nice color.
        (b) This shirt is a nice color.
        (c) Sandy bought a shirt that color yesterday.

(30)    (a) Is 6 feet 6 inches the same length as two meters?
        (b) Is your hair the same length as my hair?
        (c) A mahogany board that length would cost more than this table.

"Color" and "length" are abstract nouns that we may call "attribute" nouns, since the things that are most straightforwardly in their extensions are themselves properties, as in sentences (29a) and (30a), literal translations of which are well-formed in many languages. Native English speakers do not notice anything odd about the (b) and (c) sentences, but speakers of other languages find them initially surprising, since shirts aren't colors and hair and skirts aren't lengths. (Other languages often express (29b) by the equivalent of "is of a nice color" or "has a nice color", but in most languages it is nonsense to say that a shirt is itself a nice color, just as in English one doesn't say "Pat is an interesting occupation".) Of course as English speakers we are not asserting with (29b) that shirts can be colors in the way that blue is a color. Rather, English seems to have an idiosyncratic rule that shifts NPs headed by nouns like color, size, weight, length, shape, and other "attribute" nouns into predicates with modifier-like meanings and quasi-adjectival syntactic distribution, thus allowing them to be used not only in post-verbal predicate position as in (29b), (30b), but also in the kind of postnominal position normally reserved for adjectival phrases and prepositional phrases, as in (29c), (30c).

The fact that there are differences in productivity among various meaning-shifting rules, and especially the fact that there are such language-particular rules as the one we have just seen, make it clear that there is a complex relationship (one ripe for study) between the linguistic and the non-linguistic aspects of our propensity to shift words from one meaning to another. What we do in English clearly is not all an inevitable result of more general cognitive principles, but neither does it seem to be wholly language-specific or independent of issues of "cognitive naturalness." As progress is made in the study of systematic shifts in lexical meanings, both synchronic and historical, and in the understanding of semiproductive processes, a rich area of research can be expected to open up with important implications for language acquisition, language change, and the relationship between linguistic and other cognitive processes.

## 4.2. Natural functions.

The examples of meaning-shifting in the previous section may be assumed to represent one way that natural languages deal with the tension between, on the one hand, the advantages of a very systematic correspondence between syntactic categories and semantic types, and on the other hand, the great flexibility and versatility that seem to be essential properties of natural languages.

They are also illustrative of the fact that words and phrases can easily shift their meanings, either temporarily, as in metaphorical or figurative uses of language, or permanently, leading to families of meanings that appear to have families of types. There are analogies in formal languages as well. If one thinks about the operation(s) denoted by "+" on integers and on the rationals, which we normally casually think of as the very same operation, it is clear that given those two different domains, it (they) can't actually be exactly the same operation. But it is common for programming languages to allow the same symbol "+" for both operations (known as "overloading" the symbol), letting properties of the operands determine which interpretation is invoked (a process known as "coercion". The general phenomenon of having single expressions interpreted with multiple semantic types is known as "polymorphism".)

A broader and more speculative goal that we might try to approach from such examples is to take some steps towards sharpening up a notion of "natural" (those are scare quotes!) families of meanings, and more broadly still, a notion of cognitively natural functions. We've clearly seen that multiple meanings of a single word are not always cases of accidental homonymy, and we have seen samples of reasonably "natural" rules for shifting from a given meaning of a word to some new meaning or meanings. Our examples have mainly been drawn from adjective and noun meanings, but this enterprise can be carried out in many sorts of domains. When we concentrate on semantic types and meaning-shifting principles that involve changes in types, this leads to the search for natural families of types and natural type shifting principles, and for notions of natural functions shifting from meanings of one type to meanings of another.

Now, what's "natural"? This is a vague and loaded word, even a dangerous one to use in science, since it is easily subject to abuse and is often used as no more than a biased evaluation of one or another hypothesized principle or property or process. At the same time, if the object of our investigation is some aspect of human or animal mental functioning, then indeed we are trying to find out what sorts of principles are natural to the species, so we shouldn't dismiss the notion, we should just be cautious about rushing to conclusions about it. Is formal elegance relevant? Well, it may be. The scientist's esthetic preference for formal elegance has itself evolved in nature, and our belief in a correlation between the beauty of a hypothesis and its approximation to truth may be a symptom, if we're lucky, of the survival value of a taste for beauty. If we can find analyses with a high degree of both formal elegance and empirical generality, we can suspect that we are on the right track. It will undoubtedly take the cooperation of linguists, mathematicians and cognitive scientists to try to find some notion(s) of "natural functions" in this area that are satisfying in some degree from all those perspectives.

Let's broaden the context with an example from another domain: suppose we were to ask, "What's the most natural function from the real numbers to the integers?" (The reader is invited to think about this before reading ahead.) One logician once volunteered the answer that obviously it's the function that maps every real number onto zero. To a logician, that might be the first thought of the most natural and maybe simplest function from the reals to the integers. And indeed it is undoubtedly a simplest such function, but not for most people a most natural one. Why not? Probably because it "ignores its input", and loses all the information about the argument in passing to the value. Most people seem to agree that the most natural function from the reals to the integers is some version of a rounding off function - a function that maps a given real number to the nearest integer. So let's try to analyze why we all agree that such a function is the most natural one between those two domains. For one thing, it preserves order insofar as you can preserve order. When we're mapping from a larger domain to a smaller one we of course can't preserve everything, but the rounding-off function preserves order insofar as it's preservable. In particular, "less than or equal to" is always preserved, even though "less than" isn't always. That is, if $r_1 \leq r_2$, and $r_1$ rounds off to $n_1$ and $r_2$ to $n_2$, then it is always the case that $n_1 \leq n_2$, even though some instances of "less than" may be mapped into some instances of "equals". The rounding-off function also comes as close as is possible to preserving the various operations like addition and multiplication which are defined on both real numbers and integers. So insofar as those two domains share a certain amount of structure, one can say that the most natural function, the most natural mapping from one to the other, is the one that preserves the most of the relevant structure that they share.

We can even take this example farther and ask what if anything we can say about the choice among the various different versions of rounding off functions, which differ as to how to round off a number that ends with ".5". It is illuminating to see how the choice depends on one's purposes and goals. If you want simplicity and maximal replicability then you choose some very simple rule like "always round down". Or for instance, if you're a merchant and you want to maximize profits then when you're computing prices you always

round up.  But those aren't the most natural rounding-off functions from the point of view of physics or other sciences, since always rounding up or always rounding down is bound to magnify errors when you multiply.  So the one that is commonly taught in elementary physics is, "If it's an even number, round down, and if it's an odd number, round up."  Assuming that even and odd numbers are randomly distributed in the inputs, that will come as close as we can to minimizing the propagation of error and preserving the structure of the argument domain in applying operations that are defined on both domains. But we can pursue this example even farther. The last-mentioned function was still replicable and still a very definite algorithm. If we  really wanted to be physically "natural" and imagine that we're modelling something like balls falling down from a space that we imagine to have real-number distribution and being funnelled into discrete containers, then the rule should be that to model a ball that is falling from a ".5"-type initial location, we should flip a coin to model which way it will go.  That is, we should do something random to model the random part of the physical process; we shouldn't make a deterministic algorithm out of it.  The physics teachers don't suggest coin-flipping probably because they need to be able to check our computations and so they want us to have the same answers; and besides, we have to stretch the notion of function even to be able to call that one a function, though if it's the most natural function for modelling a class of physical phenomena, the physicists should insist on being able to work with a notion of function that allows it.  So even in a simple example like this, the question of what's the "most natural" or the "best" function of a certain general type is interestingly nontrivial. On the one hand we can often give strong arguments that converge on identifying certain properties of such functions (in the example above, that it must be a rounding-off function), but in some respects the choice of a most natural function may well depend on one's purposes and/or on empirical considerations.

