

CHAPTER 1

APPROACHES TO ANTONYMY

1.1 Introduction

This dissertation is an investigation of *antonymy*, the lexical-semantic relation¹ which unites two words which have "opposite" or "contrasting" meanings. Along with other lexical-semantic relations such as synonymy, hyponymy, and the part-whole relation, antonymy has often been taken as a primitive by linguists, psychologists, and lexicographers. These relations all have a basic quality to them in that they are found in all human languages and they seem to be characterized more easily by examples than by explanation; so, for example, if we want to teach someone what the word *antonym* means, the easiest way would be to give him or her some examples of antonyms: *hot/cold, wet/dry, old/new, active/passive, good/bad, clean/dirty* and so on. These pairs are all antonyms in the technical sense of the word as used by most linguists--they are all gradable adjectives which have opposite meanings--but in its wider sense, as it is used by some lexicographers and most lay people, the set of antonyms (or *opposites*) includes other kinds of adjectives, as well as nouns, verbs, adverbs, and prepositions, e.g., *male/female, true/false, friend/enemy, front/back, beginning/end, love/hate, pass/fail, quickly/slowly, in/out, and up/down*.

Antonymy has several characteristics which set it apart from the other lexical-semantic relations. First, as Cruse says, "Of all the relations of sense that semanticists propose, that of oppositeness is probably the more readily

¹I am following Evens, et al. (1980) in using the term *lexical-semantic relation* as a kind of neutral term, rather than either *lexical relation* or *semantic relation*, terms which have theoretical connotations for many people.

apprehended by ordinary speakers." (Cruse 1986, 197) In fact, as he points out, most languages have a non-technical term for antonyms, e.g., English *opposite*, German *gegensatz*, and Japanese *hantaigo*.

A second interesting fact about antonymy is that most (probably all) languages have morphological processes which can create antonyms. As Lyons writes:

In many languages, including English, the most commonly used opposites tend to be morphologically unrelated (*good:bad, high:low, beautiful:ugly, big:small, old:young*). But these are outnumbered in the vocabulary by such morphologically related pairs as *married:unmarried, friendly:unfriendly, formal:informal, legitimate:illegitimate*, etc. (Lyons 1977, 275)

In addition to *un-* and *in-*, English also has the prefixes *dis-* (*like/dislike, honest/dishonest*) and *de-* (*colonize/decolonize, emphasize/deemphasize*), as well as the suffixes *-less* and *-ful*, which together sometimes form pairs of antonyms (*harmless/harmful*).² However, English does not have derivational processes involving other kinds of lexical-semantic relations; that is, there are no affixes which create synonyms or hyponyms or form the name of a part from a whole.

Third, there seems to be something "mysterious" or even "mystical" about antonymy because opposites seem as different as they can possibly be, yet they still have something in common. In some measures of semantic similarity, antonyms pattern much like synonyms. For example, in word association tests, people often respond to a word with its opposite or with a synonym (see Deese 1965), yet people clearly feel that antonyms and synonyms are not the same.

²Of course, these affixes do not always create antonyms; it is easy to find examples which, due to semantic drift or some other cause, look morphologically as if they should be opposites but which actually are not, e.g., *integrate/disintegrate* and *easy/uneasy*.

Cruse describes this mysterious quality of antonyms in his book *Lexical Semantics* (1986). He writes:

Opposites possess a unique fascination, and exhibit properties which may appear paradoxical. Take, for instance, the simultaneous closeness, and distance from one another of opposites. The meanings of a pair of opposites are felt to be maximally separated. Indeed, there is a widespread idea that the power of uniting or reconciling opposites is a magical one, an attribute of the Deity... The closeness of opposites, on the other hand, manifests itself, for instance, in the fact that members of a pair have almost identical distributions, that is to say, very similar possibilities of normal and abnormal occurrence... Philosophers and others from Heraclitus to Jung have noted the tendency of things to slip into their opposites states; and many have remarked on the thin line between love and hate, genius and madness, etc. (Cruse 1986, 197)

Cruse says that this feeling of simultaneous closeness and distance can be at least partially explained by semantic dimensions--opposites are two words which share a semantic dimension, e.g., the dimension of TEMPERATURE in the case of *hot* and *cold*, but which occupy opposite poles along this dimension.

Another interesting thing about antonymy is that language learners expect that there will be pairs of antonyms available to label certain types of concepts. Miller and Fellbaum (1991) remark that students learning a second language "when given only one member of an antonymous or opposed verb pair, will insist upon being taught the other member." And Kagan says:

As the child creates categories, she is disposed to invent their complement. Soon after learning the meaning of *up*, the child learns the meaning of *down*; after learning the meaning of *high*, she learns *low*; after *good*, she develops the meaning of *bad*. (Kagan 1984, 187)

Kagan claims that children do not have to be explicitly taught that there are such things as opposites; instead they seem to expect to find them. This kind of expectation may exist for other kinds for lexical-semantic relations too, but it

seems to be especially strong in the case of antonymy.³ Even adult speakers feel the usefulness of antonyms when expanding the vocabulary of their native language. As Egan puts it:

It is good, we feel, to know the exact antonym of a word, for not only will it give us a firmer grasp of the meaning of the word to which it is opposed, but inversely, of itself. (Egan 1968, 28a)

Anthropologists (e.g., Casagrande and Hale 1967) have found that antonymy is one of the semantic relations commonly used in folk definitions; a thing or quality can often be defined in terms of what it is not, for example, *big* can be defined as 'not small.'

Native speakers have strong intuitions about which pairs of words are good examples of antonyms; for example, probably all English speakers would agree that *cold* is the opposite of *hot*, that *happy* is the opposite of *sad*, and that *down* is the opposite of *up*. This intuition has been called the "clang phenomenon." It occurs strongly with a rather small set of prototypical antonyms, but it is not so strong for less prototypical opposites; for example, some, but not all speakers would agree that *vigorous* and *feeble* or *loud* and *faint* are opposites.⁴

The research in this dissertation was prompted by a desire to learn more about the properties of antonyms in order to explain the strong intuitions that people have about them. I began with the initial questions listed in (1) below.

