Acquisition of Adjective Meanings: A Lexical Semantic Approach

by

Sandeep Prasada

B.Sc. (Honors), McGill University

Submitted to the Department of Brain and Cognitive Sciences in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

Massachusetts Institute of Technology

September, 1992

Signature of Author _____
Department of Brain and Cognitive Sciences
September 9, 1992

Certified by _____
Susan Carey
Professor, Brain and Cognitive Sciences
Thesis Supervisor

Accepted by _____
Emilio Bizzi
Chair, Department of Brain and Cognitive Sciences

Acquisition of Adjective Meanings: A Lexical Semantic Approach

Sandeep Prasada

Submitted to the Department of Brain and Cognitive Sciences in September, 1992 in partial fulfillment of the requirements for the Degree of Doctor of Philosophy in Cognitive Science.

Abstract

This thesis explored the possibility of deriving constraints on the acquisition of adjective meanings from knowledge of the semantics of adjectives. Chapter 1 briefly reviews the problem of learning the meanings of words and possible classes of solutions. It is proposed that an important source of constraint on the meanings of novel words can be derived from studying adult's knowledge of language and how it is related to nonlinguistic cognition.

In Chapter 2, I propose a semantic analysis of adjectives and how they combine with their arguments in different structural relations. It is proposed that there are two kinds of adjectives. The first kind of adjectives name properties of things. These adjectives name values along dimensions. Adjectives that name values on dimensions are proposed to combine with nouns in two ways. In predication, these adjectives name the value that the entity named by the subject has on the relevant dimension. In modification, the value named by the adjective is used to pick out the objects of the kind named by its argument. The second type of adjectives are those that do not name values on dimensions but combine with nouns to pick out kinds of things. The semantic analysis presented here is based on structural notions and is not meant to be in conflict with syntactic theories though the semantics are used to explain certain phenomena which may also have syntactic explanations. In the rest of this thesis I show how this knowledge of the semantics of adjectives can explain a number of linguistic and psycholinguistic facts as well as motivate heuristics for the acquisition of adjective meanings.

The third chapter investigated the constraints on structures and the interpretation of structures that follow from the semantics of modification. It was found that the semantic analysis of adjectives can be used to explain why adjectives cannot modify NPs (**big the dog*), pronouns (**big he*), proper names (**big John*) and quantified expressions (**the big one dog*). Furthermore, the semantic analysis of adjectives provided an explanation for when prenominal adjectives can be interpreted as being modificational and when they have to be interpreted appositively.

The fourth chapter shows that the semantics of adjectives can be used to explain when multiple adjectives have to be ordered and how they have to be ordered. In addition to explaining why there are preferred adjective orders (the big red plastic ball), the analysis provides an explanation for the conditions under which these orders can be violated. The explanation is supported by psycholinguistic data from previous research.

Chapter five investigates the restrictions on which expression can occur in the position that prenominal adjectives appear in (e.g. **the proud of his son man, * the afraid man*). Once again, the semantics of adjectives provide an explanation for the restriction that applies to expressions in prenominal position.

Having provided psycholinguistic and linguistic evidence that the semantic analysis of adjectives, chapter six investigates whether the semantic analysis can be used to derive constraints that can help the learner infer the meaning of a novel word. After reviewing the literature on children's acquisition of adjectives, three heuristics are derived from the semantics of adjectives and tested in experiments with two and three year old children and adults. The results of the experiments and previous research suggest that the semantic knowledge of adjectives that children have is the same as that of adults and that children can use this knowledge to help them learn the meanings of novel words.

Thesis Supervisor: Dr. Susan Carey
Title: Professor of Brain and Cognitive Sciences

Acknowledgements

I have had the good fortune of having the opportunity to interact with and learn from a number of extremely intelligent and interesting people. I am especially grateful to have had Susan Carey as an advisor. In addition to serving as a role model as a researcher, Susan was an ideal advisor. She listened to my ideas patiently and carefully and always provided thoughtful criticism and encouragement. Her emphasis on getting the conceptual issues clear has had a very important influence on my work and thinking. As I have often told my friends, "I wanna be Susan when I grow up!".

Conducting research with Steve Pinker has been a great experience. Steve's emphasis on the logical problems in acquisition and his careful analysis of empirical results have influenced me greatly. I am also grateful to Ken Wexler and David Pesetsky for being members of my thesis committee. In addition to reading and commenting on my thesis, they have led me to carefully think about how my work relates to other work in the field.

While I have learned a great deal from my professors, I may have learned just as much from my many interactions with the graduate students in the department. As the devil's advocate, I was willing to take a contrary position in most any discussion. My friends and fellow students patiently indulged me and in the course taught me a great deal about topics ranging from language to reasoning, vision and music. While no one who entered the graduate student lounge was safe, John Kim, Josh Stern, Daphne Bavelier, Robbie Jacobs, Paul Bloom and William Snyder were especially unsafe. I am very grateful for the opportunity to discuss and argue issues with them in a friendly and helpful manner. John Kim deserves a special award for trying desperately to educate me about linguistic theory and for being a great friend. I'd also like to thank Gary Marcus for providing me with detailed comments on an earlier draft of the thesis.

Discussions with Nancy Soja, Sandra Waxman, Geoff Hall and John Macnamara over the course of the years have proven to be very helpful. I am especially grateful to John Macnamara, who along with Laura Pettito, got me excited about the study of language and greatly influenced my approach to the study of language.

Dana Sussman, Michael Schlaefer, Jenny Welch, and Laura Klatt helped me a great deal by helping find and run subjects. I am thankful to the teachers and children at the Bright Horizon's Children Center, Lotus Children Center, Beacon Hill Nursery School, Children's Meeting Place, Cambridgeport Children's Center for participating in the experiments.

Doing research was enjoyable because of the many friends with which I could loiter, eat, drink beer, watch movies, play tennis and throw darts. Among others, Daphne Bavelier, Martyn Bracewell, Eliot Charles, Robbie Jacobs, John Kim, Laura Klatt, Corrie Lathan, Anita Moryadas and Josh Stern all helped me enjoy my time at MIT. I am especially lucky to have Laura Klatt as a friend. In addition to helping with the designing of stimuli, giving me comments on my written work and helping me find subjects, Laura has been a very special friend whose encouragement, support and company has been a great source of joy.

Table of Contents

Chapter 1

The Problem of Word Learning

1.1 Introduction.

Learning the meaning of a word requires that we learn to use the word to talk about the same range of entities, situations or relations that our speech community uses the word for. In other words, learning the meaning of a word requires the learner to guess what the speaker had in mind when the speaker used the word.

It has long been noted that in a situation of ostensive definition, there are an infinite number of hypotheses about the meaning of the novel word that are consistent with the evidence that the learner has available. Put another way, word learning presents a particular instance of the induction problem. A famous example from Quine (1960) illustrates this problem. If a speaker utters *gavagai!* in the presence of a rabbit, the evidence the learner has concerning the meaning of *gavagai* is consistent with an infinite set of meanings including rabbit, white, furry, smaller than an elephant, undetachable part of rabbit, rabbit in front of me right now, and rabbit or telephone. Furthermore, there is no way to rule out some of these hypotheses as being logically more primitive than others (Goodman, 1958). This suggests that the learner has to be constrained in the hypotheses that he or she considers. A number of researchers have suggested constraints that may be used in the process of word learning (Bloom, 1990; Brown, 1957; Clark,

1987; Landau, Smith, & Jones, 1988; Markman & Hutchinson, 1984; Markman, 1989; Soja, Carey, & Spelke, 1991). Before discussing specific constraints that have been proposed to play a role in word learning I will explore the problem of learning the meaning of a word and possible classes of solutions to this problem.

As mentioned above, the problem the learner has to solve is to guess what the speaker had in mind when the speaker used the novel word. The solution to this problem is obvious -- the learner needs to have a way to read the speaker's mind. The better the mind reader the learner is, the more successful the learner will be in guessing the state of affairs the speaker has in mind when the speaker uses the word. This ability to read the speaker's mind is presumably due to the brains of the speaker and the listener being extremely similar and processing similar inputs in similar ways. The project for researchers interested in word learning is to discover the ways in which the learner tries to guess what the speaker may be talking about. The task is an extremely difficult one -- we can, and do, speak about things and situations which are immediately observable by the listener but we also speak about events in the past, the future, mental states and possible states of affairs. The solution to the mind reading problem is likely to require many different types of knowledge.

1.2 Nonlinguistic constraints.

One kind of knowledge that will be useful in trying to infer what the speaker is talking about is that of the context in which the novel word is used. For example, if a speaker has made it known to the listener that he/she is extremely hungry then the listener may infer that a novel word in a request

such as *can you please pass me the blicket?* is likely to be the name for a kind of food.

Another way in which the learner can use contextual information to infer the meaning of a novel word is to use the meanings of the other words in an utterance. For example if a learner hears a sentence with the words *John, delicious, eat* and *blicket* in it, and if the learner knows the meaning of *John, delicious* and *eat,* then it is a good guess that *blicket* names something to eat. Note that this is not using knowledge of language structure to infer the meaning of the novel word. Rather it is using knowledge of the context which has been conveyed linguistically (Pinker, 1992).

Though it is clear that adults and possibly children can and do use knowledge of the world and context to try and infer the speaker's intention when using a particular word, there are not any well worked out theories of how we make inferences given our world knowledge or how we infer the mental states of others. This makes it difficult to propose constraints on word meaning that are derived from our world knowledge and knowledge of others' intentional states and thus constraints that are proposed on the basis of these types of knowledge are likely to be *ad-hoc* and situation specific.

1.3 Constraints from knowledge of language use.

The learner can also use knowledge of language use to try and guess what the speaker is talking about. Grice (1975) has proposed that participants in a conversation usually behave in accordance with a principle of co-operation. He argues,

> Our talk exchanges ... are characteristically, to some degree at least, cooperative efforts; and each participant recognizes in them, to some extent, a common purpose or set of purposes, or at least mutually accepted direction... at each stage, *some*

> possible conversational moves would be excluded as conversationally unsuitable. We might then formulate a rough general principle which participants participants will be expected (*ceteris paribus*) to observe, namely: Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. (Grice, 1975: pg 45)

Grice went on to specify a number of maxims that are part of the cooperative principle. The maxims included the maxim of quality and the maxim of relation. The maxim of quality states that the speaker should not say what he/she believes to be false or things which the speaker lacks adequate evidence. The maxim of relation states that the speaker should be relevant. If the learner knows these maxims and assumes that the speaker is being cooperative, then the learner can assume that the information that the speaker is providing in an utterance is relevant and believed by the speaker to be true. For example, if the speaker points to an object and says, *that is a table,* then the maxim of quality predicts that the speaker's name for the object is *table* rather than something else like *ball.* Similarly, if the learner asks the speaker what an object is, the maxim of relation predicts that the speaker should respond with a name for the object rather than some unrelated truth such as the speaker's favorite color.

Another example of a constraint on word meaning that may be a pragmatic constraint is the principle of contrast (Clark, 1987; Gathercole, 1987). The principle of contrast states that no two words can have the same meaning. It is possible that this constraint is due to the structure of the lexicon, but it may also be true that this constraint is due to the listener assuming that if the speaker wanted to convey a meaning for which a word is already known, then the speaker would have used the familiar/conventional word.

The examples given above show how a learner can use knowledge of language use to try and infer what the speaker has in mind when the speaker uses a word or utterance. Any complete theory of word learning will include constraints that are based on a learner's knowledge of language use. It is an empirical question, however, how much pragmatic knowledge the young child has. There is some evidence that young children behave in accordance to the principle of contrast (Clark, 1987; Au & Markman, 1987) but there is little work on how children may use other kinds of pragmatic knowledge in word learning.

1.4 Constraints from links between language and nonlinguistic cognition.

A number of researchers have proposed that children can make use of links between language and nonlinguistic cognition to learn the meanings of words (Bloom, 1990b,1991; Brown, 1957; Clark, 1987; Gelman & Markman, 1985; Gleitman, 1990; Hall, 1991; Hall & Waxman, in press; Katz, Baker, & Macnamara, 1974; Landau, Smith, & Jones, 1988; Markman & Hutchinson, 1984; Markman, 1989; Soja, Carey, & Spelke, 1991; Soja, 1992; Taylor & Gelman, 1988; Waxman & Gelman, 1986). While these researchers agree that children may use links between language and cognition to learn the meanings of words, the exact links that are proposed vary.

For a learner to use knowledge of language structure to infer the meanings of words it has to be the case that there are bi-directional links between grammatical knowledge and nonlinguistic cognition. This condition differs from the uni-directional connection between nonlinguistic cognition and grammar that theories of semantic bootstrapping assume (Grimshaw,

1981; Pinker, 1984). Semantic bootstrapping theories propose that part of our knowledge of language includes correlations between form and meaning; for example, objects are named by nouns and actions are named by verbs. It is proposed that we exploit this knowledge of language to identify instances of grammatical categories such as nouns and verbs. Because the correlations between form and meaning are proposed to be uni-directional it is not possible to use knowledge of the grammatical category of a novel word to learn its meaning. The motivation for proposing that the connections between grammar and cognition are uni-directional is that the links don't seem to be one-to-one. For example, though names for objects are always nouns and names for actions are always verbs, nouns don't always name objects (e.g. *idea*); similarly verbs don't always name actions (e.g. *seem*).

Bloom (1990b) argues that grammatical categories have abstract semantic properties that make it possible to have bi-directional connections between nonlinguistic cognition and grammatical categories. For example, Bloom (1990b; 1992) proposes that count nouns name kinds of entities that are construed as individuals. He proposes that this is the semantic knowledge that underlies the adult's representation of count nouns. An advantage of this account is that it helps to explain why abstract words like *day* and *inch* are lexicalized as count nouns; these words name concepts that are necessarily quantified in terms of individuals. Furthermore, Bloom proposes that children have knowledge of grammar-cognition links innately and can exploit them to learn the meanings of novel words. The constraint is supported by word learning experiments in which children take words used to talk about objects as naming the kind of object (Dickinson, 1988; Markman & Wachtel, 1988; Soja, 1987; Soja, Carey, & Spelke, 1991; Soja, 1992). Objects are usually conceptualized as individuals and thus are easily taken to be

named by count nouns. Bloom (1990) has shown that three year olds can also name individuals that are not objects when a count noun is used. He showed that children are more likely to think that a novel word names a discrete punctate sound than a continuous sound when the novel word is a count noun than when it is a mass noun. This result shows that children expect count nouns to name individuals and that the individuals do not have to be objects but can be temporally individuated sounds. This pattern of results is predicted if children, like adults, have grammar-cognition links which lead them to expect count nouns to name individuals.

Another proposal that exploits the correspondences between linguistic structure and meaning that exist in adult language has been made by Gleitman (1990). She proposed that children may be able to learn the meanings of verbs by keeping track of the syntactic contexts in which a verb can occur. Gleitman and her colleagues (Fisher, Gleitman, & Gleitman, 1991; Fisher, Hall, Rakowitz, & Gleitman, 1991; Landau & Gleitman, 1985; Naigles, 1990) have presented evidence which suggests that children and adults may be able to use the syntactic contexts in which a verb appears to learn aspects of the verb's meaning.

While one may be able to learn certain aspects of a verb's or a noun's meaning from it's syntactic contexts, Pinker (1992) argues that the amount of meaning that one can learn from purely structural information is minimal. That this is the case is evident if one looks at the proposal Bloom (1990b) makes for the semantic knowledge underlying count nouns. The only thing that a learner can infer from a speaker's use of a count noun is that the speaker is talking about something which is construed as an individual. On it's own, this information tells the learner only that the thing being spoken about is the kind of thing that can be counted and cannot be arbitrarily

divided to yield other entities named by the same word. When combined with a nonlinguistic representation of a context, the grammar-cognition links can play an important role in determining what is being spoken about. For example, if a speaker points to an object and names it with a count noun then if the learner knows that count nouns name individuals, the learner will be able to rule out those construals of the entity that do not correspond to individuals such as the substance of which the object is made or its texture or color. This is an example of the "zoom lens" effect (Fisher, Hall, Rakowitz, & Gleitman, 1991) in that it focuses the listener on the structurally encoded perspective that the speaker wants the listener to take. Furthermore, if one has knowledge of the meanings of other words in the sentence and how they are structurally related to the novel word, one can further constrain one's hypotheses about the meaning of the novel word (Pinker, 1992). For example, if one knows that *blicket* names animate entities then the hypotheses regarding the meaning of a novel verb *glip* heard in a context such as *the blicket glipped the fep* are greatly constrained. The hypotheses are further constrained if one knows the meaning of *fep*.

The current thesis investigates how one can use knowledge of the semantics of adjectives and how they combine with words in structural relations to "zoom in" on the type of meaning. Furthermore, it is proposed how this knowledge can be combined with nonlinguistic knowledge of the entities in question to help infer the meaning of novel words. The heuristics for how we may learn the meanings of novel adjectives are derived from an analysis of how adjectives are represented by adults.

The constraints that are proposed in this thesis are derived from the links that exist between language and cognition in adults because knowledge of what is acquired provides an important source of constraint on acquisition

mechanisms. In order to have a theory of how something is acquired one needs to know what is acquired. This approach to acquisition, often called the learnability approach, has been standard in research on the acquisition of syntax (Wexler & Culicover, 1980; Pinker, 1984; Hyams, 1986). Until recently relatively little research (Bloom, 1990; Clark, 1987; Gelman & Markman, 1985; Hall, 1991; Hall & Waxman, in press; Katz, Baker & Macnamara, 1974; Soja, 1992) has used knowledge of what is acquired as a source of constraint on the meanings of novel words. One reason why this may be the case is that it is not clear exactly what a theory of meaning should account for. While this is true, one cannot state *a priori* what the domain of *any* theory will turn out to be. An approach that has been successful in the study of syntax and lexical semantics is to try to discover the components of meaning that are relevant to linguistic structure (Levin & Pinker, 1992). While it is clear that there is more to the meaning of a word than those aspects of meaning that are relevant to structure, the study of these components of meaning allows one to propose acquisition mechanisms in a motivated way. The aspects of meaning that are relevant to the structure of sentences have to be acquired by the learner and thus any theory of meaning has to minimally account for how these aspects of meaning are acquired.

By taking an approach in which the constraints on word meaning are derived from the adult's knowledge of language, one avoids a number of problems associated with postulating that the child's knowledge of language is different from that of the adults. These problems include: (i) How does the child come to have the adult knowledge?, (ii) Are the differences observed between children and adults due to a difference in their knowledge about how language encodes meanings or are the differences due to other factors?, and (iii) How does the proposed difference in knowledge fit with the large

body of data suggesting that children's knowledge is the same as that of adults? (Bloom, 1990a; 1990b; Poeppel & Wexler, 1992; Naigles, 1990; see Crain, 1991 for a review). This question is important because grammatical categories and principles are often interdefined and interdependent and thus denying the child some aspect of adult knowledge may imply denying them other knowledge too.