The shift in proper noun interpretation discussed earlier can be argued to involve the formally most natural possible function from the domain of entities to the domain of generalized quantifiers; Lewis's "universal grinder" may represent the empirically most natural function from concrete count noun meanings to mass noun meanings. Studies focussing on type-shifting and other structurally-based meaning shifts have also led to a better understanding of the English definite article and the variety of interrelated meanings of different types that can be attributed to it, the basis for the ability of numerals and many other weak determiners to function either as adjectives or as determiners, the analysis of the English copula verb "be", and some explanation for the ease with which languages indicate definiteness and indefiniteness without explicit articles. The exploration of type-shifting and meaning-shifting functions might thus provide an opening wedge into a broader study of cognitively natural functions of various kinds.

One of the many domains in which language offers a "window on the mind" is the domain of metaphor. The investigation of metaphor may be at an opposite extreme in some ways from the investigation of most natural functions from real numbers to integers, but it can be seen as involving a particularly open-ended domain of meaning-shifting principles.   While formal semantics has lagged behind more explicitly cognitively- oriented approaches in contributing to the study of the extremely important and interesting area of metaphor,

there is no obstacle in principle to the integration of the study of metaphor into the model-theoretic framework, and there are potential connections between work on type-shifting and meaning-shifting principles and investigations of what one might call structural metaphors. Lewis's "Universal Grinder" underlies one example of a kind of structural metaphor; others can be found in the research area that Emmon Bach (1986) has dubbed "natural language metaphysics", where specific examples might include formal analogies that have been discovered between the mass-count distinction in the nominal domain and the process-event distinction in the domain of verbal aspect; shifts between viewing one and the same "thing" (such as a war or a thunderstorm or a random act of kindness) as an entity or an event with concomitant shifts in the use of nouns vs. verbs in talking about them; shifts between locative and temporal interpretations of adverbs, prepositions, and measure expressions; and shifts between freqency uses of adverbs like <u>often</u>, <u>seldom</u>, <u>usually</u>, etc., and their use as "unselective quantifiers" over domains with no temporal dimensions but with formal properties that make notions like frequency distributions sensible (much as in the formal study of probability and sampling), as in a famous example of David Lewis's:

    (31)   Quadratic equations usually have two distinct roots.

Further examples of structural metaphors include the imposition of spatiotemporal language and structure onto abstract domains; and likewise the extensions of so-called thematic roles (agent, patient, source, goal, etc.) from the frames of relatively concrete verbs to those of more abstract ones.

Metaphor has to do with the imposing of unfamiliar structures onto familiar domains, or describing one entity or event with language typically used for describing entities or events of a quite different sort. Understanding a metaphor as it was intended requires seeing a relevant pattern of similarity between two different domains, some structure they can be said to have in common in spite of other differences. Since any two domains have infinitely many properties in common and infinitely many differences, a successful metaphor cannot be based on just any similarity; it must be a sufficiently salient one so that the hearer will be able to identify it with some degree of confidence without the speaker having to "explain" it. It is no easy task to develop a theory of the kinds of similarities that are likely to be "salient" to creatures with our particular cognitive and perceptual propensities. The search for a characterization of "natural functions" can be seen as one part of such an effort.

While it is almost certainly impossible to define the notion of "natural function", such a notion may nevertheless be able to play a useful role in bringing formal techniques and cognitive insights closer together, and may provide one of the many potential bridges between semantics and other aspects of cognitive science.

## 5. <u>CONCLUDING REMARKS.</u>

Semantics is a field that goes back a couple of thousand years and is inherently interdisciplinary, with at least as many approaches to it as there are disciplinary entry-points and reasons for being interested in meaning. Formal semantics is a young and dynamic field that is making exciting progress on some of the many questions that can be classified as semantic, and one that is also inherently interdisciplinary, centered in linguistics, philosophy, and logic, with increasing connections to psycholinguistics and to computational linguistics, hence in principle close to the heart of cognitive science. One of the major foundational challenges that faces the fuller integration of formal semantics into cognitive science is the tension that remains between the conceptualist foundations of generative grammar and the antipsychologism inherited by contemporary logic from its Fregean roots and transmitted from there to contemporary formal semantics through the work of Montague and other philosophers and logicians. In semantics more than in other areas of grammar there may well be good reason to distinguish between language and our knowledge of it; there may be an important distinction between "what's in the head" and "what's determined by what's in the head". There is so far, for instance, no good theory of what mental representations of possible worlds might be like, and given that it is easy to argue that there must be non-denumerably many possible worlds and that they are therefore not finitely representable "one by one", this is a serious obstacle to the development of psycholinguistic models of intensionality and semantic processing. Possible worlds (or possible situations, of which there are even more) are a central notion in formal semantics, and possible worlds can't be "in the head" in any straightforward way. But there is plenty of evidence that possible worlds and other notions explored by formal semanticists are at least indirectly cognitively robust, so it is to be hoped and expected that advances on these and the many other open problems in this rapidly developing field will be made by the next generation of researchers, especially those who have benefitted from the interdisciplinary perspectives that are intrinsic to the field of cognitive science. And lest any present student reader suffer from the same fear that this author had as a student, that the interesting problems will all be solved before she has a chance to start working on any of them, let me close with the assurance that every good solution opens up interesting new problems that perhaps couldn't even be posed before, so there is little danger that a genuinely productive line of research will ever simply be "finished".

## SUGGESTIONS FOR FURTHER READING

There are several good introductions to formal semantics. Chierchia and McConnell-Ginet (1990) is a good introductory textbook especially aimed at linguistics students; Bach (1989), based on a series of lectures given in China, is a good non-technical introduction for the linguist or the interested general reader. Gamut (1991), translated from Dutch, is an excellent two-volume work which combines an introduction to logic with a solid introduction to formal semantics; it is the product of an interdisciplinary team of five Dutch co-authors whose fields are logic, philosophy, and linguistics, and the book shows the benefits of the long-standing Dutch tradition of interdisciplinary work among those fields. Dowty, Wall and Peters (1981) is a classic introduction to Montague semantics; Cresswell (1973) is a comprehensible book-length introduction to an approach to semantics quite similar to Montague's, with discussion of both philosophical foundations and many particular English constructions.