(1) 1. What makes two words antonyms? What exactly are the

³In the case of synonymy, in contrast, research suggests that children do not expect to find synonyms. In order to explain children's early word learning, a "mutual exclusivity" principle has been proposed (see Markman 1994, for example). According to this principle, when young children hear a new word, they assume that it is distinct in meaning from the words they already know.

⁴Mettinger (1994), for one, considers both of these pairs to be examples of opposites.

semantic dimensions which antonyms are said to share?

2. Why do some words have antonyms while others have none?

Why do some words have more than one antonym, e.g.,

good/bad and *good/evil* or *happy/sad* and *happy/unhappy*?⁵

3. What accounts for native speakers' strong intuitions about which words are antonyms (the "clang phenomenon")? What is the difference between pairs of words which are good, prototypical examples of antonyms, such as *hot* and *cold*, *wet* and *dry*, and pairs which are seem to contrast in meaning but which many people would not consider antonyms, e.g., *hot* and *chilly* or *loud* and *faint*? Is there a model of antonymy that can explain why some pairs are "good" antonyms and some are not?

There has been a great deal of research on antonymy in fields ranging from anthropology to lexicography to computational linguistics, but much of this research focuses on a specific aspect of antonymy (i.e., the phenomenon of markedness, described briefly in section 1.2.1 below) rather than on the basic questions listed above. The rest of this section, therefore, reviews only the literature which focuses on the kinds of larger issues I am most interested in.

1.2 Linguistic approaches to defining antonymy

One of the main concerns of linguists studying antonymy has been to try to determine the boundaries of antonymy. In its widest sense, as it is used by

⁵I am not so interested here in cases in which it is clear that one of the words has two very different senses, as in the case of *right*, which has two antonyms, *left* and *wrong*. The sense of *right* which is paired with *left* is quite distinct in meaning from the one which is paired with *wrong* and presumably this explains why *right* has two antonyms.

lay people and in some dictionaries and thesauri, the term *antonym* is the same as *opposite*, and it has been applied to many types of semantic opposition, from prototypical antonyms such as *hot/cold* and *large/small* to pairs of words which show only a vague or limited type of semantic contrast, e.g., *parent/child* or *frank/hypocritical*. Even pairs such as *coffee/tea* or *gas/electric* could be considered opposites in some contexts.

Most linguists, however, have found this use of *antonym* too broad to be useful or interesting, and so many linguists, including Palmer (1976), Lyons (1977), Lehrer and Lehrer (1982), and Cruse (1986), have tried to categorize opposites into different types. They then define the term *antonym* in a more narrow sense, using it only to refer to the set of *gradable opposites*, a set which has many interesting properties such as markedness and committedness.

Although the adjectives I focus on in the case studies all happen to be antonyms in the narrow sense, my model of antonymy is relevant to a wider range of opposites as well, so the next two sections describe some of the different types of opposites which have been discussed in the literature.

Section 1.2.1 focuses on antonyms in the narrow sense and section 1.2.2 focuses on other types of opposites (antonyms in the wider sense). Then, since understanding antonymy also necessitates understanding why some pairs of words are not antonyms, section 1.2.3 focuses on pairs that are almost-but-not-quite opposites.

1.2.1 Antonyms as gradable opposites

The set of gradable opposites includes many common and prototypical pairs of opposites including *big/little*, *good/bad*, *high/low*, *hot/cold*, *happy/sad*,

tall/short, and *wet/dry*. The defining property of this set is that the opposites name qualities which are gradable, that is, qualities which can be conceived of as 'more or less'; therefore the scale (dimension) with which each pair is associated has a neutral mid interval. Take for example *hot* and *cold*, which describe opposite ends of the scale of TEMPERATURE. *Hot* and *cold* are both gradable; for example, we can say "A is *hotter* than B," "C is *fairly cold*," "D is *very hot*," and so on.⁶ Between the opposite poles named by *hot* and *cold*, there is a mid interval, so that if something is neither hot nor cold, it might be *warm*, *cool*, or *lukewarm*, etc.⁷

Although there are also nouns (e.g., *friend/enemy*) and verbs (e.g., *love/hate* and *like/dislike*) which show properties of gradability, most attention has been given to the adjectives of this type, perhaps because the adjectives most clearly exhibit other characteristic properties of gradable opposites, such as implicit comparison, committedness, and markedness.

Implicit comparison can easily be seen in examples such as *big* and *little*, *tall* and *short*, *young* and *old*, and *hot* and *cold*. Something is described as big or tall or hot in comparison to other things of the same type. This means, for example, that a *tall child* is tall in comparison to other children of the same age, but may in fact be much shorter than a *short adult*, and that a *hot day* describes a hotter than average day, but an overall temperature that is much lower than the one described by a *hot oven*.

⁶Compare this to a typical pair of non-gradable adjectives such as *alive/dead*. Except in figurative uses, people do not say "A is *extremely dead*," or "B is *fairly alive*."

⁷The temperature scale is more lexicalized than most of the scales in that it has words to name the mid intervals; most gradable scales do not have names for mid intervals, but that does not mean that they are not there. Consider the SIZE scale for example. We know that it is possible for something to be neither big nor little but some size in between, even though we don't have a word to describe this. (We could use a phrase such as *of average size* though.)

Committedness involves an adjective's behavior in questions. An adjective is said to be *committed* if it implies a particular value when used in a question, and *impartial* or *uncommitted* if it does not have such an implication. For example, *tall* is uncommitted in a question like "How tall is Pat?" This question is neutral and can be used whether or not the speaker knows Pat's approximate height and whether Pat is tall, short or of average height. In contrast, the adjective *short* is committed; a speaker would only ask "How short is Pat?" if there is some reason to believe that Pat is shorter than average height. Many pairs of gradable antonyms contain one committed term and one uncommitted, e.g., *old/young*, *heavy/light*, *fast/slow*; many other pairs are made up of two committed terms, e.g., *innocent/guilty*, *beautiful/ugly*, *happy/sad*.