1.6 Overview of thesis.

While there has been a great deal of research done on how children learn nouns, there has been relatively little work done on how children learn the meaning of adjectives (Carey, 1978; Gelman & Taylor 1988; Gelman & Ebeling, 1986; Golinkoff, Mennuti, Lengle, & Hermon, 1992; Smith, Jones, & Landau, 1992). In this thesis, I propose a few constraints that may be used in learning the meanings of adjectives that are motivated by a semantic analysis of adjectives. In Chapter 2, I present a semantic analysis of adjectives that is proposed to underly an adult's knowledge of adjectives. In the next three chapters I examine a number of phenomenon that are explained by the proposed analysis and thus provide evidence for the analysis. In Chapter 6, I review the research on the acquisition of adjectives and how it relates to the proposal of the semantics of adjectives developed in this thesis. I also propose three constraints on the acquisition of adjective meanings that are motivated by the semantic analysis of adjectives. The heuristics are tested in experiments with adults and children. The chapter also summarizes the results and identifies questions that the current research raises.

Chapter 2

Semantic Knowledge of Adjectives

2.1 Semantic relations and structure.

Language allows speakers to express an inummerable number of thoughts and comprehend sentences that they have never heard before. This ability can only be accounted for if humans have a generative knowledge system. This knowledge system specifies the manner in which words with specific meanings can be combined to produce a sentence with a given meaning. For example, if one wants to express the thought that JOHN LIKES MARY then there is a specific way in which one combines the words *John*, *likes*, and *Mary* to express this thought. In English, one would express this thought in a sentence like (1). The combinations in (2)-(6) cannot be used to express the thought JOHN LIKES MARY, though (5) can express a thought in English. Note that the thought JOHN LIKES MARY can also be expressed in sentence (7). The relation that holds between *like*, *John* and *Mary* in (1) is the same as the relation that holds between these words in (7): *Mary* is the logical object of *like* and *John* is the logical subject of *like* in both (1) and (7). This insight motivated a transformational theory of grammar.

(1) John likes Mary.
(2) * John Mary likes.
(3) * likes John Mary.
(4) * likes Mary John.
(5) * Mary likes John.
(6) * Mary John likes.
(7) Mary is liked by John.

The point that is relevant to the current thesis is that certain semantic relations are structurally encoded in a sentence. The semantic relationships that are structurally encoded in a sentence are called *grammatical functions* (Chomsky, 1965).[1] Examples of grammatical functions include:

(8) (i) Subject-of: [NP, S]
(ii) Predicate-of: [VP, S]
(iii) Direct-Object-of: [NP, VP]
(iv) Main-Verb-of: [V, VP]

Thus, Subject-of encodes a relation that holds between an NP and an S, Predicate-of encodes a relation between a VP and an S and so on... The examples of grammatical functions given above are structurally defined semantic relations that have specific names but one can talk about many other structurally defined relations. For example, the relation between Subject and Verb can be defined as the relation between the Subject-of a Sentence and the Main-Verb-of the Predicate-of the Sentence (Chomsky, 1965).

In section 2.2, I will present an analysis of the basic ways in which adjectives can semantically combine with nouns and how this places constraints on what structures are allowed and how they are interpreted. The question of interest is how words with certain types of meaning (adjectives in the present case) semantically combine with arguments in different structural positions. Sentence (9) illustrates the interaction of word meaning and structural relations. The sentence is syntactically well formed but is

[1] The examples of how certain semantic relations are structurally represented given here are from a theory of the type discussed in Chomsky (1965). The particular semantic relations, how they are encoded in structures and the relationship between different levels of representations are different in current theories of grammar but the idea that certain semantic relations are structurally encoded remains (Chomsky, 1986).

semantically anomalous because a table is not the kind of thing that can kick. The structure indicates that *table* is the subject of the verb *kick*, and *boy* is the object of the verb. It is because the relations between *boy*, *table* and *kick* are structurally defined in the way that they are that the sentence is anomalous. If *boy* and *table* exchange positions in the structure, then the sentence is no longer anomalous. The anomaly in (9) is not due to a restriction on the type of word that can occur in subject position (e.g. must be animate). The anomaly is due to the selectional restrictions placed on the subject by the verb *kick*. It is perfectly fine to have an inanimate subject with a verb like *frighten* (10).

(9) The table kicked the boy.
(10) The table frightened the boy.

In the next few chapters, I will be investigating how certain structurally represented semantic relations between adjectives and various types of phrases are interpreted. The semantic analysis of how adjectives combine semantically with nouns will then be used to motivate some constraints on mechanisms for the acquisition of adjective meanings.

2.2 A semantic analysis of adjective.

The semantic analysis of the types of adjectives presented here is similar to many previous analyses of adjectives (Bolinger, 1967; Gross, Fischer, & Miller, 1989; Levi, 1978; Miller & Johnson-Laird, 1976). Semantically, there seem to be two kinds of adjectives. Most adjectives name properties of things. Examples of adjectives of this kind include adjectives

such as *big, red,* and *clean.* The adjectives in (11a,b) name a property of my room. The other kind of adjectives do not name a property of things but help identify certain kinds of things. Examples of this kind of adjective include *former, alleged* and *corporate*(12a,b) These adjectives do not name a property but help identify entities that are former lawyers, alleged lawyers and corporate. The entities that are named by *former lawyer* and *alleged lawyer* either are not lawyers or need not be lawyers and thus *former* and *alleged* could not name a property of lawyers. Unlike the phrases *alleged lawyer* and *former lawyer, corporate lawyer* necessarily names entities that are lawyers. This leaves open the possibility that corporate names a property of lawyers. I will argue in Section 2.3.2 that corporate does not name a property of things in the way that adjectives like *big, red,* and *clean* do.

(11a) The big/red/clean room is mine.
(11b) My room is big/red/clean.

(12a) The former/alleged/corporate lawyer is ...
(12b) * The lawyer is former/alleged/corporate.

2.2.1 Adjectives that name properties of things.

Adjectives that name properties of things do so in a particular way; they name the value that the entity named by the argument has on a given dimension (Gross, Fischer, & Miller, 1989). This aspect of the semantics of adjectives that name properties of things is reflected in the fact that most of these adjectives can have opposites and can appear in comparative constructions. In English most adjectives form their opposites through affixation (e.g. *happy/unhappy, legal/illegal, moral/immoral, desirable/undesirable*) though some adjectives like *big/small* have non-

morphologically related opposites. Another property of adjectives that name values on dimensions is that values on a dimension are mutually exclusive and a given object can only have a single value on a dimension. Thus, a given entity cannot be both big and small, wet and dry or male and female, etc... This account of adjectives that name values on dimensions seems to be undermined by the acceptability of sentences such as (13) and (14). In (13), two values on the dimension are being predicated of the subject NP. In (14), two values on the dimension COLOR are being used to modify the entities named by *house*. If it is true that an entity cannot have more than a single value on a dimension, then these sentences should be unacceptable. The sentences actually are unacceptable unless one analyzes the words *the house* or *house* as naming parts of a house. Thus, (13) and (14) are acceptable only if they are interpreted in the same way as (13a) and (14a).

(13) The house is red and green.
(14) The red and green house ...

(13a) Parts of the house are red and parts of the house are green.
(14a) The house with red and green parts/stripes/spots...

This reinterpretation of what is being spoken about when the word house is used is forced by the assumption that values on a dimension are mutually exclusive. It is perfectly fine to predicate two properties of an entity or modify a kind on the basis of values on two dimension. In these cases there is no need to reinterpret what the noun is used to talk about. In fact reinterpreting (15) and (16) as (15a) and (16a) is not legitimate. Reinterpretation of conjunctive predicates and modifiers like the ones in (13) and (14) is a consequence of the semantics of dimensions which does not

allow a single entity to have more than one value on a dimension. This rule of reinterpretation is similar to the rules of construal proposed by Jackendoff (1992).

(15) The house is red and expensive.
(16) The red and expensive house...

(15a) Parts of the house are red and parts of the house are expensive.
(16a) The house with red and expensive parts...

The fact that one is forced to re-interpret house in (14) shows that there really is a constraint against an entity having two values on a given dimension and so (13) and (14) actually present evidence for the single value constraint rather than evidence against it. Further constraints on values on dimensions will be discussed in Chapter 5.

Many of the adjectives that name values on a dimension, name values on continuous dimensions such as HEIGHT. The continuity implies that the dimension theoretically has an infinite number of values that could be named. It is obviously not a reasonable option to have names for each of the possible values on a continuous dimension. The manner in which languages deal with this problem is to name values on continuous dimension in relation to other values on the dimension rather than name particular values. One way to do this is to compare two specific values on a dimension. This leads to comparative constructions such as (17) and (18).

(17) Is this dog bigger or is that one?
(18) Which towel is wetter?

A second way in which a given value on a continuous dimension can be interpreted is with respect to a standard value. There are at least three ways in which the standard for comparison can be determined. First, the standard may be given by the kind in question. Thus, *big dog* in (19) means big for a dog or bigger than the average sized dog or most dogs.

(19) A big dog bit me.
(20) The big dog ate the food.
(21) Elephants are big.

Secondly, the standard may be given by the context. In sentence (20), *big* could mean the bigger of two dogs in the situation even though that dog may be small in comparison to most dogs. Finally, the standard may be given by consideration of values for other kinds; (21) expresses the fact that elephants are larger than most animals or kinds of things one may normally think about. This type of usage requires that the speaker and listener have a compatible contrast set; if planets, stars and galaxies are included then it no longer is true that elephants are big. For all three kinds of standards mentioned above it is necessary to look at the variation of values of the given dimension in order to determine the standard. What differs is which entities enter into the computation of the standard.

2.2.2 Adjectives that name values on dimensions and their arguments.

Adjectives that name properties of things semantically require an argument to name the entity of which the adjective names the property. How this argument is represented structurally is the subject matter of

syntactic theory; Questions such as whether something is an internal or external argument or whether the argument should be a sister of the adjective cannot be answered by looking at the semantics of adjectives. However, the fact that adjectives require an argument is due to the semantics of adjectives. Note that adjectives name properties of things rather than properties themselves. Properties themselves can be spoken about in the abstract rather than as applying to any particular entity. For example, one can talk about the fact that HEIGHT and WIDTH differ in the orientation on which they are defined without considering any entity's height or width. This is reflected in the fact that names for properties are nouns (e.g. *height, width, tallness* and *color*) and do not require arguments.

In addition to semantically requiring an argument, many adjectives are dependent on the argument for interpretation. For example, one cannot tell whether an adjective such as *big* applies to an entity such as a rat unless one knows what noun *big* is to apply to; if the entity is named as a rat then big would apply if the entity was big for a rat. If, on the other hand, the entity were named by the noun *animal*, then it would not be considered big because even the biggest rats are not big animals. Determining whether the entity can be called a rat, on the other hand, does not depend on its size with respect to different categories (is it big for a rat, is it big for an animal etc...).

The last example showed that in order to determine whether the adjective *big* applies to a particular entity one needs to know which noun the entity is named by. It has been argued that while this is true for relative or gradable adjectives such as *big*, the same is not true of absolute adjectives such as *red, male* or *dead*. Absolute adjectives are said to allow inferences of the following form (22).

(22) X is a red book.
All books are objects.
X is a red object.

X is a male dog.
All dogs are animals.
X is a male animal

This type of inference is not valid[2] for an adjective like *big*(23).

(23) X is a big book.
All books are objects.
#X is a big object.

The difference in the inferences supported by absolute and relative adjectives can be attributed to absolute adjectives' invariance in interpretations with different nouns.

The distinction between absolute and relative adjectives is conceptually clear: the former are interpreted in the same manner no matter what argument they take and the latter are interpreted only with respect to a particular argument. It turns out, however, that there are very few, if any, adjectives which are absolute. Color terms are often given as examples of absolute adjectives but it turns out that they are both graded and context dependent. Color terms are graded as can be seen from the fact that they appear in comparative constructions such as (24). In addition to being graded, color terms are also context dependent: the colors which count as red for an apple, for hair and for the sky are very different colors. Similarly, white and black skin are quite different in color from white or black paper.

(24) Which apple is redder?

[2] I will use the symbol # to indicate that an inference does not follow.

The context dependence of putatively absolute adjectives is best illustrated in the inferences that are licensed. One kind of evidence that has been offered for the context independence of absolute adjectives is that absolute adjectives support inferences of the form illustrated in (22). It turns out, however, that this type of inference is valid for only some noun contexts. For example, inferences using the adjectives *male* and *dead* in (25) are not valid.

(25) X is a dead dog.
<u>All dogs are physical objects.</u>
#X is a dead physical object.

X is a male dog.
<u>All dogs are physical objects.</u>
#X is a male physical object

(26) X is a black dog.
<u>All dogs are physical objects.</u>
X is a black physical object

The reason these inferences are not valid is that the properties named by *male* and *dead* do not apply to entities conceptualized as physical objects. Thus, even though a dead dog can be thought of as a physical object, it cannot be thought of as a dead object. Note that if an adjective names a property of dogs that can also be a property of physical objects, then the inference is fine (26). These facts indicate that even absolute adjectives are dependent on the noun to be interpreted. While the intersection of dead things and physical objects is well defined and not null (unlike the intersection of pregnant things and tables), *dead* cannot modify *physical object*. The anomaly has to be attributed to factors related to semantic interpretation because the sentences are syntactically well formed.

The preceding discussion has argued that the meanings of adjectives require arguments and that the arguments are also needed for the interpretation of most adjectives. Given the role of nouns in the interpretation of adjectives, it is important to analyze the ways in which adjectives and nouns interact semantically.

2.2.3 How adjectives combine with nouns semantically

Adjectives that name values on a dimension combine semantically with nouns in two ways: (i) Adjectives can *predicate* properties of the entity named by the subject NP, (ii) Adjectives can *modify* the kind named by the noun. The second type of adjectives combine with nouns to pick out a kind.

2.2.3.1 Predication

One of the ways in which adjectives that name values on a dimension semantically combine with nouns is by predication. In predication, an entity is stated to have the value named by the adjective on the relevant dimension. For example, sentence (27) states that a particular individual of the kind DOG has a value named by the adjective big on the dimension SIZE. The entities of which an adjective is predicated are specified independently of their value on the relevant dimension. Thus, the size of the dog is not relevant to how the dog being talked about in (27) is picked out. In addition to sentences like (27), there are other types of sentences which have the same predication relationship between the adjective and the noun. Examples of such sentences

are sentences with relative clauses (28), resultatives (29), depictives (30), and appositives (31).

(27) That dog is big.
(28) The dog that is big is mine.
(29) John scrubbed the table clean.
(30) John ate the meat raw.
(31) The flat table was a perfect place for the computer.

The adjectives in sentences (27)-(31) do not play a role in specifying which entities are being picked out. The adjectives are naming the value that the entities named by the subjects are said to have. Sentences (29) and (30) have semantics that relate the adjectives to the verbs in different ways. In (29), the direct object of the verb comes to have the property named by the adjective due to the action named by the verb being performed on the entity which is the direct object. In (30), the direct object of the verb has the property named by the adjective at the time of the action named by the verb. Despite this difference in the manner in which the adjectives and verbs are related semantically, both these sentences a predication relationship between the adjective and the noun. Finally, sentence (31) is likely to be interpreted as an appositive under normal circumstances and can be paraphrased as (32). This is clearly a predicative use of the adjective *flat*; the adjective names a value on a dimension and does not play a role in specifying which table is being talked about.

(32) The table which happens to be flat is a perfect place for the computer.

2.2.3.2 Modification

Adjectives that name values on a dimension can also be used to modify the entities picked out by the noun. Modification is the specification of individuals of the kind named by the noun on the basis of the value named by the adjective. In a sentence like (33), the noun *dog* is used to talk about individual dogs and the adjective *big* names some values on the dimension SIZE. Each of the individual dogs has one value on the dimension SIZE and the adjective and noun combine to name those dogs that have a value on the dimension SIZE which is named by the adjective *big*. Unlike predication, in modification the adjective plays a role in specifying which entities are being talked about.

(33) I like big dogs.

I want to distinguish my use of the term *modification* from the notion of *restriction* which is often used to talk about various kinds of modifiers (Siegal, 1980). Restrictive modifiers are those which help pick out a subset of the entities named by the noun they are modifying. Restrictive modifiers include prenominal adjectives and other manners of restricting the subject including relative clauses. Modification, is a special case of restriction, one in which the value on a dimension is used to pick out which entities are being spoken about. As will be seen, there are important ways in which modification and other types of restriction differ. The expressions that can name a value on a dimension are a subset of the expressions that can be used to restrict the range of entities being spoken about. Restriction is a purely semantic notion that does not make reference to any particular structural relation between the restrictive expression and the expression being restricted.

Modification, on the other hand, is defined in terms of a particular structural relation that holds between adjectives and the noun they are modifying. In English, modificational adjectives occur prenominally.

2.2.3.3 Adjectives that combine with nouns to name kinds.

In addition to adjectives that name values on a dimension, there are a small number of adjectives that combine with nouns to name kinds of things. Examples of adjectives like this are *alleged, former* and *corporate*. As discussed above, adjectives like *alleged* and *former* cannot name values on dimensions of the entities named by the noun with which they combine because the entities named by phrases like *alleged* NOUN or *former* NOUN cannot or do not have to be entities of the kind named by the noun. While this is the case for adjectives like *alleged* and *former*, it is not for adjectives like *corporate*. Adjectives like *corporate* always combine with a noun to name a subset of the entities named by the noun. In the present analysis, adjectives like *big* differ from *corporate* in that *big* is said to modify a kind when it appears in prenominal position but *corporate* helps picks out a kind when combined with a noun. The question that arises is why can't adjectives like *corporate* modify a noun?

Modification is proposed to be the specification of entities named by the argument on the basis of the value named by adjective. Given this, the question of why adjectives like *corporate* cannot modify nouns can be reformulated as why can't these adjectives name values on a dimension. Recall that values on a dimension are mutually exclusive and a given entity can have only one value on any dimension. It is unclear, however, on what

dimension *corporate* could name a value. A plausible hypothesis would be the dimension TYPE, since corporate lawyers are a type of lawyer and there are other types of lawyers like criminal lawyers and family lawyers and tax lawyers. The problem with such a proposal is that the putative values *criminal, tax,* and *family* are not mutually exclusive; a given lawyer could be all these types of lawyers. It seems like an entity named by a phrase like corporate lawyer is not restricted form having any adjective that names a value from modifying it; a corporate lawyer can also be tall, short, happy, sad, etc...

The distinction between prenominal adjectives that name values on dimensions and prenominal adjectives that help name a kind of thing is supported by a number of linguistic phenomena. First, unlike adjectives that name values on dimensions, adjectives that help name a kind of thing cannot appear as predicates (34),(35). This is because adjectives that appear as predicates state something about the entity named by the subject and do not help determine which entities are named by the subject. It follows that adjectives that help name a kind of thing should not be able to appear as predicates (35). Note that adjectives that help name kinds of things can appear *in* predicates (36), they just cannot appear *as predicates* themselves because they help name kinds of things and the predicate does not help determine what is being spoken about in a sentence like (35).