Formal semantics traces its recent roots in considerable part to the seminal work of Richard Montague, collected in Montague (1974). Its development in linguistics and philosophy is traced in such works as Lewis (1970), Partee (1973), Partee, ed., (1976), Dowty (1979), Partee (1989a), in articles in the journal Linguistics and Philosophy since its inception in 1977, in volumes of proceedings of the biennial Amsterdam conferences held since the middle 1970's, and in recent years in many other books and journals. Three recent or forthcoming handbooks give good surveys of the current state of the art in formal semantics: Von Stechow and Wunderlich, eds., (1991), Lappin, ed., (to appear), and Van Benthem and ter Meulen, eds., (to appear). Lappin's handbook also includes some other contemporary approaches to semantics.

As noted at the beginning of this chapter, there are many approaches to semantics other than the formal semantic perspective adopted in this and the previous chapter (and the boundaries of formal semantics are themselves somewhat vague.) The interested reader will find a good range of approaches to semantics presented or discussed in such works as the following: Jackendoff (1972, 1983, 1987), McCawley (1973, 1981), Lyons (1977, 1988), J.D.Fodor (1980), May (1985). The earliest attempt to construct a theory of semantics to go with a Chomskyan theory of generative grammar was Katz and J.A.Fodor (1963).

Before Montague's work became known to linguists and the development of formal semantics in linguistics took root, the so-called "linguistic wars" between "generative semantics" and "interpretive semantics" dominated the semantic scene in linguistics. Central figures on the generative semantics side were Lakoff, McCawley, Ross, and Postal; central figures on the interpretive semantics side were Chomsky and Jackendoff. For original sources, see Lakoff (1972), McCawley (1973), Jackendoff (1972), Chomsky (1971); for overviews of the dispute, see J.D.Fodor (1980) and Harris (1993).

There have been representative collections of articles published at various times that give a good snapshot of the state of the art, the interests

of semanticists, and the development of interdisciplinary cooperative efforts among scholars working from different backgrounds. Steinberg and Jakobovits, eds., (1971) is an early interdisciplinary collection with contributions by philosophers, linguists, and psychologists; the articles are in the main written from a single disciplinary perspective, and the volume therefore gives a good perspective on the disciplinary precursors of the current drive towards a more unified cognitive science. A sample of other important interdisciplinary collections would include Davidson and Harman, eds., (1972), Keenan, ed., (1975), Hintikka et al, eds., (1973), Baeuerle et al, eds., (1979), Baeuerle et al, eds., (1983), Groenendijk et al, eds., (1987).

Much of the foundation for work on lexical semantics in the context of Montague Grammar was laid in Dowty (1979), which is still an excellent point of entry for this research area. More about the semantics of adjectives in particular can be found in Kamp (1975), Siegel (1976a,b), Klein (1980), Kamp and Partee (forthcoming).

For more about the fascinating little word the, see Heim (1982) and Neale (1990).

Good starting points for a deeper understanding of the issue of intensionality and the role of possible worlds in its analysis are Carnap (1956), Lewis (1970, 1973, 1986), and Stalnaker (1984). For inquiries that attempt to bridge some of the gap between linguistic and philosophical work in formal semantics on the one hand and concerns for psychological representations on the other, see J.A.Fodor (1975, 1987), Cresswell (1978), Johnson-Laird (1983), Stalnaker (1984).

There is some literature on the very interesting problems of vagueness touched on in section 3.1, some of it rather technical; see Pinkal (1983) for a good model of a formal semantics perspective on vagueness, and see Kamp and Partee (forthcoming) and references cited therein for some discussion linking formal semanticists' and psychologists' approaches to vagueness, particularly addressed to concerns raised about vagueness and compositionality by Osherson and Smith (1981). Problems of vagueness have gained some attention recently under the banner of "fuzzy logic, a term applied to a vaguely delimited family of approaches to the analysis of vagueness starting from the work of Zadeh (1965). A critique of Zadeh's classic version of fuzzy logic is included in Kamp and Partee (forthcoming).

For more on deictics, indexicals, and the interaction of point of view with semantic interpretation, discussed in section 3.2., see Fillmore (1971, 1975) (classic studies in this area), Weissenborn and Klein (1982) (a collection that offers a fascinating look at deictic expressions in a typologically diverse array of languages), and Partee 1989b and references cited therein.

With respect to the interpretation of compounds as discussed in section 3.3., see Gleitman and Gleitman (1971) for a most provocative study of how various groups of people interpret novel 3-word combinations involving mixtures of compounding and modification and mixtures of adjectives and nouns.

For more on type-shifting and meaning-shifting (Section 4) from a variety of perspectives, see Marchand 1960 (a classic philological work with a wealth of data), Dowty 1979 (which laid the foundations for work on lexical meaning in Montague Grammar), Partee 1987 (about type-shifting), Jackendoff 1976 (particularly about basic and metaphorically extended interpretations of thematic roles), and Lakoff and Johnson 1980 (a wide-ranging and illuminating study of metaphors, both lexical and structural, marred in the present author's view by the ill-founded claim that formal semantics is incapable of contributing to the study of metaphor).

There is an increasing amount of interesting work going on at the borderline of philosophy of language, formal semantics, and literary criticism, giving rise to a new field of literary semantics drawing in part on sources in formal semantics and philosophy. See Pavel (1986), which contains a good survey of the field and some of its central issues.

**PROBLEMS**

(1) In Section 2.1. it was noted that "the French teacher" is ambiguous, with possible paraphrases (a) "the teacher who is French" and (b) "the teacher of French". Translate "the French teacher" , in both senses, into two other languages; if possible and perhaps with help from a native speaker, include translation into a non-Indo-European language. Identify nouns, adjectives, and PPs in the translations. Pool results in class, and see if the results might support or disconfirm the hypothesis that the English phrase has two different syntactic analyses, one in which the word "French" is used as an adjective modifying the head and the other a structure (perhaps a compound; see section 3.3) in which "French" is a noun which is interpreted as an argument of the head. (Some languages force you to choose a gender for "teacher"; if so, you may pick a gender at random, or include both.)

(2) (Sec. 2.1.) Although many postnominal of-PPs are arguments of the head, some are modifiers. Examples (a) and (b) below illustrate the one(s) test mentioned in Section 2.1. Assume the correctness of the thesis that one(s) can only be substituted for an N′, and not for an N which isn't also a complete N′. Then tell what conclusions can be drawn about the PPs in (a) and (b) below, and draw correspondingly different trees (selecting between tree types (7c) and (7d)) for the NPs the owners of the shop and the owners of foreign origin.

>    (a) ##The owners of the shop cooperated with the ones of the used car lot. (Anomalous if ones is interpreted as owners)
>    (b) The owners of foreign origin cooperated with the ones of domestic origin.

(3) (Sections 2.2.1., 2.2.2, 2.2.3, and 3.1.)
   (a) Classify the following adjectives as (i) intersective, (ii) non-intersective but subsective, or (iii) nonsubsective; among the nonsubsective, classify further as (iiia) privative, (iiib) plain nonsubsective. There may be unclear or debatable cases; some unclear cases may have different answers for different readings of the adjective in question: if so, suggest readings and corresponding classification.