Markedness has been used as cover term for several related phenomena which distinguish the marked member of an antonym pair from the unmarked member. Lehrer (1985) discusses several of the criteria which have been proposed to define markedness. Committedness is one of them: the uncommitted member of an antonym pair is said to be unmarked and the committed member is said to be marked, so *old* is unmarked, while *young* is marked. It has also been noted that if the name of the semantic scale is morphologically related to one of the antonyms, it is related to the unmarked member, so for example, the name of the scale of LENGTH is related to the unmarked *long* rather than the marked *short*.

Another criterion of markedness is that the unmarked antonym can generally appear in a wider range of syntactic contexts; in particular, unmarked antonyms can occur with measure phrases but marked ones usually cannot, so we can say that something is *3 feet tall* but not *3 feet short*. Similarly, ratios are

usually only possible with the unmarked antonym; we can say that Kim is twice as old as Pat, but we can't say that Pat is twice as young as Kim. Morphology also plays a role: in pairs in which one antonym is derived from the other, the derived member is said to be marked, so *happy* is unmarked and *unhappy* is marked.

Most research on antonymy has focused on gradable opposites, antonyms in the narrow sense, perhaps because the properties described above are quite subtle and fascinating, but a few people, including Lyons (1977) and Cruse (1986) have tried to characterize the other sorts of commonly occurring opposites. These other types lack the special properties found with gradable opposites, but like them, they show a "dependence on dichotomization" (Lyons 1977). In other words, like antonyms in the narrow sense, these other types of opposites are also pairs of words which share some kind of semantic dimension.

1.2.2 Other types of opposites

The type of opposite which is most similar to the gradable opposite is the *complementary* opposite, sometimes known as the *contradictory*.⁸ Examples of complementaries include adjectives such as *true/false*, *dead/alive*, and *male/female*. Like the gradable adjectives, the complementary adjectives share a semantic dimension, but it is a dimension which has no middle values. As Cruse describes it:

[T]he essence of a pair of complementaries is that between them they exhaustively divide some conceptual domain into two mutually exclusive

⁸*Contradictories* and *contraries* are terms from logic which largely overlap with *complementaries* and *antonyms* as used by linguists. For a discussion of the differences between the logical terms and the linguistic terms, see Lyons (1977, . 272).

compartments, so that what does not fall into one of the compartments must necessarily fall into the other" (Cruse 1986, 198).

For example, we know that if a person is *not dead*, s/he must be *alive*; if a statement is not *true* it is *false*, and so on.

It is sometimes hard to decide whether a pair of opposites belongs in the set of gradable adjectives or in the set of complementaries, as in the case of *clean/dirty*. *Clean* and *dirty* are both gradable adjectives: we can say that something is *fairly clean*, *very clean*, *extremely dirty*, and we can say that *X is cleaner/dirtier than Y*. However, the scale of *clean* and *dirty* does not seem to have a middle term; whenever something is not clean, it can be described as dirty, so as Cruse says, it sounds strange to say "It's neither clean nor even slightly dirty."⁹ The case of *wet* and *dry* (the subject of my third case study) is similar in that there are words such as *damp* and *moist* which appear to name midpoints of the scale. However, as I will show, it seems that *damp* and *moist* are actually just more specific terms for types of wetness. We can gloss *damp* as 'slightly wet', but we cannot gloss a true midpoint word in this way (e.g., we can't gloss *lukewarm* as "slightly hot").

In addition to adjectives, verbs such as *pass/fail* and *obey/disobey*, nouns such as *day/night*, prepositions such as *in/out*, and adverbs such as *backwards/forwards* are also sometimes considered examples of complementaries.

Although by definition, complementaries are pairs which allow no logical middle term, in actual use, complementaries are sometimes used like gradable adjectives; for example, we can say that something is *almost true*, or that

⁹As Cruse points out, it sounds fine to say "It's neither clean nor dirty," because *dirty* is more often used for things that are distinctly dirty.

someone is *barely alive*. However, as Lyons (1977) points out, in these cases it may be the "secondary implications" of the words that are being graded rather than the main sense. That is, someone who is *barely alive* is actually entirely alive, but s/he is not as lively or energetic as most people are.

Directional opposites are another type of opposite, described in Lyons (1977) and in greater detail in Cruse (1986). These are generally adverbs or prepositions and include pairs such as *up/down*, *in/out*, and *clockwise/anticlockwise*.

Reversive opposites, described in Lehrer and Lehrer (1982) and Egan (1968), are yet another type of opposite. Egan describes reversional opposites in this way:

[T]hese comprise adjectives or adverbs which signify a quality or verbs or nouns which signify an act or state that reverse or undo the quality, act, or state of the other. Although they are neither contradictory nor contrary terms, they present a clear opposition. (Egan 1968, 27a)

This class contains many verbs, for example, *tie/untie*, *marry/divorce*, *enter/leave*, *appear/disappear*. Cruse and Lyons consider the reversional verbs to be a subtype of directional opposites, because they all describe activities which result in an object undergoing a change from one state to another; the two members of the reversional pair involve the same two states, but the direction of change is different in each case; for example, the verb *tie* means roughly 'to cause something to go from the state of being untied to the state of being tied,' while *untie* means 'to cause something to go from the state of being tied to the state of being untied.' Thus Cruse says the opposition seen in pairs of reversional verbs is similar to the kind of opposition in pairs of directional prepositions such as *to/from*.

Relational opposites (Cruse 1986) (also called *relative terms* (Egan 1968) and *conversive terms* (Lyons 1977)) include pairs such as *above/below*, *predecessor/successor*, *parent/child* and *teacher/student*. Egan describes these as "pairs of words which indicate such a relationship that one of them cannot be used without suggesting the other." Cruse considers this class to also be a subclass of the directional opposites. He says that these pairs "express a relationship between two entities by specifying the direction of one relative to the other along some axis." In examples such as *above/below*, this axis is spatial, but other examples (e.g., *ancestor/descendant*) involve "an analogical or metaphorical extension of spatial dimensions" (Cruse 1986, 231). Lyons points out that many opposites of this type involve social roles (*teacher/student*, *doctor/patient*) or kinship relations (*father/mother*), and these types of reciprocal relations have been well documented in many languages in the anthropological literature.