(34) John is tall.
(35) *John is corporate.
(36) John is a corporate lawyer.

Second, *big* and *corporate* combine with other prenominal adjectives in different ways. In English, when there is more than one prenominal adjective there is a preferred order in which the adjectives occur (more on this in Chapter 4). While the order in (37) is preferred, the order in (38) is acceptable in contexts in which there are some big bats and some small bats and of the big bats some are red while others aren't. One can always construct a context in which one can change the preferred order of adjectives such as *big* and *red* but the same is not true for adjectives such as *corporate* as can be seen in (39). This difference is predicted if adjectives such as *tall* do not name a kind of thing when combines with a noun but adjectives such as *corporate* do. The relationship of corporate lawyer to lawyer is more like the relationship of *cocker spaniel* to *dog* than *small dog* to *dog*.

(37) The big red bat.
(38) The red big bat.
(39) The tall corporate lawyer. / *The corporate tall lawyer.

Third, when adjectives that name values on a dimension modify a noun, they *always* combine with the noun to pick out a subset of the entities named by the noun. This is not the case with adjectives that combine with nouns to pick out a kind. Adjectives such as *alleged* and *former* do not pick out a subset of the entities named by the noun with which they combine but an adjective such as *corporate* does pick out a subset. The reason for the difference between adjectives that name values on a dimension and those that help pick out a kind is that the adjectives that name values on dimensions are constrained to name properties of the kind named by the noun whereas adjectives which do not name values on a dimension are not

so constrained. Fourth, adjectives that combine with nouns to pick out a kind of thing cannot be nominalized whereas adjectives that name values on dimensions can (Levi, 1978). For example, *nervous* has the nominalization *nervousness* when it is used to name a property as in (40) but cannot be nominalized when it combines with nouns to pick out a kind (41).

(40) The nervous man... / The nervousness of the man...
(41) The nervous disorder... / *the nervousness of the disorder...

A final difference is that adjectives that combine with nouns to pick out kinds MUST be discrete (42). The same is not true of adjectives that modify a noun (43). The condition that adjectives that help pick out kinds be discrete makes sense because there should be some set of conditions which would make an entity a member of the subkind and membership in a kind is not a graded concept. Thus relative adjectives such as big do not combine with the noun to pick out a subkind.

(42) The corporate lawyer. / *The more corporate lawyer.
(43) The big house. / The bigger house.

This leaves open the possibility that discrete adjectives such as *dead* which are not relative but can appear both prenominally and predicatively help pick out a kind when in prenominal position. I would argue that they do not because adjectives such as *dead* pattern with the relative adjectives in that they can appear in both prenominal and predicate positions and they do not combine with adjectives that help name kinds of things as can be seen in (44) (Levi, 1978).

(44) *The dead and corporate lawyer.

As mentioned above, adjectives like *corporate* do not name a value on a dimension and thus cannot appear as predicates (35). Note that some of the adjectives that combine to name a kind of thing are homophonous with adjectives that name values on dimensions. An example of such an adjective is *criminal* which can appear prenominally as well as predicatively (45). It's prenominal use can be modificational or it may help name a kind of thing. The reading that allows one to name a kind of thing cannot appear in predicate position (46).

(45) The criminal action was..../ The action was criminal.
(46) The criminal lawyer was hungry. *The lawyer was criminal.

2.3 Semantic knowledge of adjectives and how they combine with nouns.

To summarize, I have motivated a semantic analysis of adjectives in which there are two kinds of adjectives. The first kind of adjectives name properties of things. These adjectives name values along dimensions. A dimension is proposed to consist of a set of mutually exclusive values. Adjectives that name values on dimensions are proposed to combine with nouns in two ways. In predication, these adjectives name the value that the entity named by the subject has on the relevant dimension. In modification, the value named by the adjective is used to pick out the objects of the kind named by its argument. The second type of adjectives are those that do not name values on dimensions but combine with nouns to pick out kinds of things. In the rest of this thesis I will show how this knowledge of the

semantics of adjectives can explain a number of linguistic and psycholinguistic facts as well as motivate heuristics for the acquisition of adjective meanings.

Chapter 3

Modification

3.1 Introduction

In this chapter, I will show how the semantics of modification can be used to explain a number of restrictions on structures and how they are interpreted. In the previous chapter, it was proposed that adjectives that name values on a dimension interact with nouns in two ways. In predication, adjectives state the value a given entity is said to have on the dimension on which the adjective names values. The adjective does not play a role in determining which entities are being spoken about. In modification, the value named by the adjective is used to help specify which entities are being spoken about. Whereas in English, modificational adjectives occur prenominally, the exact syntactic structure that represents this semantic relation is not relevant for the purposes of the present analysis. What is important here is that there is a structural relation between adjectives and nouns which can be interpreted as modification and that this structural relation is different from the relation that holds between an adjective and a noun when the adjective is predicating information about the entities named by the noun. An example of a modificational adjective is given in (47). In this sentence, the value named by *big* is used to identify which table is being spoken about.

(47) The big table is mine.

3.2 Restriction against the modification of single entities.

Because modification uses the value named by an adjective to determine which entities are being spoken about, it follows that it should not be possible for an adjective to modify expressions that are used to name single entities. This prediction receives support from the ungrammaticality of sentences (48)-(50).

(48) * Big [the dog]$_{NP}$ is hungry.
(49) * Big [Fido]$_{NP}$ is hungry.
(50) * Big [he]$_{NP}$ is hungry.

In each of (48)-(50), there is only one entity being named by the NP[1] and thus the value the entity has on the dimension cannot be used to identify what is being spoken about.

3.3 Semantics of plural NPs and modification.

While the semantics of modification may be able to explain the ungrammaticality of (48)-(50), it seems that it will be unable to explain why sentences (51)-(53) are ungrammatical because in each of these sentences, more than one entity is being spoken about and yet the adjective cannot modify the expression used to talk about a number of entities. This suggests that the reason why (48)-(50) are ungrammatical is because of a syntactic restriction on adjectives occurring in a structural position which would allow modification of an NP. While this is a possibility, I would like to argue that the semantics of plural NPs prevent them from being modified.

---

[1] For the purpose at hand, it does not matter whether one adopts a DP analysis or not.

(51) * Big [the dogs]$_{NP}$ are hungry.
(52) * Big [some dogs]$_{NP}$ are hungry.
(53) * Big [all dogs]$_{NP}$ are hungry.

Plural noun phrases such as *the dogs, some dogs* and *all dogs* usually can be interpreted in two ways. The first interpretation is a *distributive* interpretation. In the distributive interpretation, the noun phrase can be rewritten as a conjunction of all the entities named by the noun. For example, the distributive interpretation of (54) is that each of the dogs named by the NP *the dogs* (dog_1, dog_2, ..., dog_n) is hungry (55). In the distributive interpretation, one entity is being spoken about at a time and each of those entities has only one value on a given dimension and thus cannot be modified. This explains why (51)-(53) are ungrammatical.

(54) The dogs are hungry.
(55) Dog_1 is hungry, and dog_2 is hungry, and....., dog_n is hungry.

The second type of interpretation of plural noun phrases is a *collective* interpretation. In a collective interpretation, one is talking about the entities named by the noun as a single entity. For example, it is the dogs as a single entity that are making a lot of noise in (56), rather than each dog individually. It may be the case that each dog individually is also making a lot of noise but in the collective reading it is the group of dogs as a single entity that is making the noise. Because the collective interpretation of an NP treats the entities picked out by the noun as a single entity, the NP itself cannot be modified.

(56) The dogs were making a lot of noise.

A collective interpretation of an NP is used to speak about a *particular* or unique set. While it is true that a given NP can be used to talk about many different sets of dogs, it can be used to talk about only one particular set of dogs in a given context. For example, even though the NP *the dogs* can be used to talk about many different sets of dogs, it is used to talk about a particular set of dogs in (56) and not any of the other sets that *the dogs* could have been used to talk about.

Evidence for the claim that NPs pick out a unique set of entities named by the noun comes from the fact that one cannot quantify NPs. In each of (57)-(59) it is perfectly possible to count the dogs but it is not possible to count the particular set that the NP is used to talk about.

(57) * How many the dog(s) are in the room?
(58a) There are some dogs in the kitchen and some dogs in the bathroom.
(58b) * There are two some dogs in the house.
(58c) There are two sets of dogs in the house.
(59) * Is this the same the dog I saw yesterday?

Because the collective interpretation of NPs is used to talk about a particular set, one cannot modify NPs. One can, however, predicate adjectives of collective NPs (56). This is because predicating information does not require there to be more than one value on a dimension and thus more than one entity. Distributive NPs cannot be modified because they are used to talk about single individuals one at a time.

3.4 Semantic vs syntactic restrictions on prenominal adjectives.

The preceding discussion has provided a semantic explanation for why it is not possible for adjectives to modify NPs. NPs are used to talk about a particular set of entities or about particular individuals one at a time and thus cannot be modified. The question arises whether there is any reason to prefer the semantic explanation to a syntactic explanation. It is possible that the semantic generalization about NPs is a consequence of a syntactic restriction and that the semantic restriction does not have any independent motivation. Alternately, the semantic restriction that is used to explain the restriction on adjectives modifying NPs is independently motivated. In the following sections I will argue that the semantics of modification which are used to explain why adjectives cannot modify NPs are also needed to explain restrictions on where adjectives can occur *within* an NP and that the restrictions on modification also predict how prenominal adjectives are interpreted.

3.4.1 Restrictions on adjectives within an NP.

In addition to the restriction on adjectives modifying NPs, there are also some restrictions on where an adjective can occur within an NP. For example, (60) is perfectly acceptable, but (61) is not even though the adjective occurs within the NP. The explanation for why (61) is ungrammatical is that the phrase *three dogs*, like an NP, is used distributively to talk about three dogs one at a time or to talk about a particular set of dogs and thus cannot be modified for the same reason that adjectives cannot modify NPs. Note that the sentence would be grammatical if one could construe the phrase *three*

dogs as naming an individual of a kind. This would allow the adjective *big* to modify. A context in which (61) may be regarded as grammatical is one in which dogs come in triples and these triples are called three dogs and these triplets of dogs vary in size. It may be argued that because there are contexts in which sentences like (61) are grammatical, there is no restrictions on adjectives modifying phrases inside NPs, only strong biases against adjectives modifying number phrases in an NP. It turns out, however, that sentences like (62) and (64) are ungrammatical in any context.

(60) The three big dogs are hungry.
(61) * The big three dogs are hungry.
(62) * The big one dog is hungry.
(63) The fastest big dog is hungry.
(64) * The big fastest dog is hungry.

The ungrammaticality of (62) and (64) is due to the fact that the phrase *one dog* and *fastest dog* necessarily name a single dog and thus cannot be modified. The ungrammaticality of sentences (62) and (64) provides further evidence that there are restrictions on structure that are due to the properties of modification and that the properties of modification can account for restrictions on the positions of adjectives within NPs as well as why adjectives cannot modify NPs.

3.4.2 Interpreting adjectives as being modificational.

In addition to explaining the restriction on adjectives preceding NPs and restrictions on where an adjective can occur inside an NP, the semantics

of modification also predict when an adjective may be interpreted as modifying a noun and when it has to be interpreted otherwise.

While it is true that having more than one entity is a *necessary* condition for modification to occur, it is not a *sufficient* condition. For modification to occur there have to be at least two *values* on the dimension on which modification is taking place. If all the entities named by the argument of the adjective have the same value on the dimension being modified then modification cannot take place. This is because it is not possible to pick out an entity of the kind named by the noun on the basis of the value named by the adjective on the relevant dimension if all the entities named by the noun have the same value. An example of such a situation is given in (65). Under most circumstances, the adjective *flat* cannot be taken to be modifying *table* because all tables are flat. The adjective is interpreted as an appositive and the meaning of (65) corresponds to that of (66). Another example of adjectives being interpreted appositively is nicknames (67). A nickname like *Big Jim* in (67) has it's origins in a statement such as (68).

(65) The flat table is the best place to put the computer.
(66) The table which happens to be flat is the best place for the computer.
(67) Big Jim is a helpful fellow.
(68) Jim, who is big, is a helpful fellow.

Given that a prenominal adjective can be interpreted as an appositive when there is no more than one value represented on the dimension to be modified, the question arises as to why this isn't possible when an adjective precedes an NP? The difference between NPs in English and proper names and nouns is that NPs are used to talk about an individual, a single set (collective reading of plural NPs) or a number of individuals one at a time

(distributive reading of plural NPs). Each of these uses of an NP precludes it from the possibility of modification. Furthermore, it is not necessary to know which noun is the head of the NP or what its relation is to the adjective. The impossibility of modification is due to the semantics of NPs rather than nonlinguistic knowledge regarding the particular kinds of entities named by the noun and what value the entities usually have on the dimension named by the adjective.

The decision regarding when an adjective preceding a proper name or a noun should be interpreted appositively depends on one's knowledge of the context and nonlinguistic knowledge about the entities named by the noun or proper name. For example, one can only know that *flat* has to interpreted appositively in (65) if one knows that tables usually do not vary in flatness. It is for this reason that one cannot predict whether *flat* should be interpreted in a modificational manner or not in (69). On the other hand, it is clear that *flat* could not be modificational in (70) even though one does not know the meaning of *blicket*. Furthermore, there is no context in which *flat* could be interpreted in a modificational manner in (70).

(69) The flat blicket is a good place to put the computer.
(70) * Flat the blicket is a good place to put the computer.

Unlike NPs, proper names can be modified given the proper context (71). It is perfectly acceptable to modify a proper name when there is more than one individual with the same proper name. The reason why one usually interprets adjectives preceding proper names as appositives is that in most contexts, a proper names picks out a single individual of a given kind. This uniqueness is a pragmatic assumption and is not a semantic property of

proper names because proper names can be quantified (72)-(73). Recall that one cannot quantify NPs (74)-(75).

(71) Big John is an engineer and little John is a doctor.
(72) How many Johns are in this class?
(73) I knew two Johns in third grade.
(74) * How many the boys are in this class?
(75) * I knew two some boys in third grade.

The fact that an appositive interpretation is only possible in those contexts in which modification is possible (e.g. (69)) and not in contexts in which modification is not possible (e.g. (70)) suggests that the default interpretation of the prenominal adjective is modificational.

3.5 Cross-linguistic variation in restrictions on prenominal adjectives.

The definition of modification given in Chapter 2 was the specification of entities of the kind named by the noun on the basis of the value on the dimension named by the adjective. Given that this is a semantic definition, it should not be subject to cross-linguistic variation. The restrictions that arise from modification should be the same across languages. Thus, it should only be possible for an adjective to modify expressions which are used to speak about at least two entities which have different values on the dimension being modified.

Japanese seems to provide counter-examples to the explanation given here. In Japanese, it seems that it is possible to modify words that are translated in English as *he* and *this*.

(76) kierei-na ko-no hana
'pretty this flower'

(77) Kinoo Taroo-ni atta ka-i?
'Did you meet Taro yesterday?'

Un, demo kinoo-no kare-wa sukosi yoosu-ga hendat-ta.
'Yes, but yesterday's he was somewhat strange.'
(Fukui & Speas, 1987)

(78) ookii kare
'big he'

Though these sentences seem to provide counter-evidence to the explanation of the semantics of modification given here, it turns out that they do not. In none of (76)-(78) does the putative modifier modify the noun.[2] (76) and (78) are interpreted appositively and cannot be used to identify an entity on the basis of the value named by the adjective. Similarly, (77) has the same meaning as when *kinoo-no* doesn't have a genetive marker, indicating that *kinoo-no* is not modifying *kare-wa*. Thus these examples are consistent with the analysis of modification that states that it is not possible to modify expressions that name a single entity. One difference between Japanese and English is that adjectives can appear before pronouns and demonstratives in Japanese but not in English. Fukui and Speas (1986) suggest that lexical categories project to X' rather than X" as is the case in English. This difference in the projection of lexical categories may account for the fact that modifiers can precede expressions which are XPs in English. This syntactic difference may also have subtle semantic consequences but one would have perform careful investigation to determine if this is indeed the case.

In sum, it is predicted that the restrictions on modification are universal and putative counterexamples are likely to be interpreted as

---

2 I am grateful to Shigeru Miyagawa and Daishi Harada for these judgements.

appositives or have a different structural relation than the relation that holds between adjectives and nouns that are interpreted as being modificational.

3.6 Summary

The present chapter has shown how the semantics of modification can be used to explain a number of constraints on syntactic structure and also be used to predict when a prenominal adjective has to be interpreted appositively. The semantic of modification provide a motivation for why the structure of DP is as it is. It may turn out that there are other, syntactic, reasons why adjectives cannot precede DPs in English but the existence of these syntactic explanations does not in itself pose a problem for the semantic constraint proposed here. It is quite possible for a given structure to be unacceptable because it violates a number of constraints (e.g. the case filter and ECP). The semantic explanation was motivated for independent reasons and can explain why adjectives cannot modify certain N-bars (ones that necessarily name only one entity e.g. *fastest dog, one dog*), when prenominal adjectives are to be interpreted as appositives and restrictions on adjective ordering. The manner in which structural relations are semantically interpreted provides a source of constraint on possible structures and may lead to the simplification of syntactic theories. Furthermore, semantic constraints of this sort may be able to explain certain aspects of syntactic structure that are usually left unexplained. For example, the constraints on modification may help explain why APs do not dominate DPs.

In sum, the proposal that there may be constraints on structures due to the manner in which they are interpreted does not imply that there are

constraints on syntactic structure that are independent of the constraints imposed by the manner in which certain structural relations are interpreted. Furthermore, the semantic analysis is crucially dependent on structural notions and so it should not be seen as a proposal that is inconsistent with syntactic theories. In the next chapter, I will look at how the semantics of modification can be used to explain the ordering of multiple adjectives.

Chapter 4

Adjective Ordering

4.1 Introduction.

The analysis of adjectives presented in Chapter 2 states that adjectives that name values on dimensions can be used to predicate information about the subject or they can be used to pick out entities named by the noun on the basis of the value named by the adjective. In the predicative use, adjectives name the value the entity in question has on a particular dimension. The adjectives that are predicated of the entity do not play any role in deciding which entities are being spoken about. In the modificational use, however, the adjectives do play a role in deciding which entities are being spoken about.

Entities named by nouns have many dimensions on which they have values. It follows that it should be possible to predicate more than one adjective of entity and it should also be possible to decide which entities are being spoken about on the basis of values on more than one dimension. This is in fact the case as can be seen in (79) and (80).

(79) The dog in the corner is big and furry.
(80) I want the big and furry dog.