>    Adjectives: red, strict, new, possible, wealthy, future, audible, poor, miniature, sick, typical, counterfeit.

   (b) Add two more adjectives to each category.
   (c) Write a paragraph discussing one or two unclear cases, either from the list in (a) or from your list in (b).

(4) (Section 2.2.3.) Write a meaning postulate (analogous to (12) and (14) in sections 2.2.2 and 2.2.3) to characterize the privative meaning of an adjective like counterfeit.

(5) (Section 3.1.) Consider the following pair of sentences.

>    (a) Knives are sharp.
>    (b) This is a sharp knife.

If <u>sharp</u> is interpreted with the same positive extension in interpreting each sentence, then if (a) is true, (b) would have to be uninformative. Suggest a diagnosis of how we might adjust the positive extension of <u>sharp</u> differently in the two cases, making use of the Non-Vacuity Principle, the Head Primacy Principle, and whatever other hypotheses you need.

(6) (Section 3.2.) Illustrate by means of diagrams that show position of speaker, position and orientation (if relevant) of mentioned object, and possible intended "target" positions of hearer how "Please stand in front of the car" can be ambiguous, whereas "Please stand in front of the tree" normally is not. Optionally, draw another diagram illustrating a situation in which the second sentence would also be ambiguous. Use descriptions if you find that easier or clearer than drawing diagrams.

(7) (Section 3.2.) (i) Consider example (25) again:

(25) Most foreigners speak a foreign language.

Considering just the possible ways of "anchoring" the words <u>foreigner</u> and <u>foreign</u> to sentence-external or sentence-internal context, how many different readings does this sentence seem to have? (For simplicity, and to limit the number of relevant possibilities, imagine the sentence to be spoken in the U.S. by an English-speaking U.S. resident to another English-speaking U.S. resident, so that there is in effect only one relevant sentence-external context to anchor to, namely "the U.S.".)

(ii) Do the possibilities you found show any differences with respect to the first and second NPs and their possible anchorings? (Question for further thought: make up a hypothesis about syntactic constraints on sentence-internal anchoring of context-dependent words like "foreign", and test it against three or four further examples with different syntactic structures. This is meant as just a sample first step in what could potentially be a much larger investigation.)

(8) (Section 3.3.) Think of three different possible interpretations for the compound "bear towel" and describe scenarios in which each would be readily understood by the hearer as the intended interpretation. Try to make at least one of them relatively "easy" and at least one of them quite improbable, so that a very specific context is required to evoke that interpretation in the hearer. (The relevant context may be either current perceived context or context of shared knowledge, beliefs, past conversations, or experiences, etc.)

(9) (Section 4.2.)
(a) Describe the nature of the productive meaning-shifting pattern by which temporal expressions come to be used in expressing distances, as in "I live just 10 minutes from here." Describe the dependence of this particular meaning-shift on relevant non-linguistic facts about the context.
(b) Give examples of two other kinds of meaning-shifting phenomena that show switches in one direction or the other between locative and temporal interpretations of expressions of some kind, such as adjectival phrases, adverbial phrases, or prepositional phrases, or shifts of phrases that are

originally neither locative nor temporal into locative or temporal uses (e.g. as in "two husbands ago" or "Now we're only about twelve logging trucks from New Aiyansh".)

## QUESTIONS FOR FURTHER THOUGHT

(1) In Section 1.2. we discussed two possible syntactic analyses for the expression <u>almost half full</u>. One possibility is that both analyses are correct and that they correspond to two subtly different semantic interpretations of the expression. Can you think of any arguments for or against such a possibility?

(2) We discussed the ambiguity (in the written language) of the phrase <u>a French teacher</u> in section 2.1. and in exercise (1) above. Consider the fact that the phrase "a French teacher and two German ones" unambiguously selects the modifier reading of <u>German</u> (and by parallel structure effects also selects strongly for the modifier reading of <u>French</u>) and explore whether that fact, together with the proposal about <u>one(s)</u> mentioned in Section 2.1. might be used to argue that in addition to distinguishing <u>French</u> as an adjective and as a noun, there should be two different tree structures for <u>a French teacher</u> analogous to the two different trees (7c) and (7d) for NPs with postnominal PPs. Which tree would go with which interpretation, and why?

(3) (Sections 2.2.3, 3.1, and 4.1.) Give arguments both for and against the classification of <u>fake</u> as a privative modifier. Do the same for <u>stone</u> as in <u>stone lion</u> (discussed in Section 3.1.). Which arguments do you find stronger in each case, and/or can you think of any resolution in either case which could show each side to be somehow correct?

(4) (Section 3.1.) Suggest at least one argument for the need for and usefulness of vagueness in natural language. (Optional suggestion (just a sample): consider the difficulty or perhaps impossibility of expressing generalizations like "Scarce things are usually more expensive than abundant things", if there weren't vague words like <u>scarce</u> and <u>abundant</u>. (Note the apparent involvement of something like a Parallel Structure Effect in this example as well, although this is not a case of conjunction of the predicates.))

(5) (Section 3.2.) Sketch the beginnings of a possible experiment designed to probe the degree to which objects are conceived as having "fronts", and by virtue of what sorts of properties (of the object and/or its location or motion, or of the structure of the setting in which it occurs, or the like), e.g. by looking for relative prominence of different possible interpretations of expressions like "in front of X", "to the right of X", as noted in the text ("in front of the car" vs. "in front of the tree", and the Little League anecdote) and in Problem 6 above.

(6) (Section 3.2.) If you know a language other than English very well or can tap the knowledge of someone who does, see if you can find examples of context-sensitive expressions similar to "in front of", "to the right of", "ahead of", etc., that work a little differently from the corresponding English expressions in how the interpretation is determined on the basis of the context.

(7) (Section 3.3.) The following is a classic brain-teaser kind of question: "Why do mirrors reverse right and left but not up and down?" In the context of semantics and cognitive science, the interesting thing to try to puzzle out is what kind of a question this is: is it a question about the semantics of "right" and "left" vs "up" and "down", or about optics, about conventions of reading and writing, about gravity or some other domain of physics, or about the orientation of our eyes or how our visual system works, or what? If you lie down on your side and try to read your T-shirt then, then what? Are the presuppositions of the question correct? (Does gravity matter?) This question should be just for fun.

(8) (Section 4.1.) Consider various verbs derived from nouns, such as hand, elbow, core, seed, can, dust (the furniture), dust (the crops), and others. Explore the extent to which there are predictable subclasses of noun-to-verb meaning shifts and the extent to which they are predictable from properties of the subclasses of nouns involved.