The various types of opposites discussed so far--antonyms, complementaries, directional opposites, and so on--all illustrate the essential properties of antonymy: they are pairs of words which simultaneously seem close and yet far apart in meaning, words which share some kind of semantic dimension but denote contrasting values along that dimension. But simply looking at the different types of opposites does not go very far in explaining what makes two particular words antonyms and why some words have no antonyms. To address these questions, it is useful to look at near-opposites, pairs of words which contrast in some way but which do not seem to be "real" opposites. Understanding what prevents some words from being opposites will lead to a deeper understanding of the essential properties of antonymy. In the

next section, some examples of near-opposites are given, along with the suggestions that have been given to explain why these pairs are just near-opposites rather than antonyms.

1.2.3 Near-opposites

In the linguistic literature in general, little attention has been paid to the topic of near-opposites. Two exceptions are Cruse (1986) and Lehrer and Lehrer (1982), who describe different types of near-opposites. Their explanations of why these pairs fail to be antonyms provide valuable clues to the nature of antonymy, clues that are followed up in the case studies later in this dissertation.

One of the types of near-opposites described by Cruse is what he calls "impure" opposites. He says that these are opposites "which encapsulate, or include within their meaning, a more elementary opposition." For example, he says:

[G]iant:dwarf can be said to encapsulate the opposition between large and small (but this opposition does not exhaust their meaning); likewise, shout and whisper encapsulate loud and soft, criticize and praise encapsulate good and bad..." (Cruse 1986,198)

Cruse's discussion of this example is brief but suggestive. For example, if the semantic opposition found in impure opposites "does not exhaust their meaning", then is it the case that the opposition does somehow exhaust the meaning of a pair of "pure" (prototypical) opposites? Cruse's examples of pure opposites are adjectives and his impure opposites are basically nouns or verbs, so the "impurity" may lie in the nominal or verbal elements of meaning associated with these words. It certainly does seem that with adjective pairs

such as *large/small* or *loud/quiet*, the shared semantic dimension (SPEED or TEMPERATURE) makes up the largest part of their meaning, while nominal or verbal pairs have additional semantic elements; *shout* and *whisper*, for example, are manner-of-speaking verbs, while *dwarf* and *giant* describe people. This idea of "purity of opposition" is vague, yet it may contain the seeds of an explanation of why there are more antonyms among adjectives than nouns or verbs.

Another type of near-opposite described by Cruse includes pairs that are only weakly contrasted because of "the difficulty of establishing what the relevant dimension or axis is" (Cruse 1986, 262). He gives *work/play* and *town/country* as examples. Cruse does not analyze these pairs at all, but it seems likely that they seem somewhat like opposites because in a particular context, there are two alternatives that provide an either/or choice. At the same time, we know that the two alternatives do not exhaust the logical possibilities (in addition to being at play or at work, a person could be asleep, for example), so they are not felt to be true opposites. It is not clear, however, how this type of near-opposite is different from prototypical pairs such as *black/white* (there are also other color terms) or *happy/sad* (there are other emotions, such as *awe*, *frustration* or *boredom*).

Finally, Cruse says that "nonpropositional meaning" is important to antonymy. He says that:

[A] good pair of opposites must be closely matched in respect of their nonpropositional meaning: that is why, for instance, *tubby* and *emaciated* are not fully satisfactory opposites, although they incorporate a binary directional opposition. (Cruse 1986, 262)

He does not explain this further, but apparently "nonpropositional meaning"

refers to the connotations associated with the words. *Tubby* is an adjective that has rather positive connotations in that it is used to affectionately describe someone who is somewhat overweight but not dangerously so, while *emaciated* is more impersonal (it might be used as part of a medical description while *tubby* would not be) and also conveys the idea that the thinness is possibly life-threatening.

This idea of nonpropositional meaning seems potentially useful in explaining the behavior of many pairs of near-opposites, but there are still some questions to be answered. First, what kinds of nonpropositional meaning are relevant? Is it enough simply to speak of some kind of general positive or negative connotation (for example to say that *tubby* is somewhat positive, while *emaciated* is quite negative)? Is register important? (*Tubby* is an informal word used in conversation more than in writing, while *emaciated* is a more formal word, likely to occur in medical literature). In the first two case studies of this dissertation (in Chapters Two and Three) I will return to this question. To anticipate, the synonyms *little* and *small* appear to differ in terms of register and associated connotations--for example, *little* is often used to convey a sense of "cuteness," which is not found so often with *small*. Likewise, the synonyms of *wet* described in the second case study have strong connotations--unpleasant ones in the case of *damp*, *dank*, *humid*, and positive in the case of *moist*.

Lehrer and Lehrer (1982) also describe a few different types of near-opposites, but these types and their explanations of why these pairs are near-opposites rather than antonyms are quite different from Cruse's. One type includes pairs such as *excellent/bad* and *hot/freezing*, in which both words are associated with the same semantic dimension, but one term describes a state

that is more "extreme" than the other. They account for these cases by saying, for example, that *excellent*, *good*, *bad*, and *terrible* all lie along the same semantic dimension, but that *excellent* and *terrible* name the extreme ends of the dimension, farthest from the midpoint, while *good* and *bad* name points nearer to the midpoint.

Based on these examples, Lehrer and Lehrer say that the definition of antonymy also needs the qualification that two antonyms should be the same distance from the midpoint, which would insure that *good/bad* and *excellent/terrible* are antonyms, while *excellent/bad* and *good/terrible* are only near-opposites. Similarly, in the case of temperature words, this qualification would rule out antonym pairs such as *freezing/warm* and *cold/boiling*. Although "distance from a midpoint" does seem to work well in explaining the behavior of these two sets of words, I cannot think of other cases in which this criterion is necessary, and Lehrer and Lehrer (1982) do not provide any others, so I wonder whether it is really necessary or whether the behavior of these two sets of words can be explained in some other way. Lehrer and Lehrer themselves acknowledge that distance from a midpoint cannot explain all cases of near-opposites, and they mention the case of *large* and *little* as one example.