4.2 Types of modification.

In (79), the adjectives are naming the values the dog in the corner has on the dimensions SIZE and FURRINESS. In (80), the values named by *big*

and *furry* are used to specify which dogs are wanted. Note that in the two sentences given above, both *big* and *furry* name values on dimensions of members of the kind DOG. The determination of the values on the dimensions SIZE and FURRINESS depend on the kind DOG and are independent of each other. When an adjective modifies a noun in a manner in which values other dimensions cannot affect values on other dimensions, then the modification is said do be parallel modification.

In addition to parallel modification, two or more adjectives can modify sequentially. An example of *sequential modification* is given in (81).

(81) I want the big furry dog.

In (81), *big* can modify the entities named by *furry dog* rather than the entities named by *dog*; If the range of sizes for dogs is different than the range of sizes for furry dogs, then the entities picked out by the NP will differ when the modification is sequential. Unlike the case of parallel modification, the values on the two dimensions are not independent of each other; what counts as big can be affected by values on the FURRINESS dimension.

4.3 Constraints on adjective ordering.

The possibility of non-parallel modification means that there has to be some way in which a language encodes the order in which modification is to occur. In English, the order of non-parallel modification is encoded in the serial order in which adjectives appear in prenominal constructions. The semantics of adjectives and modification offer an explanation for (i) why

there has to be an ordering of adjectives, (ii) why the ordering is what it is, and (iii) why alternate orders are acceptable in certain contexts.

As discussed above, one reason the order of adjectives modifying a noun in a non-parallel manner must be specified is the entities named by the NP may differ depending on that order. A second reason why the order of modification is important is that modification along a given dimension is only possible if there is more than one value on the dimension in the set of entities being modified. For example, *big* can modify *red bats* in (82), only if red bats vary in size. If it turns out that all the red bats are of the same size, *big* cannot function as a modifier in (82); it would function as an appositive. Similarly, *red* can only modify *big bats* in (83) if the big bats that are being spoken about vary in color.

(82) big red bats
(83) red big bats

Because different orders of modification can lead to different entities being picked out it is also likely that different orders of modification will pick out different numbers of entities. In some cases, one of the orders may not name any entities. This would be the case in (82) if none of the red bats could be considered big. Similarly, no entities would be picked out by (83) if none of the big bats are red. If no entities are named by an expression, then it trivially doesn't allow for any further modification. For example, if there are no big bats that are red,[1] then one certainly cannot talk about the

[1] Positing that there are red big bats in some context does not mean that one can modify the term used to name them. One would only be able to modify the term used to name red big bats if they showed variation on some dimension.

expensive red big bats in contexts in which *expensive* modifies *red big bats*.[2] This implies is that there is a constraint on the order of modification which would require that:

(I) The order of two modifiers is such that if one ordering leads to no entities being picked out and the other does not, then the order that does not pick out any entities is disallowed and the order that does pick out some entities is allowed.

Of course there will be many pairs of modifiers for which both orderings will result in some entities being picked out. How should one decide between these possible orderings? I return to this question below but I first discuss other constraints on the ordering of sequentially modifying adjectives.

Clearly, one cannot modify an expression that does not name anything. Also, one cannot modify a single entity. As discussed above, a single entity has only one value on each of its dimensions and thus cannot be modified. This is the reason why sentences such as (84) are ungrammatical. Recall that modification can also fail when more than one entity is named by the argument but all the entities have the same value on the dimension in question. This is why *flat* fails to modify *table* in (85) and is interpreted as an appositive in normal circumstances. To summarize, for modification to be possible there have to be at least two entities being considered that do not have the same value on the dimension being modified.

(84) * The big one dog is mine.
(85) The flat table is mine.

---

[2] It is important to note that when I say there are no entities of a given kind I am not making a claim about the world but rather about the mental model with which one is operating. Mental models may be of actual or counterfactual situations.

The likelihood of two entities having the same value on a dimension is a function of the number of different values there are on that dimension. If there are a greater number of values on Dimension A than Dimension B, then it is more likely that two entities picked at random will have different values on Dimension A than on Dimension B. This implies that one should modify along Dimension B before one modifies on Dimension A. This is because given that A has more values than B, the likelihood that there will be more than one value of A represented after modification on B, is greater than the likelihood that there will be more than one value of B represented after modification on A. This suggests the following constraint on the order of modification.

(II) If Dimension A has a greater number of values than Dimension B, then one should modify along Dimension B before modifying along Dimension A.

This constraint on sequential modification states that the order in which adjectives occur will be constrained by the number of values there are on the dimensions in the context being considered. Note that continuous dimensions such as SIZE and TEXTURE have an infinite number of values but they differ in that at some point a surface will become too uneven to be called rough but an object can never become large enough such that the adjective big would no longer apply. This is why SIZE precedes TEXTURE even though they are both continuous dimensions.

Constraint (II), is proposed to be a constraint on the ordering of sequential modification by adjectives.[3] If the constraint is in fact a constraint on adjectives that modify sequentially rather than a more general constraint on adjective ordering, then it should be possible to violate the constraint on ordering when modification is parallel rather than sequential. Similarly, if adjectives are not modificational, then they should not be subject to the constraint on the ordering of sequential modificational adjectives. In the rest of the chapter I will present evidence supporting the analysis presented here and I will also review previous proposals for constraints on adjective ordering and how the proposals are related to the current proposal.

4.4 Ordering of modificational adjectives.

As mentioned above, there is a particular order in which prenominal adjectives typically appear when there is more than one prenominal adjective in English. For example, the order in (86) is natural and the one in (87) is odd except in a most contexts. In (86), *big* modifies (at least in some cases) *red plastic bat* rather than *bat*; thus, the big red plastic bat may be relatively small for a bat and yet big for red plastic bats.

(86) I want the big red plastic bats.
(87) I want the plastic big red bats.

[3] Because constraint (II) is a constraint on the ordering of sequential modification, it applies only to adjectives that name values on dimensions. Ordering constraints on adjectives that help name kinds of things (e.g. former corporate lawyer/*corporate former lawyer) will have to be explained in some other way.

Sproat and Shih (1986) provide cross-linguistic evidence from a range of languages[4] showing that if a language allows direct (sequential) modification then the order of the modifiers is the same as in English. In addition, they showed that the constraints on the order of modifiers does not hold when the modification is not direct. Sproat and Shih's data from Mokilese and Thai shows that distance from the head noun is what is relevant to ordering rather than linear order. This finding is supported by Martin's (1969) data on the ordering of post nominal adjectives in Indonesian.

4.5 Non-canonical orderings.

Constraint (II) orders adjectives on the basis of the number of values on a dimension. A prediction that follows from this constraint is that non-canonical adjective orders should be possible if one considers contexts in which the number of values on the relevant dimensions does not correspond to the number values usually represented on the dimensions in question. For example, it may be the case that in a given context there are more values on Dimension B than on A even though in general Dimension A has more values. In this context it would be appropriate to have the adjective ordering BA even though the standard ordering is AB.

The prediction that non-canonical adjective orders should be possible if one considers contexts in which the number of values on the relevant dimensions does not correspond to the number values usually represented on the dimensions in question is borne out in sentences like (24).

---

[4] The languages that they studied were English, Mandarin, Japanese, Greek, Kannada, Arabic, French, Thai, Mokilese, and Irish.

This sentence is acceptable in contexts in which the range of variation of values on the dimensions being considered is different from the amount of variation which is found on those dimensions in general. Specifically, it is okay to say *plastic big red bat* in contexts where there is variation in the type of material of which bats are made but the size and color of the bats is held constant (Danks & Glucksberg, 1971). Thus (87) would be appropriate if one wanted to pick the plastic bat out of big red bats in a context in which what counts as big for a red bat is different from what counts as big for a plastic bat or a red plastic bat.

4.6 Ordering of non-modificational adjectives.

The previous section provided evidence that when there is more than one adjective modifying in a sequential mannner in English, the ordering of the adjectives is subject to constraint (II). If the constraint is a constraint on the ordering of sequential modificational adjectives rather than a more general constraint on adjective ordering it should follow that adjectives that do not have a modificational role or modify in a parallel manner should not be subject to this constraint. The current section presents evidence that the constraints due to the nature of sequential modification are not present in the ordering of non-modificational adjectives and parallel modifiers in English.

The most obvious instance of non-modificational adjectives are adjectives that appear in predicate position as in (88) and (89). Both orders of the adjectives are equally acceptable and it seems that considerations of focus determine which of the adjectives is mentioned first. Crucially,

considerations of the number of values on the various dimensions does not determine whether a given ordering of adjectives is acceptable or not.

(88) That bat is big, red, and plastic.
(89) That bat is plastic, red, and big.

Another context in which adjectives are non-modificational is when they are applied to proper names. Because proper names usually cannot be modified, the adjectives are interpreted as appositives and the order of the adjectives is determined by considerations other than those relevant to modification, possibly considerations of focus or salience (90)-(91). The number of values on the dimensions in question represented in the current context *cannot* be relevant to the ordering of the adjectives because a given individual can only have one value on each dimension.

(90) Big Old Jim is coming this way.
(91) Old Big Jim is coming this way.

Finally, another manner in which English indicates that prenominal adjectives are not modifying sequentially is in the phonology. Prenominal adjectives which have pauses between them are interpreted in either a parallel modificational manner (Beckman & Pierrehumbert, 1986) or as appositives. This is illustrated in (92) and (93).

(92) The big, red, plastic bat is mine.
(93) The plastic, big, red bat is mine.

4.7 Other accounts of adjective ordering.

Previous research on the constraints on the ordering of prenominal adjectives has suggested that various semantic (Bever, 1970; Danks & Glucksberg, 1971; Martin, 1969; Sproat & Shih, 1986), syntactic (Annear, 1964; Vendler, 1968) and processing variables (Bever, 1970; Martin, 1968; Crain & Hamburger, 1992) may be relevant. In the current section I review these proposals and how they relate to the proposal being made in this chapter.

Bever (1970) proposes a processing explanation for the ordering of prenominal adjectives in English. He observes that the order of prenominal adjectives is such that the closer the adjective is to the noun the more "nounlike" it is. The more "substantive" or "concrete" the quality that an adjective refers to, the more "nounlike" the adjective is. Following Vendler (1968), Bever argues that adjectives that name substances such as *plastic* are more "nounlike" than color adjectives like *red* in that substance terms occur in more constructions as nouns than do color terms. In support of this claim he offers the following type of evidence showing that substance adjectives are more "nounlike" than color adjectives and that *red* is more "nounlike" than *large*.

(94) a) Red is a color; redness is nice.
b) Plastic is a substance; plasticity is nice.
c) *That is made out of red.
d) That is made out of plastic.
e) *The red broke.
f) The plastic broke.
g) ?Reds are of variable quality.
h) Plastics are of variable quality.

(95) a) Red is my favorite color.
b) *Large is my favorite size.
c) He splattered some red on me.
d) *He splattered some large on me.

e) Red and blue and green are colors.
f) ?Large and enormous and tiny are sizes. (Bever, 1970; pg 324-25)

Bever argues that the ordering of prenominal adjectives in terms of "nounlikeness" is due to a perceptual strategy that is used to segment sentences. Specifically, one can tell that the noun phrase is complete as soon as one encounters a word which is less "nounlike" than the previous word. There are a number of problems with this proposal, however. First, the proposal does not give an account for why alternate orderings of sequentially modificational adjectives are sometimes acceptable and the contexts in which they are acceptable. Second, the proposal does not give an account for why it is possible to violate the canonical ordering of adjectives if the adjectives are modifying the noun in parallel rather than sequentially. Finally, the notion of "nounlikeness" does not seem to play a role in explaining any other linguistic phenomenon. Furthermore, the strategy cannot be used to predict in which contexts a prenominal adjective has to be interpreted as an appositive rather than a modifier nor can it be used to explain other restrictions on the positions in which adjectives can occur (e.g. why **big the dog* and **the big one dog* are unacceptable). Thus while the strategy of looking at the relative "nounlikeness" of adjectives may be useful as a processing strategy, it leaves a number of seemingly related phenomena unexplained.

Martin (1969) investigated a number of semantic factors that may be relevant to the ordering of prenominal adjectives. He asked subjects to rate a number of adjectives on the semantic dimensions, definiteness of denotation, absoluteness, imagery, and substantiveness. The dimensions were defined as follows: *definiteness of denotation* is the degree to which

adjectives are definite in meaning to the class of all nouns. This notion is similar to the notion of the number of values an adjective names on a dimension. Those adjectives which are definite in denotation have the same value for a large class of nouns where as those that are less definite in denotation have may different values for the class of all nouns. *Absoluteness* was defined as the number of comparisons needed in order to determine whether a given adjective applies to an entity. It is proposed that an adjective like *black* requires no comparisons but that an adjective such as *big* requires many comparisons. *Imagery* was defined in terms of the adjective's ability to invoke images and *substantiveness* was defined in terms of how essential subjects thought the property named by the adjective was to the noun they modified's meaning. The dimension that correlated most highly with preferred adjective order was definiteness of denotation: r = .73. This result supports the account of adjective ordering presented here. As mentioned above definiteness of denotation is equivalent to the number of values that are named by an adjective. Adjectives, that are higher in definiteness of denotation (have fewer values named by the adjective) are preferred closer to the noun. Martin, however, posits a processing constraint to account for the ordering of prenominal adjectives. He proposed that the accessibility of an adjective is related to the order in which prenominal adjectives occur. He proposes that accessability is related to reaction time and that adjectives that are easily accessed will occur closer to the noun because the noun phrase is built from the noun. The plausibility of this account of processing aside, there are a number of facts that such a theory will be unable to account for. First, it does not provide an account of why alternate orders are sometimes acceptable and the conditions under which they are acceptable. Second, such an account does not explain why the preferred order can be violated in

parallel modification. Third, there is no explanation for why the preferred order can be violated with predicational adjectives and parallel modification without a change in context relating to the number of values represented on the given dimensions. The account of adjective ordering presented here accounts for all these facts as well as the experimental results.

Sproat and Shih (1986) propose that the ordering restrictions on direct (sequential in the terminology used here) modification are cognitively based and seem to relate to the "apparentness" of the adjective. Apparentness is said to be related to the number of comparisons that are necessary to determine whether a particular adjective applies to an entity. It is suggested that adjectives like *red* can be attributed to an entity by only looking at low level properties of the object such as the surface reflectance and that the kind of object does not have to be identified. Adjectives like *large*, however, are less apparent because they require identifying the kind of entity and making comparisons to other entities of that kind. In support of this model, Sproat and Shih (1986) found that the strongest preferences in adjective orders occur between adjectives that differ in predicativeness. Predicative adjectives are those adjectives that allow inferences such as the one shown in syllogism (96) (Kamp, 1975). Sproat and Shih found that when there is a predicative and a non-predicative adjective modifying a noun there is a strong preference for the predicative adjective being closer to the head noun. This fact is from the absoluteness of predicative adjectives. The data also support the account of adjective ordering presented here; predicative adjectives are relatively invariant across contexts and thus have a single or very few values that are named by them, whereas non-predicative adjectives vary in their interpretation in different contexts and thus have many different values that they name.

(96) All mice are mammals.
Freddy is a male mouse.
Freddy is a male mammal.

Even if one accepts this analysis of how adjectives are processed there are a number of problems with this proposal. First, it gives no explanation for why apparentness is relevant to adjective ordering though it is suggested that the differences in the amount of computations may be responsible for the order in which the adjectives can appear. Second, there is no explanation for why the ordering does not have to hold for parallel modification and predication. Third, there is no explanation for why alternate orders are sometimes acceptable and how they are related to the number of values on the dimensions in the context being considered. The account of adjective ordering presented here accounts for all these facts as well as the ordering preferences between adjectives of different predicativeness.

Annear (1964) and Vendler (1968) both propose syntactic accounts of the ordering phenomenon. Annear (1964) divided adjectives into three classes and predicted the order in terms of the relative positions of these classes of adjectives. As Annear notes, this classification is post-hoc and does not help explain any other phenomena. Vendler (1968) presents a theory or adjective ordering in terms of the order in which transformations that apply to different adjectives apply. Once again the theory is purely descriptive because the order of transformations isn't relevant to any other linguistic phenomena.

Danks and Glucksberg (1971) propose an account of adjective ordering that is similar to the account given here. They propose that the ordering of adjectives is based on their discriminative function in a given context. The canonical order of adjectives is proposed to correspond to the most frequent

cases of the discriminative uses of adjectives. While this proposal shares some aspects with the current proposal, there are a couple of key differences. First, the current proposal provides an explanation for why some adjectives may more often play a discriminative role. Adjectives that have more values on a dimension are more likely to be able to discriminate between two entities than dimensions that have fewer values on a dimension. Second, the notion of the number of values on a dimension is relevant to phenomena other than the ordering of adjectives (e.g. why one cannot modify certain types of expressions **big he, *the big one dog*; when one has to interpret an adjective as modificational) while the notion of discriminability is a function of the context and it is not clear that it plays a role in explaining other linguistic phenomena.

4.8 Summary.

The semantics of modification was used to motivate an account of (i) why there has to be an ordering of adjectives when there is sequential modification, (ii) why the canonical order is what it is, and (iii) when the order can be violated. It was also shown that this order does not have to hold for parallel modification or predicative adjectives. The account of adjective ordering received support from psycholinguistic (Martin, 1969) and linguistic data (Sproat & Shih, 1986) and provides further evidence for the proposed semantics of adjectives.

Chapter 5

Differences Between Modification and Predication

5.1 Introduction.

In previous chapters, it was proposed that adjectives in predicate position name the value that the entity named by the subject is stated to have on the dimension named by the adjective. While this definition of predication is adequate to describe the relationship between the subject and the adjective in sentences such as (97), it fails to capture the similarity of this relationship to the relationship between the subjects and the predicates in (98) - (100). In (97), the adjective *tall* states the value that the subject has on a the dimension HEIGHT. In each of (98)-(100), the predicate is stating something about the subject but this information is not a value a dimension. Given the parallel structure between the subject and the predicate phrases in (97)-(100), it seems reasonable to generalize the notion of predication to one in which the predicate states something about the subject. This is the standard analysis of a predicate.

(97) The man is tall.
(98) The man is proud of his son.
(99) The man is in the corner.
(100) The man is running faster than a horse.

This more general analysis of predication acknowledges the similarities between different types of predicates. The specific interpretation regarding the relationship of the entity named by the subject and what is stated by the

predicate will depend on the semantics of the specific type of predicate. When the predicate is simply an adjective then the predicate names a value on a dimension.

Given that the structural relation between the subject and the predicate allows one to state things about the subject that are not the value that the subject has on a dimension, one can ask whether the same is true of the structural relation that holds between prenominal adjectives and the nouns they modify. Can expressions which are not adjectives appear in the same structural position that prenominal adjectives do?

As can be seen from (101)-(104), not all the phrases that can appear as predicates can appear prenominally. Given the difference in the acceptability of (101) and (102)-(104), the question arises as to what kinds of elements can appear prenominally. Why is it that the predicate in (97) can appear in a prenominal position (101) but the predicates in (98)-(100) cannot?