(9). (Section 4.2.) This problem could be for the whole class together, or could be the subject of an informal survey carried out by students. It is traditionally said that what makes an expression like "keep tabs on" an idiom is the fact that the meaning of the whole is not a function of the meanings of the parts, i.e. that idioms are expressions whose meanings are not compositionally derived, and which therefore must be treated like "phrasal lexical items". Write down your own (relatively unedited) understanding of the what the "tabs" in "keep tabs on" are, and what (possibly metaphorical) action is involved in keeping them on somebody. Then compare responses with classmates, or survey a group of people and compare responses.[10]

## NOTES

1. The spoken versions of this phrase are normally not ambiguous, being distinguished by two different stress and intonation patterns, the "modifier pattern" and the "compound pattern" discussed below in section 3.3. In this section we will continue to speak of (4c) as a single phrase with two interpretations.

2. When the modifier or argument comes before the noun, we have to choose among appropriately reordered versions of trees (7a,c,d); we defer adjective semantics to the next section.

3. To be safe, we need to question whether the semantics of the combination is indeed set intersection 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. Try the snowman test as in (17).

4. The distinction between demonstratives and indexicals is not always sharp; the intended distinction is that demonstratives are typically accompanied by an explicit or implicit pointing or demonstration that fixes their reference ("deixis" is the Greek word for "pointing"), whereas indexicals have their reference fixed directly by the context of the utterance in which they occur: "I" must be the speaker of the utterance, etc.

5. Or more accurately as functions from contexts to intensions, i.e. as functions from contexts to functions from possible situations to referents. But since, like names, demonstratives always have constant functions as their intensions, it is a benign simplification to think of them as functions from contexts to referents.

6. This last use illustrates that context-dependence interacts with quantification and variable-binding, which we aren't going into here, but which provides one of the arguments for having to take account of this kind of context-dependence in an integrated account of the syntax and semantics of "sentence-grammar" rather than trying to leave it to a separate module of pragmatics where the use of language in context is the focus of attention.

7. This (non) generalization holds only for "free compounds", not for "argument compounds" like the earlier example of French teacher (section 2.1), where the first element is interpreted as an argument of a relational head noun.

8. Formally, one might suppose that set union should be just as natural an option as set intersection; in fact, possibly more natural insofar as "addition" is intuitively more natural than "subtraction". One would need to argue from more than purely formal principles that in a context in which one is trying to "add information", intersection, which narrows down the interpretation, is more natural than union, which broadens it.

9. "Productive" is a term used to characterize those patterns which can be freely applied to novel cases; "semiproductive" is a vague term that signals patterns that seem to be somewhat productive but not fully so. Semiproductivity has always been something of a problem for generative grammar and for theoretical accounts of linguistic competence more generally.

10. This suggested project comes from the author's experience with an undergraduate semantics class in which virtually every student had a compositional though metaphorical interpretation for the expression and rejected the idea that "tabs" in that idiom is meaningless, but no two students had the same compositional interpretation. If this phenomenon is general, it provides interesting support for the robustness of the principle of compositionality.

**REFERENCES**

Allén, Sture, ed. (1989) <u>Possible Worlds in Humanities, Arts, and Sciences:</u>
<u>Proceedings of Nobel Symposium 65</u>. Berlin & New York: Walter de Gruyter.

Bach, Emmon (1986a) "Natural language metaphysics", in Barcan Marcus et al
(eds), <u>Logic, Methodology and Philosophy of Science VII</u>, Elsevier Publ.,
573-595.

Bach, Emmon (1989) <u>Informal Lectures on Formal Semantics</u>, State University of
New York Press

Bach, Emmon, & Robin Cooper (1978) "The NP-S Analysis of Relative Clauses and
Compositional Semantics," <u>Linguistics and Philosophy 2.1</u>, 145-150.

Baker, C.L. (1978) <u>Introduction to generative transformational syntax</u>.
Englewood Cliffs, NJ: Prentice-Hall.

Carnap, Rudolf (1952) "Meaning Postulates" in <u>Phil. Studies</u> III, 65-73.

Carnap, Rudolf (1956) <u>Meaning and necessity</u>, 2d. edition with supplements.
Chicago: Chicago University Press.

Chierchia, Gennaro and Sally McConnell-Ginet (1990) <u>Meaning and Grammar: An</u>
<u>Introduction to Semantics</u>, MIT Press 1990.

Chomsky, Noam (1971) "Deep structure, surface structure, and semantic
interpretation", in D. Steinberg and L. Jakobovits, eds., <u>Semantics: An</u>
<u>Interdisciplinary Reader</u>. Cambridge:Cambridge University Press
(183-216).

Cresswell, M.J. (1973) <u>Logics and Languages</u>. London:Methuen.

Cresswell, M.J. (1978) "Semantic competence", in M. Guenthner-Reutter and F.
Guenthner (eds.), <u>Meaning and Translation: Philosophical and Linguistic</u>
<u>Approaches</u>, London:Duckworth.

Cresswell, M.J. (1982) "The autonomy of semantics", in S. Peters & E.
Saarinen (eds.), <u>Processes, Beliefs, and Questions</u>. Dordrecht (Reidel)
69-86.

Davidson, Donald and Gilbert Harman, eds. (1972), <u>Semantics of Natural</u>
<u>Language</u>. Dordrecht: Reidel

Dowty, David (1979) <u>Word Meaning and Montague Grammar</u>, Dordrecht:Reidel.

Dowty, David, R. Wall and S. Peters (1981), <u>Introduction to Montague</u>
<u>Semantics</u>, Dordrecht: Reidel.

Fillmore, Charles (1971) "Types of Lexical Information" in D. Steinberg and L. Jakobovits, eds., Semantics. An Interdisciplinary Reader in Philosophy, Linguistics, and Psychology. Cambridge: Cambridge University Press, 370-392.

Fillmore, Charles (1975), Santa Cruz lectures on deixis. Bloomington:Indiana Univ. Linguistic Club.

Fodor, Janet Dean (1980) Semantics: Theories of Meaning in Generative Grammar. Harvard University Press, Cambridge.

Fodor, Jerry A. (1975) The Language of Thought. New York: Thomas Y. Crowell.

Fodor, Jerry A. (1987). Psychosemantics: The Problem of Meaning in the Philosophy of Mind. Cambridge, Mass.: Bradford Books, MIT Press.

Frege, Gottlob (1892) "Ueber Sinn und Bedeutung", Zeitschrift fuer Philosophie und philosophische Kritik 100, 25-50. Translated as "On sense and reference", in P.T. Geach and M.Black, eds., Translations from the Philosophical Writings of Gottlob Frege. Oxford: Blackwell (1952), 56-78.

Friederici, A. (1985). Levels of processing and vocabulary types: Evidence from on-line comprehension in normals and agrammatics. Cognition 19, 1-34.

Gamut, L.T.F. (1991), Logic, Language, and Meaning. Vol I: Introduction to Logic; Vol II: Intensional Logic and Logical Grammar. University of Chicago Press, Chicago and London.

Gleitman, Lila R. and Henry Gleitman (1971) Phrase and Paraphrase. New York:Norton.

Hamburger, Henry and Stephen Crain (1987) Plans and semantics in human processing of language. Cognitive Science 11: 101-136.

Hamburger, Henry and Stephen Crain (1984) Acquisition of cognitive compiling. Cognition 17: 85-136.