Lehrer and Lehrer point out that *big* and *large* are near synonyms, associated with the same pole of the dimension of SIZE, with neither one more "extreme" than the other; the same is true of *small* and *little*. *Big* and *little* are antonyms, of course, as are *large* and *small*, but *big/small* and *large/little* might also be expected to form pairs of antonyms. Many speakers do in fact regard *big/small* as antonyms (I saw this pair in many of the lists of antonyms in the papers I read while doing this study, for example), but it seems that native

speakers do not feel that *large/little* are antonyms. Lehrer and Lehrer do not analyze this case in any detail, but they do make a useful suggestion. They note, "their distribution is not identical, since *big* and *little* but not *large* and *small* overlap with *older* and *younger*. *My big sister* is not equivalent to *my large sister* on one interpretation" (Lehrer and Lehrer 1982, 496). In the first case study, I look at the distribution patterns of these four adjectives by looking at the types of nouns each adjective typically modifies. I show that distribution patterns suggest an explanation of why *little* and *large* are not antonyms. This notion of difference in distribution may also explain the behavior of many other pairs of near-opposites, for example, pairs examined in the second case study such as *wet/arid* and *dry/damp*.

Lehrer and Lehrer have one more thought about near-opposites, suggesting that some near-opposites may be a result of multi-dimensional scales. They give an example involving words related to intelligence, saying that dictionaries of antonyms list many different words as possible antonyms of *clever*, including *dull*, *stupid*, *unintelligent*, *unskillful*, *slow*, *dense*, *simple*, and *foolish*. These "antonyms" of *clever* are all partial synonyms of each other, and *clever* itself has many partial synonyms, including *smart*, *able*, *intelligent*, *skillful*, *ingenious*. All of the synonyms of *clever* contrast at least loosely with all of the synonyms of *stupid*, but only a few of the possible pairs form good examples of antonyms, e.g., *skillful/unskillful* and *intelligent/unintelligent*. Most of the combinations seem only to form near-opposites, e.g., *clever/unskillful* or *ingenious/dense*. Lehrer and Lehrer suggest that while all these words are related to a scale of intelligence, this scale is not a simple one. They say:

Each of these words may be connected to the various ways of being clever. *Smart* and *intelligent* might be more appropriate for mathematical skills than for interpersonal relationships, whereas *skillful* may be more applicable to interpersonal relationships than to mathematical abilities. However, we do not have very strong intuitions on that point. (Lehrer and Lehrer 1982, 496)

The third case study, on *happy*, *unhappy* and *sad*, looks at a similar case, showing that each of the antonyms *sad* and *unhappy* corresponds to a different way of being happy.

1.2.4 Summary of the linguistic view of antonymy

Linguists studying antonymy have identified many different types of opposites, and they have also generally agreed to apply the term *antonymy* only to gradable opposites, a type of opposite which has many interesting properties such as markedness and committedness. Although these properties are interesting in themselves, they do not seem to provide any strong clues to the nature of antonymy.

A few linguists, namely Cruse (1986) and Lehrer and Lehrer (1982) have tried to identify the factors which determine whether two words are antonyms or just near-opposites. Each of these factors, in identifying what keeps two words from being antonyms, also makes an implication about the nature of antonymy, and together these factors seem likely to be useful in answering some of the "big questions" about antonymy, in particular, the questions of what makes two words antonyms and why some words have no antonym. These factors are listed in (a) to (f) of (2) below, and they will be examined in more detail in the case studies.

- (2) Factors which affect the "goodness" of a pair of opposites:
- a. *The purity of the semantic opposition:* In some pairs of near-opposites (e.g. *whisper/shout*) the semantic opposition does not exhaust the meaning of the words. The implication is that in prototypical pairs of opposites, the semantic contrast does in some sense exhaust the meanings of the words.
 - b. *The ease with which a semantic dimension can be imagined:* With near-opposite pairs such as *town/country*, it is hard to determine what the relevant semantic dimension could be. The implication is that for prototypical opposites, the semantic dimension can be easily identified.
 - c. *Correspondence of nonpropositional meaning:* Some near-opposites (e.g., *tubby/emaciated*) have very different connotations. The implication is that prototypical opposites are very similar in terms of non-propositional meaning.
 - d. *Distance from the midpoint of a semantic dimension:* In some cases, one member of a pair of near-opposites (e.g., *terrible*) seems to name a more extreme value than the other member (*good*). The implication is that prototypical opposites lie at equally distant points from a midpoint.
 - e. *Similarity in distribution:* For example, both *big* and its antonym *little* can refer to relative age rather than physical size in constructions such as *big sister* and *little brother*, but the near-opposite of *little*, that is, *large*, can only refer to physical size with the nouns *sister* and *brother*. The implication is that

prototypical opposites are similar in distribution.

- f. *Whether the semantic contrast involves a single dimension or multiple dimensions:* Some near-opposites seem to be located along different dimensions, although the dimensions seem to be related to the same general concept. The implication is that prototypical opposites share a single dimension, even if that dimension is related to a concept associated with multiple dimensions.

Linguists are not the only ones who have worked on antonymy. The next section presents the work of a lexicographer who takes a somewhat different approach to the nature of antonyms and near-opposite pairs.

1.3 A lexicographical approach to antonymy

Like the linguists mentioned above, Egan (1968), in her "Survey of the History of English Synonymy" which forms the *Introduction to Webster's New Dictionary of Synonyms*, is concerned with finding a useful definition of *antonym*, but where the linguists concentrate first on the properties of the more prototypical examples of antonyms, Egan focuses on pairs of near-opposites (which she calls "contrasted terms") to see how they are different from pairs of more prototypical antonyms. Thus Egan builds her definition of antonymy around phenomena that linguists consider only as an afterthought.