(101) The tall man ...
(102) * The proud of his son man ...
(103) * The in the corner man ...
(104) * The running faster than a horse man ...

5.2 Restrictions on prenominal phrases.

At first glance it seems that the explanation for the difference between (101) and (102)-(104) is syntactic and that there is a constraint against phrases appearing before the noun. Pinker (1984) has noted that this explanation is inadequate because of the acceptability of sentences like (105) in which the AP *very tall* appears prenominally. In addition to sentences like (105) which have an intensifier in the AP, sentences with prenominal measure phrases

are also acceptable (106). Note that even if one has an analysis in which *very tall* or *ten foot tall* are not APs but are some other kind of XP, the XP would have scope over the AP and thus the new analysis would not help distinguish the difference in acceptability of the APs in (101) and (102).

(105) The very tall man is a basketball player.
(106) The ten foot tall man should be a basketball player.

The difference between (101), (105)-(106) and (102) is that the AP in (102) has an internal argument but the ones in (101), (9105 and (106) do not. One possibility concerning the correct generalization is that a phrase can occur prenominally if the head of the phrase does not have an internal argument. This generalization would allow (105) and (106) to be acceptable while ruling out (102)-(104). It is completely unclear, however, why there should be such a restriction. Syntactic theory attempts to specify the constraints on how syntactic categories such as N, V, A, NP, VP and AP can be combined. In doing so, the internal structure of the category is not relevant other than features of the head, though there may be additional constraints on what the internal structure of a maximal projection may be. Furthermore, the restriction as stated above is stipulative. It would be nice if the restriction were due to some independently motivated principles.

The restriction also does not account for the acceptability of prenominal phrasal idioms like those in (107)-(109) in a natural way. The acceptability of the prenominal phrases in (107)-(109) depends on their being interpreted idiomatically; when the same phrases are interpreted compositionally the sentence that result are unacceptable (110)-(112).

(107) The under-the-table dealings of the mafia were known to many.
(108) The over-the-hill hockey players were easily injured.
(109) The off-the-wall behaviour of the linguist was appalling.

(110) * Despite the warning not to eat the food that had fallen under the table, John ate the under the table food.
(111) * To get to the bus stop you have to go over the hill. The over the hill bus stop is easy to find.
(112) * The calender fell off the wall, the off the wall calender is out-dated.

The prenominal phrases in (107)-(109) are similar in meaning to adjectives such as *illegal, old* and *crazy* and if one assumes that the semantics of an idiom determine its argument structure then the idioms would be acceptable because they would not have an internal argument. While this assumption about idioms would allow one to explain the acceptability of (107)-(109), the restriction on prenominal phrases remains stipulative and descriptive. Furthermore, there are also restrictions on what kinds of adjectives can appear in this position. I would like to argue that the restriction on which kinds of phrases and which adjectives can appear prenominally follows naturally from the semantics of dimensions.

5.3 A semantic constraint on prenominal modifiers.

While many different types of expressions can be used to predicate information about the entities named by the subject, the kinds of expressions that can appear in the position that modificational adjectives appear in is limited to expressions that name values on a dimension. The question that arises is why values on dimensions cannot be named by phrases that have heads with internal arguments? For example, why is it the case that *proud of his son* cannot be a value on a dimension? Those men who are proud of

their sons would be said to have that value on the dimension and different men could have the value to varying degrees.

The reason that phrases that have heads with internal arguments cannot name values on a dimension is that they name a property of an entity only in relation to an external entity. Because these phrases name a property of an entity only in relation to an external entity they allow for the possibility of a given entity having more than one value on the dimension named by the head. This is because the entity may have a different value on the dimension in question if different external entities (internal arguments) are considered. For example, the phrase *proud of his son* names a property of a man in relation to an external entity (his son). This allows the man to have more than one value on the dimension named by the adjective *proud*. Suppose that he is proud of his son but not of his brother, then the man would have two values on the dimension named by the adjective *proud*, one in relation to his son and one in relation to his brother. If the prenominal position is limited to expressions that name values on a dimension then it is clear that the phrase *proud of his son* cannot appear prenominally because it allows for the possibility that the noun being modified has more than one value on a dimension.

In the example given above, the phrases *proud of his son* and *proud of his brother* were said to potentially name different values on the dimension PROUD. One may argue, however, that neither phrase names a value on the dimension PROUD and that they name values on independent dimensions PROUD OF HIS SON and PROUD OF HIS BROTHER. If this were the case then the phrases would not violate the single value on a dimension constraint. The problem with such a proposal is that it misses the generalization that each of the phrases is naming the same property but with respect to different entities.

A related question is why *proud of his son* cannot name a value but *proud* can rather than the other way around? The answer to this question is that the former is dependent on the latter for its meaning. *Proud of his son* names a value on the same dimension as *proud* but only in relationship to an external entity. The dimension named by *proud* would exist even if the external entity didn't or had no relationship to the dimension whereas the putative dimension on which *proud of his son* names values can only exist if the external entity exists and is related to values on the dimension named by *proud*.

In sum, the reason why phrases whose heads have internal arguments cannot appear prenominally is because the structural position is limited to expressions that name values of the entity named by the noun. Phrases whose heads have internal arguments name properties only in relation to external entities (named by the internal argument) and thus allow for the possibility of having more than one value on the dimension named by the head and a given entity can only have one value on any dimension.

5.4 The relation between phrasal modifiers and relative adjectives.

Adjectives like *tall* name a property of an entity in relation to some standard and the standard often depends on entities other than the entity which is being spoken about. If interpreting adjectives like *tall* and phrases like *proud of his son* both depend on entities other than the entities being spoken about then why can *tall* appear prenominally but *proud of his son* cannot?

Relative adjectives name values on a dimension in relation to a standard. The standard itself is a value on the dimension and relative adjectives name values that are more or less than the standard value. The standard value is often determined by looking at other entities of the kind named by the noun and taking an average. While one may look at other entities to determine what the value of the standard should be, the adjective names a value on the dimension with respect to the standard value, not with respect to the entities that are used to determine what the standard value is. This differs from phrases with heads with internal arguments like *proud of his son*, in which the value named by *proud* is with respect to an external entity (his son) rather than with respect to another value on the dimension. Because relative adjectives name values on a dimension with respect to a standard value, an entity can only have one value on the dimension. By considering different sets of entities one might change the standard value but the entity being modified cannot have two values on the dimension for any particular standard value. The same is not true for phrasal modifiers, for any set of entities the entity being modified may have more than one value on a given dimension and thus modifiers that depend on external entities for their interpretation cannot name values on a dimension.

5.5 Restrictions on prenominal adjectives.

The constraint that restricts expressions in prenominal position to be expressions that name values on dimensions also explains why certain adjectives cannot appear prenominally. There are a few adjectives like *afraid*

that cannot appear prenominally. A list of the adjectives in Francis and Kucera (1982) that cannot appear prenominally[1] is given below in Table 5.1.[2]

Table 5.1 Adjectives that cannot appear prenominally

accustomed	acquainted	ablaze	alive
afoul	afraid	afire	asleep
devoid	aghast	aflame	awake
fond	ashamed	afloat	
laden	aware	aglitter	
mindful	awash	aglow	
sure	read	agleem	
unable	ready	aflutter	
		alone	

On the face of it, it seems unlikely that the constraint that prohibits these adjectives from appearing prenominally could be semantic. The list of adjectives contains many adjectives which have near synonyms which can appear prenominally. Examples of such synonym pairs include *afraid/scared, ablaze/blazing,* and *aghast/shocked.* While these adjectives are semantically very similar at an intuitive level, I will argue that they are semantically different in subtle ways and that these subtle differences in semantics are responsible for the different distributions of these adjectives.[3]

---

[1] Note that there are some adjectives on the list that can appear prenominally but when they do, they combine with the noun to pick out a kind of thing rather than modify the entity named by the noun. For example, a sure hand is not a hand that is sure but a steady hand and a ready wit is not a wit that is ready.

[2] I am grateful to Dana Sussman for her help in compiling this list.

[3] Inspection of the list shows that an extremely large percentage of the adjectives begin with *a*. One may be tempted to think that the restriction is phonologically based except that there are adjectives that cannot appear prenominally which do not begin with *a* and there are adjectives like *abrupt* and *absurd* that begin with *a* and can appear prenominally. Another tempting generalization is that adjectives whose initial morpheme is *a* cannot appear prenominally. This hypothesis is better supported than the phonological hypothesis because there aren't any adjectives in which the initial morpheme is *a* and can appear prenominally.

5.5.1 Adjectives with obligatory internal arguments.

Looking at the list of adjectives that cannot appear prenominally one finds that there are a small number of adjectives like *fond* and *unable* that obligatorily take an internal argument. These adjectives are shown in the first column of Table 5.1. Given that these adjectives cannot appear without an internal argument even when in predicate position it seems reasonable to assume that the reason they cannot appear prenominally is the same reason that phrases with heads that have internal arguments cannot appear prenominally. These adjectives allow an entity to have more than one value on a dimension depending on the external entity considered (named by the internal argument).

5.5.2 Adjectives with implicit obligatory internal arguments.

In addition to the adjectives that obligatorily require internal arguments, there is a set of adjectives that do not obligatorily require an

---

The main problem with such a restriction is that it is completely unclear how a child would learn such a generalization. The restriction is a language specific one that would apply to particular morphological forms and thus is unlikely to be part of a child's innate knowledge of language. This means that children would have to learn this generalization. If children are productive and extend adjectives they have heard in the predicate position to the prenominal position they would require negative evidence to realize that certain adjectives cannot appear in the prenominal position. There is now a great deal of evidence that children do not receive systematic feedback on the grammaticality of their utterances (Brown & Hanlon, 1970; Marcus, in press) and thus they could not rely on negative evidence to limit which adjectives can appear prenominally. Another undesirable property of the morphological restriction is that there is no explanation for why the restriction is what it is. Finally, even if there were some way in which the child could learn the morphological restriction, the child would require some other restriction to rule out the few adjectives that do not start with *a* but still cannot appear prenominally. Ideally one would like to have a single restriction account for all the adjectives.

internal argument when they appear in a predicate position. These adjectives are shown in the second column of Table 5.1. While these adjectives can appear without an overt internal argument when they are predicates (113), I will argue that semantically these adjectives require an internal argument. The meaning of an adjective such as *afraid* requires that there is something that makes the subject afraid. The same is not true of adjectives like *hungry*.

(113) The man is afraid (of x).

If adjectives like *afraid* implicitly require internal arguments, then the reason why they cannot appear prenominally is the same reason why phrases with heads with internal arguments cannot appear prenominally. However, *scared* and *frightened* are nearly synonymous with *afraid* and thus should require internal arguments also and should not be able to appear prenominally but they can appear prenominally (115), (117). This suggests that the explanation for why *afraid* cannot appear prenominally is incorrect but I will argue that the restriction due to internal arguments is correct but that there are two homophonous adjectives *scared* and *frightened* and that one of these is barred from appearing prenominally but the other one is not.

(114) The man is scared.
(115) The scared man ...
(116) The man is frightened.
(117) The frightened man ...

What evidence is there that there are homophonous forms of *scared* and *frightened* but only a single form of *afraid*? When these adjectives appear prenominally (115), (117), the entity named by the noun is necessarily

in a state of being *scared* or *frightened*. This is not the case when these adjectives appear in predicate position. When these adjectives appear in predicate position, then the entity named by the subject may be in a state of being *scared* or *frightened* but it may also be the case that the subject is not currently in a state of being *scared* or *frightened* but that it is a property of the subject that they are *scared* or *frightened* of something. This reading of the adjectives is similar to *afraid* in that it requires an internal argument (which may not be phonologically overt) (118)-(120) and this reading of the adjective cannot appear prenominally.

(118) The man is scared (of rabbits).
(119) The man is frightened (of rabbits).
(120) The man is afraid (of rabbits).

The reason for the difference between the adjectives *scared* and *frightened* and *afraid* becomes clear when one looks at the verbs from which they are derived. *Afraid* is derived from *fear* which has a single argument structure, while *scared* and *frightened* are derived form *scare* and *frighten* each of which have two argument structures. One of the argument structures of *scare* and *frighten* allows the derived adjectives to appear prenominally while the other one, which is similar to the argument structure of *fear*, prevents the derived adjectives form appearing prenominally. In one of the argument structures of *scare* and *frighten*, the subject has the thematic role AGENT and the direct object is the PATIENT. In this reading the PATIENT is affected and comes to have the state that is the result of the verb (e.g. the man in (121), (122) comes to be in a state of being scared or frightened). This is the argument structure from which the adjectives which can appear

prenominally are derived (115), (117); the noun being modified by the adjective has the state named by the adjective.

(121) Rabbits scared the man.
(122) Rabbits frightened the man.

In the second argument structure of *scare* and *frighten*, the subject is the THEME and the direct object is the PATIENT. Unlike the first argument structure, in this argument structure the PATIENT is not affected (e.g. the man in (121), (122) does not come to be in a state of being scared or frightened). It is from this reading that the adjectives *scared of* and *frightened of* are derived. These adjectives do not imply that the man is in a state of being scared or frightened. The adjective *afraid* is derived from the verb *to fear*, this verb has only one argument structure (123) which is similar to the second argument structure of *scare* and *frighten* of in that the PATIENT is not affected (e.g. the man in (123) is not in a state of fear).

(123) The man feared rabbits.

To summarize, *afraid* and one of the argument structures of *scared* and *frightened* name a property of the entity named by the noun (e.g. the man) only in relation to some other entity (e.g. rabbits). Furthermore, it is not necessary for the subject to be in the state named by the adjective. The other argument structure of *scared* and *frightened* names a property of the entity named by the noun. This realization of the adjectives can occur prenominally.

For the adjectives that cannot appear prenominally examined so far one can argue that the semantics of the adjectives require an internal argument and thus the adjectives cannot appear in a prenominal position cannot for the same reason that phrases with internal arguments cannot: they allow for the possibility of a given entity having more than one value on a given dimension because they name properties only in relation to external entities.

5.5.3 Adjectives that name properties in relation to an external entity.

A third group of adjectives that cannot appear in a prenominal position are those that are listed in the third column of Table 5.1. I argue that these adjectives do not name properties of the entity named by the noun *per se,* rather they name properties of the entity named by the noun only in relation to an external entity. These adjectives differ from those in the previous section in that they do not have internal arguments but they are dependent on external entities for their meaning.

These adjectives are similar to the adjectives discussed in the previous section in that they state something about the entity named by the noun only in relation to an external entity. The most obvious example of this is the adjective *alone*. It is clear that this adjective states something about the subject only in relation to other entities: there is no way to tell whether an entity is alone without considering its relation to other entities. For other adjectives in this group this semantic property is not as apparent at first glance. Many of the adjectives have close synonyms which can appear prenominally. For example, while the adjectives *aglow,* and *aglitter* cannot

appear prenominally *glowing*, and *glittering* can do so. I argue that the first set of adjectives state something that is true of the subject only in virtue of the subject's relationship to some other entity. This relationship between the subject entity and an external source is illustrated by the fact that *aglow*, and *aglitter* usually occur with a complement which names the source (124), (126). Furthermore, these adjectives cannot be used when the entity of which they are being predicated possess the property on their own (128). This contrasts with the closely related adjective *glowing* which can be used to describe the state of an entity independently of whether the state is a property of the entity or if the state is a property of the entity by virtue of a relationship to an external source (125), (129).

(124) His face was aglow (with pleasure).
(125) His face was glowing (with pleasure).
(126) The lake was aglitter in the moonlight.
(127) The lake was glittering in the moonlight.
(128) * The sun was aglow.
(129) The sun was glowing.

Another set of adjectives in this group includes adjectives like *ablaze*, *afire* and *aflame*. These adjectives also name a property of the subject entity which is true by virtue of a relationship to an external entity. These adjectives can only be use as resultatives and they have an implied external source (130). The implied external source is evident in their use with the verb *set* and the inability of these adjectives to appear with entities that possess these properties on their own (131). The related adjectives *blazing* and *flaming* cannot be used with the verb *set* (132) and can be used to describe entities that intrinsically possess these properties (133).

(130) John set the house ablaze/aflame/afire.
(131) * The fire was ablaze/aflame/afire.
(132) * John set the fire blazing/flaming.
(133) The fire was blazing/flaming out of control.

The difference between these resultatives and adjectives like *aglow* and *aglitter* is that the latter name properties of an entity that are true of the entity while it is in a particular relationship to the external source whereas resultatives name properties that the entity comes to have due to an external source. The continued existence of the property named by a resultative does not require a continued relationship between the subject entity and the source.

Finally, *asleep* and *awake* are similar to the resultatives in column three in that they are both resultatives but unlike the other resultatives the source is not external. *Alive* is the only adjective on the list that is not semantically similar to the other adjectives in any obvious way.

In sum, there is a semantic property that adjectives that cannot appear prenominally have in common: they name a property of the subject that is not intrinsic to the entity and is true only in relation to some external source. In some cases the external source has to be expressed overtly but in others the source is implicit. Because these adjectives name properties of entities in relation to external entities, they allow for the possibility that a given entity will have more than one value on a dimension depending on which external entities are considered. This prevents these adjectives from appearing prenominally.

5.6 Further evidence for the constraint on values on dimensions.

The proposal that modifiers that cannot appear prenominally are those that name properties in relation to external entities is supported by the pattern of acceptability exhibited by sentences like (134)-(141). Exactly those predicates that cannot appear prenominally are the predicates that cannot be predicated of *manner*. This makes sense because manner is observable behaviour. The behaviour is what it is and whether it is judged to be proud or not does not depend on any other entities such as one's son. It is perfectly possible to talk about John's manner towards his son (140) but once again the only question that arises is whether this behaviour can be called proud or not, independently of considering any other entities. (141) shows that the unacceptability of (137) and (139) is not due to a constraint against predicating phrasal modifiers of a genitive subject.[4]

(134) John is proud.
(135) John's manner is proud.
(136) John is afraid.
(137) * John's manner is afraid.
(138) John is proud of his son.
(139) * John's manner is proud of his son.
(140) John's manner towards his son is proud.
(141) John's sister is proud of his son.

Adjectives like *isolated, unread, much read* and *unique* provide counterevidence to the generalization that adjectives that depend on an external entity cannot appear prenominally. Why can these adjectives appear

[4] Pesetsky (1991) explains why phrases like *proud of his son* cannot be predicated of expressions like *John's manner* in terms of a morphological restriction. He explains a range of phenomena not discussed here with this restriction. It is unclear, however, how this explanation would account for why adjectives like *afraid* cannot be predicated of *manner* or why the expressions that cannot be predicated of *manner* are the same expressions that cannot appear prenominally.

prenominally? The reason is that each of these adjectives name a value with respect to the *whole* set of entities under consideration, rather than with respect to any particular entity or subset of entities under consideration. Because these adjectives name a property with respect to the whole set of entities under consideration, an entity cannot have more than one value on the relevant dimension. The entity being modified has a single value for the entities with respect to which the adjective is defined. For example, an unread book is unread by everyone being considered. Similarly a much read book is a book that has been read by many of the people being considered and thus the book cannot also not be much read.