Harris, Randy Allen (1993) The Linguistics Wars, Oxford University Press, New York and Oxford

Heim, Irene (1982) The Semantics of Definite and Indefinite NP's, Ph.D. dissertation, UMass, Amherst: GLSA, UMass.

Hintikka, K.J.J. (1969) Models for Modalities. Dordrecht: Reidel.

Jackendoff, Ray S. (1972) Semantic Interpretation in Generative Grammar, Cambridge, MA: MIT Press.

Jackendoff, Ray S. (1976) "Toward an explanatory semantic representation." Linguistic Inquiry 7, 89-150.

Jackendoff, Ray S. (1977) <u>X-Bar Syntax: A Study of Phrase Structure.</u> Cambridge, MA: MIT Press.

Jackendoff, Ray S. (1983) <u>Semantics and Cognition</u>, MIT Press, Cambridge, MA.

Jackendoff, Ray S. (1987) <u>Consciousness and the Computational Mind</u>, Cambridge, MA: MIT Press.

Janssen, T.M.V. (1983). <u>Foundations and Applications of Montague Grammar</u>, Amsterdam: Mathematisch Centrum.

Johnson-Laird, Philip N. (1983) <u>Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness</u>, Harvard University Press, Cambridge, MA.

Kamp, J.A.W. (1975) 'Two theories about adjectives.' In: E.L.Keenan (ed.), <u>Formal Semantics for Natural Languages</u>. Cambridge University Press, Cambridge, 123--155.

Kamp, Hans and Barbara Partee (forthcoming) "Prototype theory and compositionality", to appear in <u>Cognition</u>.

Katz, Jerrold J., and Fodor, Jerry A. (1963) 'The structure of a semantic theory.' <u>Language</u> 39, 170--210.

Keenan, Edward L. (1974) 'The Functional Principle: Generalizing the Notion of "Subject of".' <u>CLS</u> 10, 298--309.

Keenan, Edward L. (ed.), <u>Formal Semantics for Natural Languages</u>. Cambridge: Cambridge University Press.

Klein, Ewan (1980) 'A Semantics for Positive and Comparative Adjectives.' <u>Linguistics and Philosophy</u> 4, 1--45.

Klein, Ewan and Ivan Sag (1985), "Type-driven translation", <u>Linguistics and Philosophy</u> 8, 163-201.

Kripke, Saul (1963) "Semantical considerations on modal logic", <u>Acta Philosophica Fennica</u> 16, 83--94.

Lakoff, George (1972) "Linguistics and natural logic", in D. Davidson and G. Harman, eds., <u>Semantics of Natural Language</u>. Dordrecht: Reidel, 545-665.

Lakoff, George and Mark Johnson (1980) <u>Metaphors We Live By</u>. Chicago: University of Chicago Press.

Lappin, Shalom, ed. (to appear), <u>Handbook of Contemporary Semantic Theory</u>, Blackwell Handbooks in Linguistics Series, Oxford: Blackwell.

Lewis, David (1970) "General semantics" Synthese 22, 18-67; reprinted in D.Davidson and G.Harman (eds.), <u>Semantics of Natural Language</u>. Dordrecht: Reidel (1972), 169-218; and in B.H.Partee (ed.) (1976), 1-50.

Lewis, David (1973) <u>Counterfactuals</u>, Harvard University Press, Cambridge, MA

Lewis, David (1975) "Adverbs of Quantification," in E. L. Keenan, ed., <u>Formal Semantics of Natural Language</u>, Cambridge University Press.

Lewis, David (1986) <u>On the Plurality of Worlds</u>. Oxford: Basil Blackwell.

Lyons, John (1977) <u>Semantics</u>, Vols. 1 and 2., Cambridge University Press, Cambridge.

Lyons, John (1988) <u>Principles of Linguistic Semantics</u>. Cambridge: Cambridge University Press.

Marchand, Hans (1960) <u>The Categories and Types of Present-Day English Word-Formation. A Synchronic-Diachronic Approach.</u> Wiesbaden: Harrassowitz.

Matthei, Edward H. (1979) <u>The Acquisition of Prenominal Modifier Sequences: Stalking the Second Green Ball.</u> Unpublished Ph.D. dissertation, University of Massachusetts. Distributed by GLSA, Linguistics Dept., UMass Amherst.

May, Robert (1985) Logical Form: Its Structure and Derivation, <u>Linguistic Inquiry Monograph 12</u>, MIT Press, Cambridge, Massachusetts.

McCawley, James D. (1973) <u>Grammar and Meaning: Papers on Syntactic and Semantic Topics</u>. Tokyo: Taishukan. Reprinted (1976): New York: Academic Press.

McCawley, James D. (1981) <u>Everything that Linguists Have Always Wanted to Know About Logic But Were Ashamed to Ask</u>, University of Chicago Press, Chicago; Second edition, 199x.

Montague, Richard (1970) "English as a Formal Language", in B. Visentini et al, eds. <u>Linguaggi nella Societa e nella Tecnica</u>. Milan: Edizioni di Comunita; reprinted in Montague (1974) 188-221.

Montague, Richard (1973) "The Proper Treatment of Quantification in Ordinary English," in K.J.J. Hintikka, J.M.E. Moravcsik, and P. Suppes, eds., <u>Approaches to Natural Language</u>, Reidel, Dordrecht, 221-242, reprinted in Montague (1974) 247-270.

Montague, Richard (1974) <u>Formal Philosophy: Selected Papers of Richard Montague</u>, edited and with an introduction by Richmond Thomason, New Haven: Yale Univ. Press.

Neale, Stephen (1990) <u>Descriptions</u>. Cambridge, MA: Bradford Books, MIT Press

Osherson, Daniel N. and Smith, Edward E. (1981) 'On the adequacy of prototype theory as a theory of concepts.' <u>Cognition</u> 15, 237--262.

Parsons, Terence (1972) "Some problems concerning the logic of grammatical modifiers" in D. Davidson and G. Harman, eds., (1972), 127-141.

Partee, Barbara H. (1973) "Some transformational extensions of Montague grammar", Journal of Philosophical Logic 2, 509-534; reprinted in Partee, B., ed., (1976), 51-76.

Partee, Barbara H., ed. (1976) Montague Grammar. Academic Press, New York.

Partee, Barbara H. (1979), "Semantics - Mathematics or Psychology?" in R. Bauerle, U. Egli, and A. von Stechow (eds.), Semantics from Different Points of View. Berlin:Springer-Verlag, 1-14.

Partee, Barbara 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, GRASS 8, Dordrecht:Foris, 115-143.

Partee, Barbara H. (1989a) "Possible Worlds in Model-Theoretic Semantics: A Linguistic Perspective", in S.Allén (ed.), Possible Worlds in Humanities, Arts, and Sciences: Proceedings of Nobel Symposium 65, Berlin & New York: Walter de Gruyter, 93-123.

Partee, Barbara H. (1989b) "Binding Implicit Variables in Quantified Contexts", in C. Wiltshire, B. Music, & R. Graczyk, eds., Papers from CLS 25, Chicago: Chicago Linguistic Society, 342-365.