Egan begins by noting that the term *antonym* was coined as an opposite for *synonym* by C.J. Smith in 1867, who explained it in this way: "Words which agree in expressing one or more characteristic ideas in common he has called

Synonyms, those which negative one or more ideas he has called Antonyms."¹⁰ Egan points out that "negative" is quite vague, making it hard to know exactly what Smith had in mind, so she goes on to consider modern dictionary definitions of *antonym*. She finds that all of them contain the word *opposite*, but she points out that *opposite* itself has a wide range of meanings. She says, for example,

Opposition is a relation involved when two things are placed that: (1) they may be connected in a straight line...(as, *opposite* windows); (2) they lie at either ends of an axis, diameter, or the like (as, *opposite* points of the globe) [Egan lists 5 more senses, ending with] (8) they represent the obverse and the reverse (as, the *opposite* sides of a coin). (Egan 1968, 26a)

Egan says of this wide range of things which can be considered opposite in some way, "One can go no further than to say that *opposite* represents a setting of one thing against another so as to sharpen their differences or reveal their differences."

Looking at the kinds of words that are generally listed as antonyms in dictionaries and thesauri, Egan identifies several types of opposites, and they generally correspond to the types of opposites described above in sections 1.2.1 and 1.2.2. However, rather than focusing on the features which distinguish the various types, Egan looks for what is common to all of them by looking at what distinguishes all of these typical opposites from the near-opposites (contrasted terms). She finds that near-opposites "differ sharply in some part, but not all parts, of their meaning. They do not clash full force." More specifically, she says:

Put more simply, they [contrasted terms, i.e., near-opposites] differ (1) in their range of application or applicability, one being general, the other

¹⁰Egan quotes this from *Synonyms and Antonyms* by C.J. Smith, 1867.

specific, or one being more inclusive or less inclusive than the other, and (2) in their depth of meaning--that is, in the number and quality of implications contained in the terms. (Egan 1968, 29a)

She gives a few examples of this. One is the pair *rich/destitute*. She says that the meaning of *rich* is much broader and vaguer than that of the much more explicit *destitute*. She says "*rich* suggests more possessions than one needs," and thus the best opposite for *rich* is *poor*, which, she says, "suggests fewer possessions than one requires and so negates the full meaning of the other [that is, *rich*]. Unlike *poor*, though, *destitute* has some further implications, and these implications are not matched by contrasting implications in *rich*. She says:

Opulent could be opposed to *destitute* in narrowness and explicitness of meaning, for *destitute* suggests the miserable condition where one is deprived of all that is needed for bare existence and *opulent* the felicitous condition where everything that is desired is possessed in abundance. Though *rich* and *poor* come close together (the dividing line being marked by such a word as *comfortable*) and *destitute* and *opulent* are very far apart, being in fact "diametrical opposites," each represents the negation of the other. (Egan 1968, 29a)

Egan's observations are similar to Cruse's comments on purity of opposition and Lehrer and Lehrer's ideas of similarity of distribution, but unlike them, she takes these observations as central to understanding the nature of antonymy and bases her definition of antonymy on them. Her definition is:

An antonym is a word so opposed in meaning to another word, its equal in breadth or range of application, that it negates or nullifies every single one of its implications. (Egan 1968, 30a)

This definition gives a clear answer to the question of what makes two words antonyms--they are equal in range of application and yet are opposed in meaning--and at the same time provides a natural explanation of why two

words can contrast in meaning yet still not be antonyms--they may be different in their range of application. Thus, Egan's definition may be easier to apply to actual cases than the theories of antonymy described in section 1.2, and I have tried to do so in the three case studies that make up the main part of this dissertation.

In the case studies, I take up the notion of antonyms as words which are "equal in breadth or range of application", looking carefully to see if there are in fact differences between pairs of antonyms and pairs of near-opposites in terms of semantic range. For example, in the first case study of *big*, *little*, *large* and *small*, if Egan is right, it would be expected that the antonyms *large* and *small* are equal (which apparently means "identical") in their semantic range, while the pair of near-opposites *large* and *little* are quite different in their semantic ranges. But when applying Egan's definition in these terms, however, it immediately becomes clear that the definition as it stands is unlikely to be entirely correct.

First of all, given the fact that there are no perfect synonyms (words that absolutely identical in terms of register, range of application and implications), it would seem unlikely that there are any antonyms that contrast perfectly. In fact, the case of *happy* and its two antonyms, *sad* and *unhappy* (the subject of the third case study), demonstrates this point. *Unhappy* and *sad* are near, but not perfect synonyms, so they are not exactly the same in terms of range of application or implications. *Happy* cannot be "equal in breadth or range of application" to both *unhappy* and *sad* simultaneously when these two words themselves are not exactly the same. However, I think that Egan has identified some very important factors which, with some modification, can go a long way

toward explaining the phenomena of antonymy. In the case studies, I show how "range of application" can be used to explain many of the phenomena associated with antonymy; specifically, I show that antonyms, while not completely identical in their range of application, share a great deal of semantic range and have more shared range in common than near-opposites do.

The clang phenomenon (which she calls "clash") also may require some modification of Egan's definition. She says that two words "clash satisfactorily" when they are equal in range of application and all of their implications contrast, and she gives a few examples of this, including the pairs *rich/poor* and *destitute/opulent* in the quotation above. Although it seems that for Egan, both of these pairs are equally contrasting, I believe that most native speakers (including myself) feel a strong sense of "clash" with the first pair, which seem to be prototypical antonyms, but only a weak sense with the second. *Destitute* and *opulent* seem more like near-opposites to me.¹¹ Before the case studies then, it will be useful to look at the research into antonymy done by psychologists and psycholinguists since they have looked more deeply into the clang phenomenon.