The acceptability of these adjectives as prenominal modifiers suggests that the constraint on prenominal modifiers is that they cannot have a form that would allow the entity being modified to have more than one value on a given dimension. In general, properties that are true of the entity being modified with respect to external entities lead to the possibility of the entity having more than one value on the dimension in question. If, however, the property is defined by the whole set of entities under consideration, then it is no longer possible for the entity being modified to have more than one value on the relevant dimension.

5.8 Summary

The current chapter provided evidence for the claim that the structural position in which prenominal adjectives appear is limited to naming values on dimensions. Dimensions consist of mutually exclusive values and thus any phrases or adjectives whose meanings allow for the possibility of the entity being modified having more than one value is excluded from the

prenominal position. These included phrases that have heads with internal arguments and adjectives which name properties in relation to some external entities. It was also suggested that predication should be viewed as the general relation that states something about the subject. When the predicate consists of an adjective then the predicate names the value the subject is said to have on the given dimension.

Chapter 6

Semantics of Adjectives and Constraints on Word Learning

6.1 Introduction.

The previous chapters have provided psycholinguistic and linguistic data that suggest that adults represent adjectives as either naming values on dimensions or as helping pick out kinds and that adjectives can either predicate information about an entity or they can be used to talk about a set of entities on the basis of the value named by the adjective. The evidence presented thus far shows what the adult's semantic knowledge of adjectives is. In this chapter, I will show that the data from studies on the acquisition of adjectives suggest that children's semantic knowledge of adjectives is the same as that of adults. I will also show how knowledge of the semantics of adjectives can be used to motivate constraints on word learning. Finally, three of the constraints are tested in experiments with children and adults.

6.2 Evidence of children's knowledge of the semantics of adjectives.

Taylor and Gelman (1988) have shown that young two year olds (mean age 26 months) differentiate adjective and noun syntax in learning the meanings of novel words. In addition to differentiating the two types of words syntactically, children showed that they knew the semantic correlates of these syntactic categories. Taylor and Gelman presented children with two

instances of two kinds of familiar and unfamiliar objects. One of the objects in each kind had green fur and the other had a yellow and black plaid pattern on it. Children were taught either an adjective or a noun for one of the objects and then tested on their interpretation of the novel word. They found that children were more likely to have a category interpretation when a noun was used to name a familiar or unfamiliar object than when an adjective was used. Furthermore, children were more likely to make a property interpretation when the novel word was an adjective than when it was a noun. This result is predicted if children know that adjectives name properties of things but nouns do not. Golinkoff, Mennuti, Lengle, & Hermon (1992) find a similar pattern of results in two year olds interpretations of adjectives and nouns.

In the semantics developed here it is proposed that relative adjectives name a value on a dimension in relation to another (usually standard) value. Research by Carey (1978) and Gelman and Ebeling (1988) shows that two and three year olds use the words *big* and *little* in relation to a standard value and that they can use different types of information to determine the standard value. Gelman and Ebeling (1988) showed that two and three year olds use normative standards for deciding whether a given object is big or little; when they were asked whether an object was big or small, they gave a response that was based on knowledge of the range of sizes for that kind of object. In addition, children were able to change the standard they used for deciding whether an object should be called big or not. Children who called an object big in isolation called the same object little if it were presented in a context in which it was the smaller of two objects of that kind. Gelman and Ebeling also showed that three year olds can use information about intended use to determine the standard value on the dimension SIZE. Children who called a

hat little in isolation said that the same hat was big for a doll. These data suggest that two and three year olds, like adults, use words like *big* and *little* to name values in relation to a standard value.

Gelman and Markman (1985) propose that adjectives imply a contrast between members of a single noun category. Furthermore, they showed that 3-5 year olds can use this knowledge when learning novel words. In the experiment, they showed children four pictures of novel objects, three of which had the same shape and thus were meant to be seen as the same kind of thing and a fourth that differed in overall shape suggesting that it was a different kind of thing. Of the three objects of the same kind, two had the same marking and the third one had a different marking. Children were shown the pictures and asked to give the experimenter *the fep* or *the fep one.* They found that children who heard the novel noun were likely to give the object that differed in shape while those that heard the novel adjective picked the object with the distinctive marking. These results were interpreted as showing that 3-5 year olds know that adjectives imply a contrast within a noun category. This would follow if children know that adjectives name properties of things but that nouns do not. Markman and Gelman's (1985) result is supported by Waxman (1990) who found that adjectives help three year olds in classifying things at a subordinate level but nouns do not.

A prediction that follows from the semantics of adjectives is that, in general adjectives should be learned later than noun. This is because the default interpretation of prenominal adjectives is modificational not appositive and one cannot know which dimensions the kinds of entities named by the noun vary on without some experience with that kind. The prediction that adjectives are learned later than nouns is confirmed by cross-linguistic evidence (Gentner, 1982). Of course there are likely to be many

factors contributing to this difference in the order of acquisition but one factor that has independent motivation is the semantics of adjectives.

Bloom (1990a) investigated one and two year olds' knowledge of restrictions on modifiers in English. He looked to see whether young children know that adjectives can precede some kinds of nominals (nouns) but not others (proper names and pronouns which are proposed to be lexical NPs). Bloom found that while children produce both adjective-nominal and nominal-adjective sequences, the former are only produced when the nominal is a noun but nominal-adjective sequences are produced for nouns, proper names and pronouns. Bloom interprets this data as evidence that one and two year olds know the restriction in English that prevents prenominal adjectives from modifying NPs. While this data is consistent with a restriction based on the N/NP distinction, it is also consistent with a restriction that follows from the semantics of modification. As discussed in Chapter 3, the restriction based on the semantics of modification seems to be more explanatory in that it explains why there are certain restrictions on adjectives within NPs and when prenominal adjectives are interpreted as appositives. In any case, Bloom's data suggest that very young children's utterances are constrained in the way predicted by the semantics of modification. Note that for children to use a restriction on the basis of the semantics of modification they need to have represented adjectives as naming values on dimensions and so the data implicitly suggests that children represent adjectives as naming values on dimensions from the very begining.

Further evidence that young children's knowledge of the semantics of adjectives is the same as that of adults is that children as young as three order multiple prenominal adjectives in the same manner as adults (Bever &

Epstein, reported in Bever, 1970; Martin & Molfese, 1972). Martin and Molfese (1972) asked three and four year olds to repeat back sequences of two adjectives and a noun. Half of the sequences had the adjectives in the order preferred by adults and the other half had the adjectives in the non-preferred order. It was hypothesized that if children had the same restrictions on adjective orderings as adults then they would make more errors in repeating the sequences which had the adjectives in the non-preferrred order. They found that the three year olds made the same number of mistakes on the two kinds of sequences but that four year olds made more errors on the non-preferred orders than the preferred order. Martin and Molfese interpret this as evidence that children's sensitivity to the restrictions on adjective ordering develops between the ages of three and four. It turns out, however, that if one looks at the kinds of errors made by the three year olds there is evidence that even three year olds are sensitive to the order in which prenominal adjectives appear. Errors included many deletions of an adjective, deletions of the noun and reversals of the adjective order presented. The deletion errors may be due to processing limitations and do not necessarily bear upon the question of whether children know the constraints on the ordering of adjectives. The reversal errors on the other hand bear directly on this question. Ten out of the 13 reversal errors made by three year olds were on the non-preferred order; the proportion for four year olds was 15 out of 18. This data suggests that by the age of three children are sensitive to the constraints on adjective ordering. This conclusion was supported by the results of a production task in which three and four year olds were asked to describe pictures using two adjectives. Children ordered the adjectives as adults do though some pairs of adjectives were more likely to be ordered in this manner than others (Martin & Molfese, 1972). The data on adjective

ordering are consistent with the predictions made by the semantics of modification.

In summary, the data on the acquisition of adjectives suggests that young children have the same semantic knowledge of adjectives that adults have. In particular, they know that adjectives name properties of entities named by nouns; they know that adjectives imply a contrast (i.e. values on a dimension are mutually exclusive); adjectives cannot modify pronouns or proper names, a fact that follows from the semantics of modification; and they know the constraints on adjective ordering which also follow form the semantics of modification. In the rest of the chapter I will motivate certain heuristics that children may use to learn the meanings of words if they have the adult knowledge of the semantics of adjectives and modification.

6.3 Constraints due to knowledge of the semantics of adjectives.

Knowledge of the semantics of modification can be used to motivate constraints for the acquisition of adjective meanings. If a speaker hears an utterance in which a novel adjective is modifying a noun, then the speaker should be able to rule out those properties that do not vary for the kind named by the noun as potential meanings of the novel adjective. For example, if a speaker of English hears the utterance *that is a blicket table,* the speaker should be able to rule out the hypothesis that *blicket* means flat because *blicket* appears in a modificational position and all tables are flat and thus cannot be modified for flatness. Of course it is possible to interpret the prenominal adjective as an appositive but this is not the default

interpretation (see Chapter 3). Knowledge of the semantics of modification motivates the following heuristic for learning the meanings of adjectives.

(I) If a sentence with a prenominal adjective is used to talk about an entity of the kind named by the noun, then assume that the adjective names a value on a dimension on which that kind of entity has more than one value.

6.3.1 Study 1: Semantics of modification as a constraint on adjective meaning

Study 1 investigated whether adults and children as young as two-and-a-half years of age can use heuristic (I) to learn the meanings of novel adjectives. In the experiment subjects were presented with three instances of a novel kind. The three instances had the same value on two of the three dimensions (color, texture, marking) tested in the current study and differed in their values on the third dimension. After naming the kind of object, one of the three objects was named using a novel prenominal adjective. It was predicted that if children and adults can use their knowledge of the semantics of modification to constrain word meanings, then they should think that the novel adjective names the value on the dimension on which the objects varied.

Subjects.

Subjects were 30 undergraduate MIT students and 29 children between the ages of 2;4 and 3;10 with a mean age of 2;11 from daycare centers in the Boston-Cambridge area.

Materials.

The materials consisted of six different kinds of nonsense objects. There were three instances of each kind of nonsense object. Each kind of nonsense object was defined by a unique shape. Each of the three kinds of the objects in the teaching set had the same value on two dimensions (e.g. color and texture) and differed on the third (e.g. kind of marking). One of the objects in each of the teaching sets was designated the teaching object. The response set for a given teaching set contained three instances of a kind that was different from the teaching objects. Each of the objects in the response set matched the teaching object on one of the dimensions but differed on the other two dimensions of the teaching object. An example of a teaching set and the corresponding response set is shown in Figure 6.1. For the adults, the response set consisted of three different kinds of objects each matching the teaching object on only one dimension. The response set was changed for the children to avoid noise from choices based on a preference for one of the objects on the basis of it's object kind.

Procedure.

There were two conditions in the experiment, an adjective condition and a noun condition. The children were asked to play a game with Cookie Monster in which they were asked to help Cookie Monster pick up the things he asked for. They were then shown all the objects used in the study and were asked to help Cookie Monster put the objects into a box. This was done to give the children a chance to manipulate and play with the novel objects so that later choices would be less likely to be due to a desire to play with a novel object. After the objects were put away, Cookie Monster brought out one of

the triplets from the teaching set and taught the child a name for the kind of object in the following manner. *See these, these are blickets. This one is a blicket, and this one is a blicket and this one is a blicket. All of these are blickets, there are three blickets here.* The child was then asked to perform three actions with a blicket. This was to insure that the children thought that all three objects were blickets.

Adjective condition: In the adjective condition, Cookie Monster then pointed to one of the teaching objects and said: *You see this? This is a fep blicket. This is a fep blicket. This blicket is fep. These are all blickets and this blicket is a fep blicket.* The child was asked to give Cookie Monster the fep blicket out of the blickets and then the blickets were put to the side but remained in sight of the child. The child was then shown the response set and asked for *a fep thing, something which is fep.* After the child had made a choice, the child was asked to give Cookie Monster the fep blicket from the teaching set. This was to insure that children remembered to which object the novel adjective had been applied. Each of the subjects participated in three trials, each of the trials had a different dimension on which the teaching objects varied. The order in which children received these trials was counter-balanced across subjects.

Noun condition: A noun condition was included so that one could know whether the pattern of responses was due to children's knowledge of the semantics of adjectives or whether children would respond in a similar manner no matter what type of word was used. The noun condition was identical except that children were taught a second noun for the teaching

object rather than an adjective. Children were taught the noun in the following manner. *You see this? This blicket is a fep. This blicket is a fep. These are all blickets and this blicket is a fep. This blicket is a fep.* Children heard the noun the same number of times children in the adjective condition heard the adjective. Children were asked to give Cookie Monster *the fep* out of the response set.

Results.

The measure of interest is whether children in the adjective condition are more likely to pick the object in the response set that matches the teaching object on the dimension of variation than are children in the noun condition. Only the trials on which children successfully identified the teaching object at the end of the trial were analyzed. The percentage of responses in which children picked the object that matched the teaching object's value on the dimension of variation was 56% in the adjective condition and 33% in the noun condition. A one-way ANOVA was done with condition as the independent variable and the number of responses in which children picked the object that matched the teaching object's value on the dimension of variation as the dependent variable. The difference in conditions was significant $t(27) = 2.34$, $p<.01$ (one-tailed). The percentage of responses matching the teaching object's value on the dimension was greater than would be expected by chance $t(13) = 2.97$, $p< .01$. These result suggest that children can make use of the semantics of modification to decide which dimension is being named by a novel adjective. Unlike the children, the adults performed in a similar manner in the adjective and the noun condition. In both conditions, the adults picked the object that matched the

teaching object's value on the dimension of variation on practically all the trials (98% in the adjective condition and 96% in the noun condition). The responses on the noun trials indicate that they interpreted the novel noun as meaning glittery object or green object or striped object. This interpretation was confirmed by subject's report of what they thought the novel word meant. Unlike children, the adults seem to have been able to resort to a meta-linguistic strategy to pick the only response that could make sense, even though this response is unlike any real nouns. The adults reported consciously figuring out the meaning of the words in both the adjective and noun condition. The use of a meta-linguistic strategy makes it difficult to interpret the results as either supporting or rejecting the hypothesis that adjectives are more likely to name values on dimensions that have more than one value than dimensions that do not show any variation.

Discussion.

The results of this experiment show that children as young as two-and-a-half years of age can use use their knowledge of the semantics of adjectives to constrain their hypotheses about the meanings of a novel adjective. Children expect an adjective that modifies a noun to name a property that has more than one value for the kind of entity being spoken about. The study confirmed Taylor and Gelman's (1988) finding that two year olds differentiate adjective and noun syntax. The present study goes beyond the results of Taylor and Gelman (1988) and Gelman and Markman (1985) by showing that in addition to knowing that adjectives name properties and that adjectives name contrasts between members of a noun category two-and-a-half year olds

know that prenominal adjectives usually name values on dimensions which have more than one value for the kind named by the noun.

The results also show that there isn't a hierarchy for how nameable properties are. The nameability of a property is a function of the property's variation for the kind in question. Of course whether a property shows variation for a given kind cannot be the only factor that is relevant to the property's nameability. Most kinds of entities vary on more than one dimension and thus other constraints will be needed to choose among these dimensions. One possibility for how one might choose between dimensions that vary is discussed in Study 2 below. These results of the present study also help explain the difficulty children in Smith, Jones, & Landau (1992) had in mapping a novel adjective onto a property interpretation. In their study, children were taught a novel adjective on a single instance of an unfamiliar kind. In addition to interference form a strong object bias, children had no information about which dimensions vary and thus are likely to be named.

Finally, the study shed some light on children's expectations about the semantics of nouns. Children in the noun condition could have taken the noun to mean green object, or striped object, or glittery object but they did not. This is consistent with the fact that languages do not have nouns with meanings of this kind and suggests that the semantic knowledge of nouns that is responsible for this fact is present in children as young as two and a half years of age.

6.3.2 Study 2: Semantic constraints on choosing which dimension is being named by a novel adjective.

In most cases, objects of a given kind vary on more than one dimension. How does one decide which dimension is being named by a novel adjective? One possibility is motivated by the semantics of relative adjectives such as *tall*. As discussed in Chapter 2, these adjectives name a value in relation to another, often standard, value. This knowledge of the semantics of relative adjectives suggests that a novel adjective will be taken to name a value in relation to the standard value rather than the standard value itself. If this is the case, then one might expect that when a novel adjective is used with a particular entity, the adjective is more likely to name a value on a dimension on which the entity does not have a standard value, than a dimension on which the entity has a standard value. For example, if a given entity has the standard value on the dimension HEIGHT but has a value that differs from the standard on the dimension WIDTH, then one would expect that the novel adjective names a value on the dimension WIDTH.

The current experiment investigates whether values that are different from the standard or typical value are more likely to be named than values that are typical for the entities in question. The semantic analysis of adjectives clearly motivates such a heuristic for learning adjectives that name values on continuous dimensions because adjectives that name values on continuous dimensions do so only in relation to another, usually standard, value. However, the same heuristic may be used to learn adjectives that name values on discrete dimensions.

While discrete adjectives do not name values with respect to a standard value, knowledge of the standard value is relevant for whether a prenominal adjective is to be interpreted as being modificational or appositive. For example, the adjectives in phrases like *the spotted leopard* or *the feathered bird* are likely to be interpreted as appositives because leopards are usually spotted and most bird have feathers. Thus knowledge of the standard value on discrete dimensions is relevant for how prenominal adjectives are interpreted. Furthermore, there is evidence that suggests that we represent the standard or typical values that different kinds of things have on dimensions (Rips, 1989). Given that we represent standard values for discrete dimensions, it may be the case that a novel adjective would be more likely to name a value that is different from the standard for continuous as well as discrete dimensions.

The current experiment investigated whether values that are different from the standard are more likely to be named by a novel adjective than the standard value. Discrete dimensions were used because it is difficult to create the appropriate stimuli for relative dimensions for use with children as young as three years. Relative adjectives like *tall* require the presentation of multiple objects so that children can determine whether a given object is named by the adjective. Discrete adjectives like *spotted*, on the other hand, can be portrayed using a single object.

Subjects.

Subjects were 20 children between the ages of 3;1 and 3;9 with a mean age of 3;5 from daycare centers in the Boston-Cambridge area.

Materials.