Partee, Barbara and Mats Rooth (1983) "Generalized conjunction and type ambiguity", in R. Bauerle, C. Schwarze, and A. von Stechow, eds., Meaning, Use, and Interpretation of Language, Walter de Gruyter, Berlin, 361-383.

Pavel, Thomas G. (1986) Fictional Worlds. Cambridge, MA: Harvard Univ. Press

Pinkal, Manfred (1983) "Towards a semantics of precization", in T. Ballmer and M. Pinkal, eds., Approaching Vagueness. Amsterdam: North-Holland. 13-57.

Rodman, Robert (1976) "Scope phenomena, 'movement transformations', and relative clauses", in B.H.Partee (ed.), Montague Grammar, Academic Press, New York, 165-176.

Siegel, E.A. (1976a) Capturing the Adjective, Ph.D. dissertation, University of Massachusetts, Amherst.

Siegel, Muffy (1976b) 'Capturing the Russian Adjective.' in B.H.Partee (ed.), Montague Grammar, Academic Press, New York, 293-309.

Stalnaker, Robert C. (1984) Inquiry, Bradford Books/MIT Press, Cambridge.

Stechow, Arnim von and Dieter Wunderlich, eds.(1991). Semantik/ Semantics: An International Handbook of Contemporary Research. Berlin: Walter de Gruyter.

Steinberg, Danny and Leon Jakobovits, eds. (1971), <u>Semantics. An Interdisciplinary Reader in Philosophy, Linguistics, and Psychology.</u> Cambridge: Cambridge University Press.

Van Benthem, J.F.A.K. and Alice ter Meulen, eds. (to appear) <u>Handbook of Logic and Language</u>, Amsterdam: Elsevier.

Weissenborn, Juergen and Wolfgang Klein, eds. 1982. <u>Here and There: Cross-Linguistic Studies on Deixis and Demonstration</u>. Amsterdam:John Benjamins.

Zadeh, Lofti (1965) "Fuzzy sets". <u>Information and Control</u> 8, 338-353.
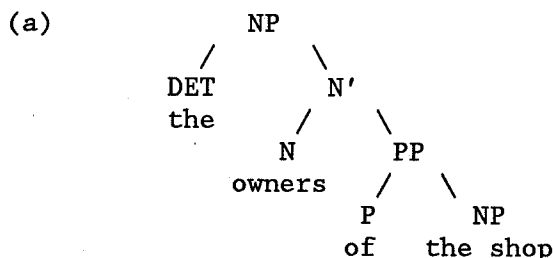
ANSWERS TO PROBLEMS

(1)  For instance:

Spanish: (Note: "teacher" may be masculine or feminine; adjectives agree with nouns.)

    (a) "the French teacher" in the sense of "the teacher who is French":
        (i)    el   maestro              francés
               the teacher (N, masc) French (Adj, masc)
        (ii)   la   maestra              francesa
               the teacher (N, fem) french (Adj, fem)
    (b) "the French teacher" in the sense of "the teacher of French":
        (i)    el   maestro              de francés
               the teacher (N, masc) [of French (N)] (PP)
        (ii)   la   maestra              de francés
               the teacher (N, fem) [of French (N)] (PP)

Czech: (Gender marked on nouns and adjectives, as for Spanish; no definite or indefinite article.)
    (a)    "(the/a) French teacher" in the sense of "teacher who is French":
        (i)    učitel               francouzký
               teacher (N, masc) French (Adj, masc)
        (ii)   učitelka             francouzká
               teacher (N, fem) French (Adj, fem)
    (b)    "(the/a) French teacher" in the sense of "teacher of French":
        (i)    učitel               francouzštiny
               teacher (N, masc) of-French-language (N, fem, genitive)
        (ii)   učitelka             francouzštiny
               teacher (N, fem) of-French-language (N, fem, genitive)

Many languages, like the two illustrated above, have two structurally different translations for the two interpretations of the English phrase. For the (a) reading, one frequently finds an adjective "French" which may show agreement with the gender of the noun "teacher". For the (b) reading, one often finds a noun meaning "French" or "the French language" used either in a PP "of French", as in Spanish above, or put into the genitive case, as in Czech above, both common structures for an "argument" of a head noun. The structural disambiguation found in other languages offers support for the plausibility of a structural ambiguity for the ambiguous English phrase, although such evidence is only suggestive and by itself not at all conclusive.
(2) In "owners of the shop", "of the shop" is an argument. "Owners" by itself is not an N' if it can't be substituted for by "ones". So the structure is like that of (7d):

```
(a)                      NP
                       /    \
                  DET        N'
                  the       /  \
                           N    PP
                        owners  /  \
                               P    NP
                               of  the shop
```

In "owners of foreign origin": if we accept the thesis that "ones" substitution shows that "owners" by itself can be an N', then "of foreign origin" is a modifier. (Semantically it is an intersective modifier: the construction denotes the intersection of the set of owners with the set of individuals of foreign origin, i.e. the set of people who have both properties.) In this case the structure is like that of (7c):

```
(b)                      NP
                       /    \
                  DET        N'
                  the       /  \
                          N'    PP
                          |    /  \
                          N    P    NP
                        owners of foreign origin
```
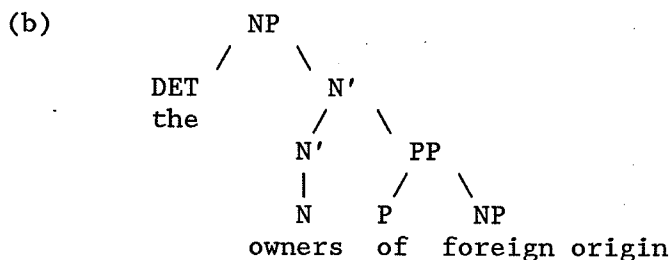
(3) (a) Classifications with some notes about unclear cases that could be used as answers to part (c).

red: (i): intersective. Possible disputes about red wine not being red liquid, but that might better be classified as context-dependent vagueness resolution (Section 3.1.)

strict: (ii): non-intersective but subsective. (example: someone may be a strict teacher but not a strict parent; and that can be true even when strict is used in the same sense and with the same resolution of vagueness.)

new: ambiguous. In one sense, "has not existed for very long", it belongs in class (i): intersective, though it is clearly vague and context-dependent. But it also has a sense in which a "new N" is something that hasn't been an N for very long, as in "new teacher", "new pop star", and in that sense it is in group (ii): non-intersective subsective. (There are interesting ambiguities lurking here, as illustrated in "My new car is an old car", contradictory on one reading and reasonable on another. That's a tricky example, though, since there is an interesting compositionality puzzle in analyzing "my new car", where "new" seems to apply to the possessive relation "car-of".)

possible: (iiib): plain nonsubsective. A possible solution may or may not be a solution; similarly for a possible new virus, a possible forgery.

wealthy: (i): intersective.

future (iiia) nonsubsective, privative. A future king is not a king, a future champion is not a champion, a future outbreak is not (now) an outbreak. This one is subject to debate, perhaps, on the grounds that it may not be strictly entailed that a future N is not now an N; it does not seem contradictory to say "my present and future best friend", and there is a novel by T.H. White with the intriguing title "The Once and Future King" (a novel which plays games with the directionality of time.)

audible (i): intersective.

poor: ambiguous. As the opposite of wealthy, (i): intersective. As an opposite of good, (ii): non-intersective but subsective.

miniature: (iiia): nonsubsective, privative. A miniature city, understood as a model or toy, is not a city. But there may also be a sense in which it could be classified as (i or ii), intersective or non-intersective but at least subsective, if it simply means very small, as in miniature poodle, which is a kind of poodle. Uncertainty about (i) vs. (ii) on this sense relates to whether something could be a miniature St. Bernard but not a miniature dog, and whether the term when applied to dogs simply means "very small", classifiable as intersective but vague and context-dependent, or has a more specific meaning like "a particular small kind of", which could make it non- intersective.

sick: (i): intersective.

typical: (ii): non-intersective but subsective.

counterfeit: (iiia): privative.