1.4 Psycholinguistic approaches to antonymy

As shown above, linguists and lexicographers studying antonymy have

¹¹Actually, it would be possible to explain the difference in intensity of "clash" in Egan's own terms by saying that she just has not considered the range of application and implications of these two words carefully enough; for example, *destitute* is used to describe people, but *opulent* can be (and I think more often is) applied to people's clothes, homes, lifestyles and so on rather than to people themselves. (Cf. *destitute man*, *?opulent man*; **destitute home*, *opulent home*) Secondly, it does not seem that *opulent* exactly negates all the implications of *destitute*. If *destitute* means "lacking the basic essentials", shouldn't its antonym mean "having more than is needed of the basic essentials?" Instead, *opulent* implies the possession of luxuries in addition to the basic essentials.

focused mainly on the semantic aspects of antonymy and on identifying different types of opposites. Psychologists and psycholinguists, in contrast, seem to be more interested in the fact that antonyms have a strong association, an association which seems to be the basis of the "clang" phenomenon. In this way, their work provides a counterpoint to the linguistic approaches described earlier, answering some of the questions which linguists have not dealt with.

1.4.1 Antonymy as association

One of the first psychologists to study antonymy systematically was John Deese. In his 1965 book, *The Structure of Associations in Language and Thought*, Deese seeks to understand how word meanings are organized in the mind. He uses the method of free-association, presenting informants with words and asking them to respond. By looking at the pattern of responses to a particular word, he arrives at a psychological "definition" of the word. As he explains it:

The distribution of responses evoked by a particular word as stimulus defines the meaning of that word. The only ways in which such meaning can be specified are (1) by the nature of the distribution itself and (2) by the relation that distribution has to distributions of responses to other linguistic forms as stimuli. (Deese 1965, 43)

In other words, Deese investigates the meaning of a word by asking many different respondents to give a response (a free association) to it. He then looks at the range of responses and compares it to the range of responses given for other target words. Deese feels that he can measure the similarity of two words by comparing the sets of responses, so that "...two stimuli may be said to have the same associative meaning when the distributions are identical, and they may be said to share meaning to the extent that these distributions agree"

(Deese 1965, 45-46)

Since Deese is examining the "associational" meaning of words, his definitions are somewhat different than the kinds of definitions usually given by linguists and lexicographers. Consider, for example, some of the definitions from the Associative Dictionary at the end of his book (Deese 1965): "Rabbit: Animal, Dog" and "Bronze: Metal. Gold, Copper." While *rabbit* may cause people to think of dogs because dogs are often used in hunting rabbits or because dogs like to chase rabbits, it is unlikely that a dictionary definition of *rabbit* would include this information. Likewise, *gold* is not likely to be listed in a dictionary definition of *bronze*.

However, Deese does uncover some of the same kinds of semantic relations linguists describe, including antonymy. A few nouns and verbs produce antonyms as responses, but he notes that many adjectives, especially the common ones, are strongly associated with their antonyms. This fact is so striking to Deese that he proposes that the entire adjectival vocabulary of English might be organized in the mind around the relation of antonymy. Frequently occurring adjectives are associated with their opposites, and the remaining adjectives can be characterized in terms of their associations with the common adjectives. As an example of the latter, in the Associative Dictionary, *famous* is said to share the meaning of *good* and *rich*, which are in turn defined in terms of their opposites *bad* and *poor*.

Given this proposal, Deese says that it is important to find a way of determining whether or not two words are antonyms which does not rely on subjective judgments. He makes this suggestion:

The notion of contrast implies that one member of a pair of words should have its associative meaning most strongly determined by the other member of the pair and that the meaning should be reciprocal. Contrasting words, in brief, are words that are defined in terms of each other... (Deese 1965, 122)

In other words, Deese suggests, because opposites are so strongly associated, if a person is given one member of an antonym pair and asked to give another word in response, he/she is highly likely respond with the antonym. This kind of behavioral response can serve as an identification procedure for antonyms.

Deese used a simple word association test to identify a set of basic antonyms which he could then use to define the meanings of the less frequent adjectives which do not necessarily have antonyms. He says that in this kind of association test, "[A]ll pairs of words in which the stimulus to one is the most frequently occurring response to the other are the best candidates for being reciprocal pairs." That is, if most people respond with *hot* when they are given *cold* and if most respond with *cold* when they are given *hot*, then *hot* and *cold* would be identified as a reciprocal, contrasting pair, i.e., as antonyms.¹² Starting with a set of 278 commonly occurring adjectives,¹³ Deese used a free association test given to undergraduate students to arrive at a set of 39 basic contrasting pairs. These are listed below in Figure 1.

Most of the pairs on this list conform to the standard described above; that is, each word was the most frequent response to the other, but in a few cases (those marked with an asterisk), one of the words was the second

¹²Pairs which contrast in meaning in some way but which do not produce this kind of reciprocal response are not considered to be antonyms in Deese's model.

¹³The list contained adjectives that occurred more than 50 times per million words in the Thorndike-Lorge vocabulary, with some participles weeded out.

most common response; this occurred when one word had two opposites, e.g., *easy/hard* and *hard/soft*. *Hard* was the most common response for both *easy* and *soft*, *soft* was the most frequent response for *hard* and *easy* was the second most frequent.

alone/together	dark/light	hard/soft	old*/young
active/passive	deep/shallow	heavy/light*	poor/rich
alive/dead	dry/wet	high/low	pretty/ugly
back/front	easy/hard*	inside/outside	right/wrong
bad/good	empty/full	large/small	rough/smooth
big/little	far/near	left/right*	short/tall
black/white	fast/slow	long/short*	sour/sweet
bottom/top	few/many	married/single	strong/weak
clean/dirty	first/last	narrow/wide	thick/thin
cold/hot	happy/sad	new/old	

Figure 1. Deese's list of 39 contrasting pairs

There is nothing surprising in the list of opposites produced by this association test; that is, all of the pairs seem like very good, prototypical, opposites.¹⁴ Indeed, as Deese says,

The most important aspect of this table [of the 39 contrasting pairs] is the extent to which the principle which generated it produces pairs which do not violate intuitive expectation. (Deese 1965, 122)

Deese notes that his method of identifying antonyms may fail to produce all the English pairs which seem to be opposites, however. He gives the example of

¹⁴Although all of the reciprocal pairs of adjectives identified by Deese's method are antonyms, this is not true of the reciprocal pairs found in other parts of speech. Deese notes that some nouns form reciprocal pairs (that is, the two nouns each occur as the most common response to the other), although this is quite rare. Sometimes, these nouns show some kind of semantic contrast as in *mother/father* and *hand/foot*, but with other pairs, the reciprocal association does not involve much contrast at all, as in *house/home* and *clock/time*.