Three sets of stimuli were created in the following manner. Each set consisted of two kinds of objects that three year olds are familiar with. There were three instances of each of the two kinds of objects. One of the instances, the teaching object, had a standard value on one dimension (D1) and a value that is different from the standard on another dimension (D2). The second instance of the kind had the standard value from the first instance on D1 but not the value that was different from the standard (D2). The third instance had the value that was different from the standard (D2) but not the value that was standard (D1). For example, a transparent spotted glass served as one of the teaching objects. The property of transparency is standard for a glass but being spotted is not. The second object in the set was a transparent glass that was not spotted and the third object was a spotted glass that was not transparent. The kinds of objects in a set were chosen such that the value that was considered standard for one kind of object was considered different from standard for the other kind of object. For example, the transparent spotted glass was paired with a transparent spotted dress. Unlike glasses, it is strange for a dress to be transparent but it is not strange for a dress to be spotted. The list of objects and properties used in the experiment are given in Table 6.1 below.

Table 6.1 List of objects and properties used in the experiment.

Teaching object	Same kind of object; matching standard value on teaching object	Same kind of object; matching non-standard value on teaching object	Different kind of object; matching standard value on teaching object	Different kind of object; matching non-standard value on teaching object
transparent spotted dress	opaque spotted dress	transparent plain dress	opaque spotted glass	transparent plain glass
transparent spotted glass	transparent plain glass	opaque spotted glass	transparent plain dress	opaque spotted dress
3-legged furry bear	4-legged furry bear	3-legged plain bear	4-legged furry stool	3-legged plain stool
3-legged furry stool	3-legged plain stool	4-legged furry stool	3-legged plain bear	4-legged furry bear
red bent leg	flesh tone bent leg	red straight leg	white bent fork	red straight fork
red bent fork	red straight fork	white bent fork	red straight leg	flesh tone bent leg

Procedure.

There were two conditions in the experiment, an adjective condition and a noun condition. Twelve children participated in the adjective condition and 8 participated in the noun condition. The children were asked to play a game with Cookie Monster in which they were asked to help Cookie Monster pick up the things he asked for. They were told that he was really picky and wanted them to give him only the objects that he had asked for. Children were then given two practice trials in which they were asked to give

Cookie Monster only those objects that were named by a familiar noun or adjective.

Adjective condition: Half the children were taught a novel adjective on one of the teaching objects in each set. Cookie Monster showed the child the teaching object and then described various properties of the object including the two properties of interest. The English words for the properties of interest were not used. The order in which the properties were spoken about was random. The following is a sample protocol.

Look at this. This is a glass. Let's look at it. It is round. it is small. You can drink out of it. It has these on it (pointing to spots). And you can see through it (holding glass up to child's eyes). And you know what? This is a fep glass. This glass is fep. This is a fep glass. I don't know what fep means but Cookie Monster says that this is a fep glass and wants you to help him find the fep things. Remember to give him only the things which are fep; the fep things.

The child was then shown a response set which contained (i) the teaching object, (ii) the other two objects of the same kind: one which had the same value as the standard but not the value that differed from the standard on the teaching object, and (iii) the two non-teaching objects from the matched kind: one which had the same value as the standard but not the value that differed from the standard on the teaching object. A sample teaching object and response set is shown in Figure 6.2. After the child made a choice, the experimenter would ask *are there anymore things which are fep, any more fep things or is that it?* This question was repeated after each choice.

Noun condition: The other half of the children were taught a novel noun for one of the teaching objects in each set. As in the first experiment, the

noun condition was included to insure that the responses children made were due to their knowledge of the semantics of adjectives rather than responses they would make for any kind of word. The procedure was identical to the adjective condition except that the experimenter taught the child a noun rather than an adjective: *You see this, this glass is a fep. This glass is a fep. This glass is a fep.* The experimenter then said, *I don't know what fep means but Cookie Monster says that this is a fep and wants you to help him find the feps. Remember to only give him the feps; give him the feps.*

Results.

The measure of interest in this experiment was whether children in the adjective condition would be more likely to think that a novel adjective names a value on a dimension on which an entity has a non-standard value than a dimension on which the entity has a standard value. The analyses of the data were performed in two ways. First, each trial was categorized as to whether the pattern of responses showed that: (i) the novel word was naming the property for which the teaching object had a non-standard value, (ii) the novel word was naming the property for which the teaching object had a standard value, or (iii) neither. A word was considered to be naming a property if the child picked only those items that had the relevant property and at least one of the items was not the teaching object. The trials were also categorized as to whether they fit a category interpretation for the same kind as the teaching object or for the other kind of object. Table 6.2 shows the percentage of trials that fell into these categories in the two conditions.

Table 6.2 Percentage of types of responses

Condition	Property matching standard value on teaching object	Property matching non-standard value on teaching object	Category matching the teaching object	Category not matching the teaching object
Adjective	.17	0	.31	.03
Noun	0	0	.42	0

As can be seen, the majority of children could not be classified as having named a value on either of the dimensions being considered or falling into any category. A one-way ANOVA was performed to see if children in the adjective condition were more likely to interpret an adjective as naming the property which was nonstandard on the teaching object than children in the noun condition. Similarly, a one-way ANOVA was performed to see if children in the noun condition were more likely to interpret a noun as naming the category of the teaching object than children in the adjective condition. Neither analysis yielded significant results but this may be due to the small number of subjects and the even smaller number of trials that could be interpreted as belonging to any of these categories. It seems likely that both effects would become significant with slightly larger sample sizes.

One problem with the analyses reported above is that there was a large tendency for children to give the experimenter all or many of the objects and it is possible that a number of the objects were given as part of a routine or a bias. For this reason, the next measure compared the number of times the

first non-teaching item picked had the standard value on the teaching object and the number of times it had the non-standard value in the two conditions. It was postulated that children's first choices were less likely to be due to a task bias and more likely to reflect children's intertpretation of the novel words. The first non-teaching object picked in the adjective condition had the value that was non-standard on the teaching object on 73% of trials in the adjective condition and 29% of trials in the noun condition. This difference was significant, t (17) = 2.56, $p < .035$ (one-tailed). The percentage of trials on which the first non-teaching object matched the nonstandard value of the teaching object in the adjective condition was marginally greater than would be expected by chance, $t(10)=2.304$, $p < .06$ (one-tailed).

Discussion.

The results of this experiment provide support for the hypothesis that children will take a novel adjective to name a value on a dimension that is different from the standard for the entity under consideration. Though there was a great deal of noise in the data, the pattern of results are consistent with the heuristic derived from the semantics of adjectives. A major source of noise in the present experiment was the bias children had for giving the experimenter all the objects. In addition to this task related source of noise, there were probably a number of other sources of noise. For one thing, it was assumed that children were familiar with the kinds of objects used in the experiment and that they knew what the standard values for the dimensions in question are. To the extent that children are unsure of this knowledge, they will be unable to make use of the proposed heuristic. Furthermore, children did not always attend to the syntactic frame in which the novel

words were taught. At least two children in the adjective condition spoke about the objects using the novel adjective as a count noun even though they had been taught an adjective and had been asked to repeat the novel word in an adjectival frame. Despite these multiple sources of noise three year olds were more likely to think that a novel adjective named a nonstandard value than a standard value.

The present experiment along with the previous experiment shows that three year olds can use their knowledge of the semantics of adjectives to constrain their hypotheses about the meanings of novel adjectives. Three year olds know that prenominal novel adjectives are more likely to name values on dimensions on which there is variation than dimensions on which there is no variation, Furthermore, those dimensions that have values different from the standard value are more likely to be named than dimensions on which the entity under consideration has the standard value. This second heuristic allows children to decide which of two dimensions that show variation is likely to be named by a novel adjective.

However, there are likely to be many situations in which two dimensions have values that are different from the standard. How does one know which dimension a novel adjective is likely to name in this kind of a situation? One possibility is that one can make use of a notion such as "distance from the standard" and then propose that dimensions on which the adjective is further from the standard are more likely to be named than dimensions on which the value is closer to the standard. While this may turn out to be true, it is unclear what aspect of the semantics of adjectives would motivate such a heuristic. The notion of "distance from the standard" does not play a role in the semantics developed here. One way in which the

notion could play a role in learning novel adjectives is in the application of the heuristic tested in the present experiment. It may be the case that the more different a given value is from the standard, the more likely it is that it is seen as differing from the standard. If this is the case, then the present heuristic may, in some cases, be able to help a learner choose between two dimensions that have values that differ from the standard.

The fact that children use this heuristic when learning novel adjectives shows that the notion of standard value is relevant for discrete adjectives as well as relative adjectives. It was proposed that the notion of a standard value on discrete dimensions is relevant for deciding whether an adjective should be interpreted as being modificational or as an appositive. The results of Study 2 suggest that conditions on interpretation in addition to manner of representation can serve as an important source of constraint in word learning.

6.3.3 Study 3: Constraints on noun meanings from the semantics of adjectives.

The previous two studies have investigated constraints on adjective meanings that are motivated by knowledge of the semantics of adjectives and modification. The current study investigates a heuristic that is motivated by the semantics of modification that constrains the meanings of nouns. The default interpretation of prenominal adjectives is modificational. In order for an adjective to be interpreted as modifying a noun, it has to be the case that the entities named by the noun have a value that is different than the one named by the prenominal adjective. If a learner has this knowledge of

modification then they would know that entities named by the noun vary on the dimension on which the prenominal adjective names a value. This prediction was tested in Study 3.

Subjects.

Subjects were 80 undergraduate MIT students.

Materials & Procedure.

Eight versions of a questionnaire were prepared in the following manner. On the first page of the questionnaire subjects were presented with a picture of a novel object. Half the questionnaires with a given novel object were labeled with a novel noun (e.g. *This is a blicket*). The other half were labeled with a familiar prenominal adjective and a novel noun (e.g. *This is a hairy blicket*). Two kinds of objects were used in the questionnaires. Each kind of object was paired with two adjectives. In all the versions, the next two pages contained the response set which included five distortions of the original object and another object with a completely different shape. One of the distortions had the property named by the adjective while the others had other values on the same dimension.[1] The object with the completely different shape also had the value named by the adjective. Subjects were asked to indicate whether each of the objects in the response set could be named by the novel noun. They were also asked to rate how certain they were of their response on a scale of 1-10.

---

[1] In the case of the item with the adjective *six-legged*, 3 of the distortions were six-legged and two weren't.

Results.

The variable of interest was whether subjects in the adjective condition would extend the novel noun to more instances than in the no adjective condition. A 2X4 ANOVA with condition (adjective/no adjective) and adjective as the independent variables and the number of items chosen in the response set as the dependent variable was conducted. The only significant effect was of condition $F(1,72) = 7.76$, $p < .007$. Subjects picked more items in the adjective condition than in the no adjective condition. No difference was found in the ratings of the items that were picked as being named by the noun in the two conditions. This suggests that the presence of the adjective affected whether subjects thought that a given entity was named by the novel noun but not the degree to which they are certain that a given entity is named by the novel noun.

These results show that subjects were willing to extend a novel noun to a greater number of entities if they had seen the noun modified by a familiar adjective than if they hadn't. While this is consistent with the predictions made by the semantics of modification, the predictions of the semantics of modification are more specific. The prediction is that subjects should be more likely to name items of the kind named by the noun which have other values on the dimension named by the adjective if they have heard the noun modified by the adjective. In order to test this more specific prediction a parallel set of analyses was performed without the items in the response set that matched the teaching item on the value of the teaching item and without the distractor. Excluding these items insured that a

difference between the two conditions would be due to items with other values on the relevant dimensions being chosen rather than a tendency to pick the items that match the value named by the adjective in the adjective condition. The results were identical to the previous analyses; subjects picked more items in the adjective condition than the no adjective condition, $F(1,72) = 10.879$, $p < .002$; and no difference in the ratings. The percentage of novel instances to which the novel noun was extended for the four adjectives in each condition is given below in Table 6.3.

Table 6.3 Percentage of novel instances to which novel noun was extended.

	Adjective condition	No Adjective condition
spotted	62.5	47.5
striped	65.0	55.0
hairy	90.0	40.0
six-legged	56.7	40.0

Discussion.

These results confirm the hypothesis that subjects are likely to extend a novel noun to items with different values on a given dimension if they hear the noun used with a prenominal adjective that names values on that dimension. These results extend those of Landau, Smith, & Jones (1988) and Jones, Smith, & Landau (1991) who have found that adults and two and three year olds extend count nouns used to name unfamiliar objects on the basis of shape and that they extend the novel nouns on the basis of shape and texture when the object has eyes, suggesting animacy. They propose that children can use information about the co-occurrance of different types of features (e.g. eyes and texture) to make certain features of objects more relevant (e.g. texture) to the generalization of count nouns for that object. The present experiment provides evidence for a linguistic source of information that can lead the learner to generalize a count noun for an object on the basis of a property other than shape. The advantage of the heuristic proposed here is that it does not require the learner to have nonlinguistic knowledge of the co-

occurrance of different types of features. The learner only needs to know the meaning of the adjective and the semantics of modification. Future research will determine whether young children can make use of this source of information when learning the meanings of novel nouns.

6.4 Conclusions from the acquisition data.

The results of the experiments reported here show that knowledge of the semantics of adjectives can be used to learn the meanings of words. In addition to providing information about how meanings of adjectives may be learned. The present experiments present acquisition data that supports the analysis of the semantics of adjectives proposed in this thesis. Finally, given that children are able to exploit knowledge of the semantics of adjectives to learn the meanings of novel adjectives, these data suggest that at least some links between grammatical categories and nonlinguistic cognition are bi-directional.

6.5 Questions raised by the present analysis.

In this section, I discuss a number of questions that are raised by the semantic analysis of adjectives that is proposed in this thesis. I will also discuss possible solutions to these questions that will have to be investigated in future research.

The analysis of adjectives proposes that there are two types of adjectives, those that name values on dimensions and those that combine with nouns to name certain kinds of things. An obvious question that arises is how does the learner know to which of these two kinds of adjectives a

novel adjective belongs? In addition to raising this question, the analysis suggests some answers. First, because adjectives that help name kinds of things do not name properties of things, if one hears an adjective in a predicate one can conclude that the adjective names a value on a dimension. A second way in which one could learn that a novel adjective names a value on a dimension is if one hears a pair of sentences like (143) and (144). It is clear from these sentences that if *zav* is an adjective then it has to be an adjective that names a value on a dimension. This is because one would know from (2) that *fep* names a value on a dimension and adjectives that help name a kind can only modify a noun or an adjective noun combination in which the adjective is an adjective that helps name a kind. Third, if an adjective appears with an intensifier or in a comparative form, then the adjective has to name a value on a dimension because adjectives that help name kinds have to be discrete.

(143) The blicket is fep.
(144) The zav fep blicket is expensive.

While the heuristics mentioned above provide a way of deciding which adjectives name values on dimensions, they do not provide a way of determining which adjectives help name a kind. One possibility is that the default meaning of an adjective is that they help name kinds of things and only if one gets information of the kind mentioned above do we decide that the adjective names a value on a dimension. There are number of problems with this proposal. First, though we do not have the relevant data it seems likely that we do not have to hear an adjective in predicate position or with an intensifier or as a comparative in order to know that it names a value on a dimension. Furthermore, it is not clear how an adjective would ever come to

be used with a meaning other than one which helps pick out a kind. The point is that ultimately the difference between the two kinds of adjectives is a semantic difference and while syntactic and morphological cues may help one decide whether particular adjectives belong to one or the other semantic class, the syntactic and morphological cues are dependent on a semantic distinction. The question then becomes, how does the child initially know that a given word names a dimension or helps pick out a kind?

The heuristics presented in this thesis assume that the learner conceptualizes the entities in question as having values on various dimensions. It is not obvious, however, what should be conceptualized as a dimension. Notions that are conceptualized in English as dimensions include readily perceivable physical attributes like size shape, color and texture as well as mental attributes like happiness, nervousness, and jealousy; internal states like hungriness and queasiness; and abstract properties like legality, appropriateness and absurdity. Furthermore, not all languages lexicalize these properties as adjectives (Dixon, 1982; Gentner, 1982) and within a language one can often talk about the same situation in more than one way. For example, In English, we can convey the notion that an apple is red by saying something like the apple has redness. This leads to a very difficult learning situation unless there are certain notions that are always lexicalized as adjectives. If this is the case then one can learn the syntactic properties of adjectives with a given meaning by attending to how these canonical dimensions are named and then learn to extend this conceptualization to other domains through the use of language. It turns that there is a small set of notions that are always lexicalized as adjectives (Dixon, 1982). In a cross-linguistic survey of seventeen languages, Dixon (1982) found that words that name values on the dimensions of AGE (e.g. young/old),

PHYSICAL DIMENSION (e.g. big/small), VALUE (e.g. good/bad) and COLOR are always lexicalized as adjectives; even in languages with very small adjective classes. This suggests that children can learn the syntactic and morphological properties of adjectives that name dimensions when they learn names for these dimensions and can then rely on linguistic information to identify other notions that are conceptualized as dimensions in their speech community.

This proposal is similar to semantic bootstrapping theories (Grimshaw, 1981; Pinker, 1984) in that it makes use of canonical mappings between semantics and syntax to learn aspects of syntax but it differs from these theories in that it proposes that the mappings are bi-directional and thus the syntax can be used to learn aspects of the word's meaning. The current proposal is an example of the semantic competence hypothesis proposed by Bloom (1990).

6.6 Conclusions.

This thesis explored the possibility of deriving constraints on the acquisition of adjective meanings from knowledge of the semantics of adjectives. The assumption that underlies such a project is that there are at least some bi-directional links between grammar and cognition. This assumption was supported by the results of the word learning experiments, which showed that one can use knowledge of how grammatical categories are related to nonlinguistic cognition to help learn the meanings of novel adjective. The links themselves are not sufficient to allow the learner to infer the meanings of the novel word but when knowledge of the links is combined with nonlinguistic knowledge about the entity being considered,

the semantics of adjectives provide a powerful constraint on the meaning of the novel adjective. In conclusion, the experiments in the current thesis add to a large body of data suggesting that children's knowledge of language is the same as that of adults. If this is so, then looking at adult knowledge of language is likely to provide and extremely important and useful source of constraint on the kinds of mechanisms that are used to read speakers' minds.

References

Annear, S. S. (1964). The ordering of pre-nominal modifiers in English. *Project on linguistic analysis report*, **8**, Ohio State University Research Foundation.

Au, T. K., and Markman, E. M. (1987). Acquiring word meanings via linguistic contrast. *Cognitive Development*, **2**, 217-236.

Beckman, M., and Pierrehumbert, J. (1986). Intonational structure in Japanese and English. *Phonology Textbook*, **3**, 255-309.

Bever, T.G. (1970). The cognitive basis for linguistic structures. In J.R. Hayes (Ed.) *Cognition and the development of language.* New York: Wiley.

Bloom, P. (1990a). Syntactic distinctions in child language. *Journal of Child Language*, **17**, 343-355.

Bloom, P. (1990b). Semantic stucture and language development. Unpublished doctoral dissertation, MIT.

Bloom, P. (1992). A theory of constraints on word meaning. Unpublished manuscript, University of Arizona, Tuscon, AZ.

Bolinger, D. (1967). Adjectives in English: attribution and predication. *Lingua*, **18**, 1-34.