(b) Additional examples of each type:

(i):     carnivorous, blond, rectangular, French.
(ii):    recent, good, perfect, legendary.
(iiia):  would-be, past, spurious, imaginary, fictitious, fabricated
         (in one sense), mythical (maybe debatable); there are prefixes
         that have this property too, like ex-, pseudo-, and of course
         non-.
(iiib)   potential, alleged, arguable, likely, predicted, putative,
         questionable, disputed.

(4) Meaning postulate for privative adjectives:

$$||\text{counterfeit N}|| \cap ||N|| = \phi$$

(5) Suppose the normal meaning of sharp is such that (a) "Knives are sharp" is true. Then using that interpretation, (b) would be uninformatively redundant, because every knife would be a sharp knife. The Head Primacy Principle says to keep the interpretation of knife fixed and "recalibrate" sharp. The Non-Vacuity Principle says that we recalibrate sharp in the context of the

denotation of the head noun <u>knife</u> in such a way that some knives are sharp and some are not, so we interpret <u>sharp</u> more narrowly or strictly, raising the threshhold for what counts as sharp.

(6) Discussion can be found in Fillmore (1971). Basically, "in front of the car" can be interpreted either (i) with respect to a line of orientation involving speaker, hearer, and car: the hearer is in front of the car if the hearer is between the speaker and the car; or (ii) with respect to a line of orientation involving the front-back axis of the car itself: the hearer is in front of the car if the hearer is "in front of the front of the car", regardless of where the speaker is.

"In front of the tree" normally has a reading of type (i) only, since trees don't usually have intrinsic "fronts"; however, a tree could have a front if it were shaped or pruned in a particular way, and a tree can indirectly come to be seen as having a front and a back if, for instance, it is adjacent to the front of a house.

(7) (i) Sentence (25) has two readings, roughly paraphrasable as follows:

    (a) Most non-U.S. residents speak a language other than English.

    (b) Most non-U.S. residents speak a language other than the language of their country.

(I would argue that the sentence actually says nothing at all about whether they speak a foreign language rather than or in addition to the non-foreign language in question; others may perceive that as an additional source of ambiguity.)

(ii) Yes, there seems to be an asymmetry between the first and second occurrences of "foreign." You can anchor the subject to the "external context" and then anchor the object either to the external context or to the context associated with the interpretation of the subject. But you can't anchor the object to the external context and then anchor the subject to the context associated with the object. What the missing interpretation would be like can perhaps be seen by thinking about the possible interpretations of "some foreign language is understood by most foreigners".

Sometimes we do get interesting "symmetrical" anchorings, as in "Friends don't let friends drive drunk", where each anchors to the other (analogous perhaps to "Bach-Peters" pronouns.)

The constraints on anchoring are not fully understood (see Partee (1989)), but seem approximately parallel to the sorts of constraints that govern the interpretation of pronouns: roughly, the "antecedent" (or "anchor") must be in a position that is "higher than" and/or "to the left of" the pronoun or context-dependent expression. Non-linguistic antecedents can (almost) always count as "higher".

(8) Some among the indefinitely many possible interpretations for the compound <u>bear towel</u>.

"Easy":
- (a)　towel for drying or rubbing down bears.
- (b)　towel with pictures of bears on it
- (c)　towel made out of a bearskin

"Improbable", needing more support from particulars of shared context:
- (d)　(the) towel that has a hole in it from the time the bear got into our tent
- (e)　towel to be worn around the waist to keep bears away
- (f)　towel for the baby to take to bed with him as a substitute for his lost teddy bear
- (g)　towel to carry bears in

Possible point for debate: Is there a possibility of an interpretation in which "bear" is a verb, analogous to "pickpocket", "killjoy", "carryall"? It might then have a meaning like "towel bearer". It is debatable whether that kind of compound formation is really productive in English. (In Spanish, on the other hand, that sort of Verb + Noun compound formation is very productive, and Noun + Noun compounds, so common in English, are absent.)

(9)　(a) If the primary meaning of an expression $\alpha$ is a measure of temporal duration, such as "10 minutes", then the expression $\alpha$ can have as a derivative meaning a measure of distance, namely the distance typically or conventionally traveled in that amount of time. The precondition for such a shift is the existence of a conventionally understood means of travel in the given situation at a more-or-less predictable speed.

(b)　(i) "from place A to place B" can be used as a measure of time if there is travel going on along a path; on a train trip, one can say "We played cards from Pittsburgh to Cleveland". This shift is similar to that in (i); both trade on the fact that if there is motion along a fixed path at a known rate, then distance can be computed from elapsed time and vice versa.

(ii) Similarly with frequency expressions: "We stopped every 50 miles" uses a distance expression rather than a temporal expression to measure frequency. (This can lead to interesting discussion, since in a sense the whole frequency notion may be shifted into the spatial dimension. If one says "There were signs every fifty miles", that is similarly a kind of "frequency" expression, but purely spatial, with no "translation" back into the temporal domain. The line between meaning-shifting and structural metaphor is a fuzzy one.)

(iii) The example "Now we're only about twelve logging trucks from New Aiyansh" comes from an experience of regularly travelling a hundred-mile road between two towns in Northern British Columbia on which you meet logging trucks coming the other way at regular intervals; thus counting logging trucks, though they are moving, can become a basis for measuring both time and distance from the destination as long as both you and the logging trucks travel at more or less known rates.

**KEY-WORDS FOR THE INDEX**
adjectives, intersective
   , subsective
   , nonintersective
   , privative
arguments (of a head; of a function)
compositionality
compounds, non-compositionality of
context-dependence
demonstratives
DET (determiner)
extension
Frege's principle
functions [as in function-argument application]
    , adjectives as
idioms
indexicals
intension
Lewis's Advice
meaning
meaning-shifting
meaning postulates
modifier, non-restrictive
     , restrictive
Montague
noun, common
noun, proper
NP (Noun Phrase)
point of view
possible worlds
properties (vs. sets)
recursive
semantic features
sets
semantics
    , formal
truth conditions
type-shifting
vagueness