Brown, R. (1957). Linguistic determinism and the part of speech. *Journal of Abnormal and Social Psychology*, **55**, 1-5.

Brown, R., and Hanlon, C. (1970). Derivational complexity and order of acquisition in child speech. In J. R. Hayes, (Ed.), Cognition and the development of language. New York: Wiley.

Carey, S. (1978). The child as word learner. IN M. Halle, J. Bresnan, and G. Miller (Eds.), *Linguistic theory and psychological reality.* Cambridge, MA: MIT Press.

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambride, MA: MIT Press.

Chomsky, N. (1986). *Knowledge of language:Its nature, origin, and use.* New York: Prager.

Clark, E. V. (1987). The principle of contrast: A constraint on language acquisition. In B. MacWhinney, (Ed.), *Mechanisms of language acquisition*. Hollsdale, NJ: Earlbaum.

Crain, S. (1991). Language acquisition in the absence of experience. *Brain and Behavioral Siences*, **14**, 597-650.

Danks, J. H., and Glucksberg, S. (1971). Psychological scaling of adjective orders. *Journal of Verbal Learning and Verbal Behavior*, **10**, 63-67.

Dickinson, D. K. (1988). Learning names for materials: Factors constraining and limiting hypotheses about word meaning. *Cognitive Development*, **3**, 15-35.

Dixon, R. M. W. (1982). Where have all the adjectives gone? in R. M. W. Dixon (Ed.) *Where have all the adjectives gone? and other essays in semantics and syntax*. New York: Mouton.

Fisher, C., Gleitman, H., and Gleitman, L. R. (1991). On the semantic content of subcategorization frames, *Cognitive Psychology*, **23**, 331-392.

Fisher, C., Hall, D. G., Rakowitz, S., & Gleitman, L. R. (1991). When is it better to receive than give: Syntactic and conceptual constraints on vocabulary growth. Unpublished manuscript, Department of Psychology, University of Illinois, Urbana.

Francis, W. N., and Kucera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar*. Boston: Houghton Mifflin.

Fukui, N., and Speas, M. (1986). Specifiers and projections. In N. Fukui, Rapaport, T., and Sagey, E. (Eds.) *MIT Working Papers in Linguistics*, **8**, Cambridge.

Gathercole, V. C. The contrastive hypothesis for the acquisition of word meaning: a reconsideration of the theory. Journal of Child Language, 14, 493-531.

Gelman, S. A., and Ebeling, K. S. (1988). Nonegocentric uses of *big* and *little* by preschool children. *Papers and Reports on Child Language Development*, **27**, 47-54.

Gelman, S. A., and Markman, E. M. (1985). Implicit contrast in adjectives vs nouns: Implications for word learning in pre-schoolers. *Journal of Child Language*, **12**, 125-143.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. A. Kuczaj II (Ed.), *Language development, Vol II: Language, thought, and culture.* Hillsdale, NJ: Earlbaum.

Gleitman, L.R. (1990). The structural sources of word meaning. *Language Acquisition*, **1**, 3-55.

Golinkoff, R. M., Mennuti, T., Lengle, C., and Hermon, G. (1992). Is *glorpy* a noun or an adjective?: Identifying the part of speech of a novel word. Paper presented at the International Conference on Infant Studies, Miami, Florida, May, 1992.

Goodman, N. (1983). *Fact, Fiction, and Forecast.* Cambridge, MA: Harvard University Press.

Grice, H. P. (1975). Logic and Conversation. In P. Cole, and Morgan, J., (Eds.), *Syntax and semantics 3: Speech acts.* New York: Academic Press.

Grimshaw, J. (1981). Form, function, and the language acquisition device. In C. L. Baker and J. McCarthy, (Eds.), *The logical problem of language acquisition.* Cambridge, MA: MIT Press.

Gross, D., Fischer, U., and Miller, G. A. (1989). The organization of adjectival meanings. *Journal of Memory and Language*, **28**, 92-106.

Hall, D.G. (1991). Acquiring proper names for familiar and unfamiliar animate objects: Two-year-olds' word-learning biases. *Child Development*, **62**, 1142-1154.

Hall, D. G., and Waxman, S. R. (in press). The acquisition of unrestricted and restricted count nouns. *Child Development.*

Hyams, N. (1986). *Language acquisition and the theory of parameters.* Dordrechht: Reidel.

Jones, S., Smith, L., and Landau, B. (1991). Object properties and knowledge in early lexical learning. *Child Development*, **62**, 499-516.

Katz, N., Baker, E., and Macnamara, J. (1974). What's in a name? A study of how children learn common and proper names. *Child Development*, **45**, 469-473.

Landau, B., Smith, L., and Jones, S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, **3**, 299-321.

Levi, J. N. (1978). *The syntax and semantics of complex nominals.* New York: Academic Press.

Levin, B., and Pinker, S. (1992). Introduction to special issue of Cognition on lexical and conceptual semantics. *Cognition,* **41,** 1-7.

Marcus, G. (in press). Negative evidence in language acquisition. *Cognition.*

Markman, E. M., and Hutchinson, J. E., (1984). Children's sensitivity to constraints on word meaning: Taxonomic vs. thematic relations. *Cognitive Psychology,* **16,** 1-27.

Markman, E. M., and Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology,* **20,** 121-157.

Markman, E. M. (1989). *Categorization and naming in children.* Cambridge, MA: MIT Press.

Martin, J. E. (1969). Semantic determinants of preferred adjective order. *Journal of Verbal Learning and Verbal Behaviour,* **8,** 697-704.

Martin, J. E., and Molfese, D. (1972). Preferred adjective ordering in very young children. *Journal of Verbal Learning and Verbal Behavior,* **11,** 287-292.

Miller, G. A., and Johnson-Laird, P. N. (1976). *Language and perception.* Cambridge, MA: Harvard University Press.

Naigles, N. (1990). Children use syntax to learn verb meanings. *Journal of Child Language,* **17,** 357-374.

Pesetsky, D. (1991). Linking and selection: The role of the zero morpheme. Unpublished manuscript, Department of Linguistics and Philosophy, MIT, Cambridge, MA.

Pinker, S. (1984). *Language learnability and language development.* Cambridge MA: Harvard University Press.

Pinker, S. (1992). How could a child use verb syntax to learn verb semantics?

Quine, W. V. O. (1960). Word and object. Cambridge, MA: MIT Press.

Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Voisiniadou and A. Ortony, (Eds.), *Similarity, analogy, and thought.* New York: Cambridge University Press.

Siegal, M. E. A. (1980). *Capturing the adjective.* New York: Garland.

Smith, L., Jones, S., and Landau, B. (1992). Count nouns, adjectives, and perceptual properties in children's novel word interpretations. *Developmental Psychology,* **28,** 273-286.

Soja, N. N. (1987). Ontological constraints on 2-year-olds' induction of word meanings. Unpublished doctoral dissertation, MIT.

Soja, N. N. (1992). Inferences about the meanings of nouns: The relationship between perception and syntax. *Cognitive Development,* **7,** 29-46.

Soja, N. N., Carey, S., and Spelke, E.S. (1991). Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition,* **38,** 179-211.

Sproat, R., and Shih, C. Prenominal adjectival ordering in English and Mandarin. *Procedings of the North-eastern Linguistics Society,* **19,** 465-489.

Taylor, M., and Gelman, S. A. (1988). Adjectives and nouns: Children's strategies for learning new words. *Child Development,* **59,** 411-419.

Vendler, Z. (1968). *Adjectives and Nominalizations.* The Hague: Mouton.

Waxman, S. R., and Gelman, R. (1986). Preschoolers' use of superordinate relations in classification and language. *Cognitive Development,* **1,** 139-156.

Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual heirarchies: Evidence form preschool children. *Cognitive Development,* **5,** 123-150.

Wexler, K., and Culicover, P. (1980). *Formal principles of language acquisition.* Cambridge, MA: MIT Press.

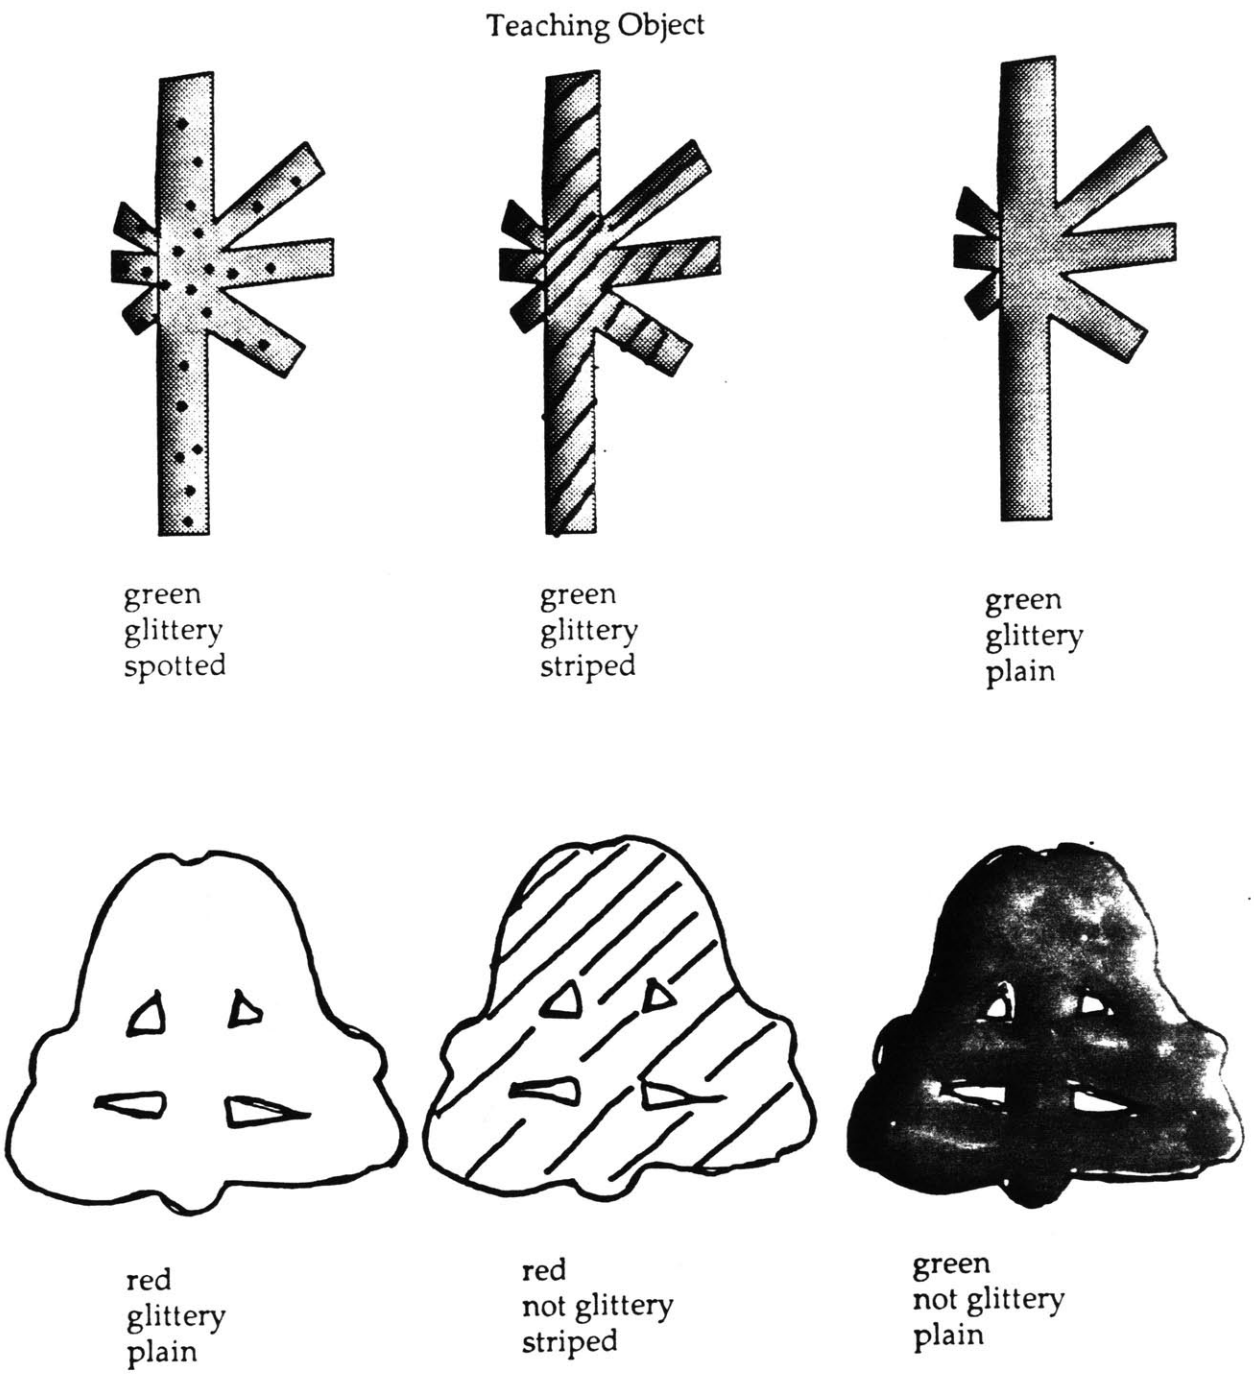

Figure 6.1. Sample stimuli for Study 1.

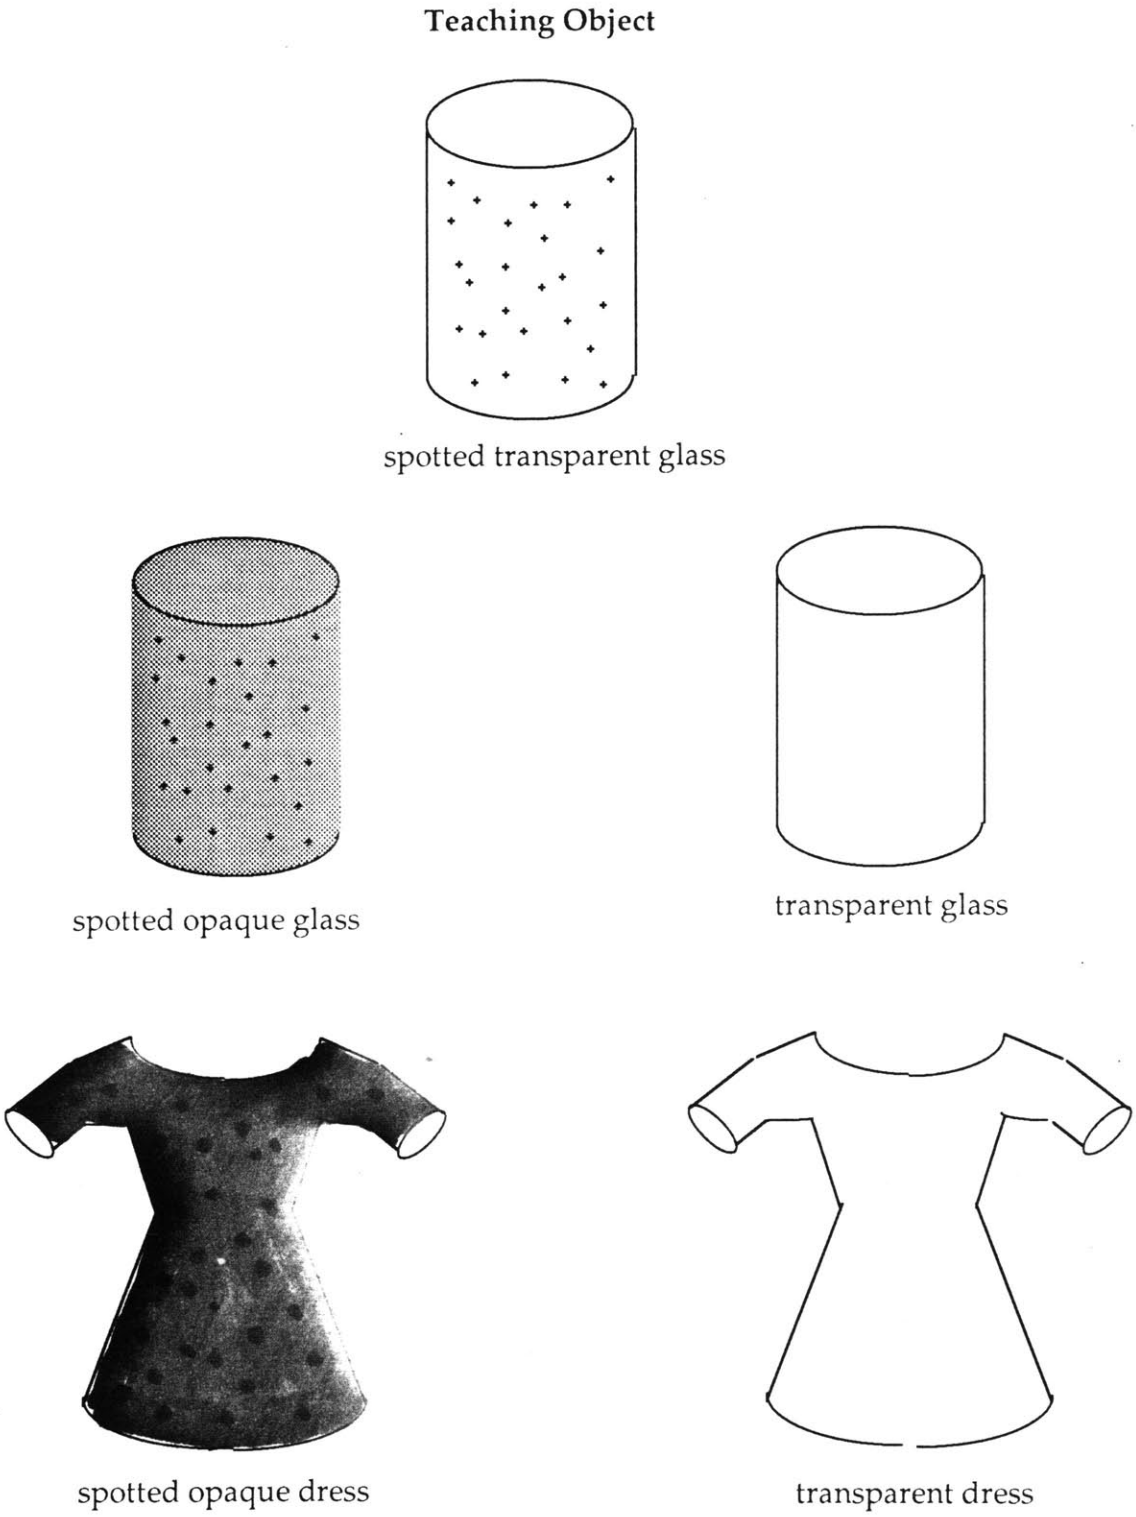

Figure 6.2 Sample teaching object and response set.
All the items were used in the response set.