brave and *cowardly*; both of these words were in the original sample, but the undergraduates did not respond to *brave* with *cowardly* or vice-versa. Deese gives this explanation:

It is very possible that those antonyms not appearing by the criterion--allowing for the possibility of sampling error--are not in the cognitive sense defined by each other. We do not think of *brave* as being the opposite of *cowardly*; we think of *brave* as being *strong* and *cowardly* as being *afraid*. *Cowardly*, *brave*, and *afraid* are all related in a basic way to the pair *strong-weak*. (Deese 1965, 123)

Brave and *cowardly* seem to me to be examples of near-opposites, words which certainly contrast in meaning but lack the "clang" effect; this suggests that the strong association between antonyms which Deese uncovered may be behind the clang phenomenon. All of the adjectives on Deese's list of 39 "sound" like antonyms, and most native speakers would readily identify them as such.

In this work, Deese's concern is simply associations between words. He does not look any deeper into the semantics of the kinds of associations he discovers, and so although his work goes a long way toward answering one of my initial questions, the question about the basis of the clang phenomenon, it does not really touch on the others. Deese does not try to explain why two particular words are associated as antonyms, or why some words have more than one antonym while others have none. But some psycholinguists, in particular, George Miller and his colleagues working on the WordNet project at Princeton, have tried to deal with these questions. They take Deese's idea that all adjectival meanings can be organized around a set of basic antonyms and develop it into a more complete model of antonymy. This work is presented in the next section.

1.4.2 The WordNet model of antonymy

WordNet is an on-line lexical database developed by George Miller and his colleagues at Princeton University. Although the organization of this lexical database is intended to reflect the organization of words in the mind, it is not meant to be an exact model of the mental lexicon. In the case of antonymy and adjectival meanings, however, Miller and colleagues have done some experimental research designed to test the extent to which the WordNet model is consistent with the actual structure of the mind. The WordNet project is described in many papers, including Miller (1990) and Miller and Fellbaum (1991), but in this section, I will focus on an earlier research report from the Princeton Cognitive Science Laboratory (Gross, Fischer and Miller 1988) because this paper focuses on antonymy in WordNet and so it most fully explains Miller and his colleagues' view of antonymy as the basis for the organization of the meanings of adjectives. This report also deals with some of the implications and questions which arise from this model of antonymy but which are not discussed much in the later papers.

1.4.2.1 The role of antonymy in organizing adjectival meanings

In the 1988 report entitled "Antonymy and the Representation of Adjectival Meanings," Derek Gross, Ute Fischer, and George Miller set out the view of antonymy and adjectival meaning which is adopted in WordNet, and they also report on the results of two experiments designed to test the psychological plausibility of this model. They start with Deese's idea that the meanings of adjectives are organized on the basis of antonymy; like Deese, they propose that there is a basic set of antonym pairs, and that most adjectives

are either a member of one of these basic sets or a synonym¹⁵ of one of them.

They begin by pointing out that although speakers can easily recognize antonyms, it is hard to define exactly what an antonym is. They say that although lexicographers such as Egan identify many different types of opposites, in general, they can account for the various types of antonyms with just two assumptions:

- (1) the semantic role of adjectives is to express values of attributes (properties, qualities, features, or dimensions); and (2) attributes are bipolar in nature. (Gross, Fischer and Miller 1988, 2)

It seems from these assumptions, then, that Gross, Fischer, and Miller believe that antonymy in a sense reflects the world--the bipolarity found in antonymy reflects the fact that many attributes in the world are bipolar. They give as examples the attribute of size (with the bipolar values *large* and *small*), the attribute of sex (which has two values, *female* and *male*) and what they call the "evaluative" attribute, which ranges from *good* to *bad*. They note that there are some exceptional kinds of attributes, e.g., color, which has many values (*red*, *yellow*, *green*, etc.), but they say that by and large, most attributes are bipolar. They note that several different types of opposites, e.g., Egan's complementaries, contradictories and reverse terms, all fit these two assumptions. In other words, all types of opposites involve bipolar attributes.

Gross, Fischer, and Miller go on to say that in building a model which will organize the meanings of most adjectives on the basis of antonymy, they will need to account for the fact that many adjectives do not seem to have antonyms.

¹⁵Gross, Fischer, and Miller (1988) use *synonym* rather widely, both to describe words that are quite close in meaning (e.g., *lifeless* and *dead*) and to describe words which are not so similar in meaning but which could be considered synonymous in a particular context (e.g., *alive* and *bouncy*).

This includes the kinds of adjectives which Egan called "contrasted" or "loosely contrasted." As an example, they say that the adjective *musty* contrasts in meaning with *fresh*, but that the antonym of *fresh* is *stale*, not *musty*. However, they note, *musty* is similar in meaning to *stale*, so they propose a third assumption to account for this. They say:

(3) Any adjective with no direct antonym will be similar in meaning to some adjective with a direct antonym. (Gross, Fischer and Miller 1988, 3)

Thus, in theory, all adjectives can be accounted for. Many adjectives have clear examples of antonyms--these are called *direct antonyms*--and the adjectives which do not have a direct antonym are synonyms of adjectives which do, and thus can participate in antonymy indirectly. Adjective pairs which are thus "mediated" by a pair of direct antonyms are called *indirect antonyms*. As another example, Gross, Fischer, and Miller say, "The incompatible pair *vigilant/careless* is mediated, presumably, by the direct antonym *careful*." That is, in their model, the antonyms *careful* and *careless* are directly linked as antonyms, and *vigilant* is linked to *careful* as a synonym, thereby being indirectly linked to *careless*.¹⁶

Although synonymy and antonymy are both ways of linking words together in the lexicon, Gross, Fischer, and Miller argue that there is a crucial difference between them. They say that while synonymy is "a relation between lexical concepts," antonymy is "a relation between words, not concepts." This point is somewhat difficult to understand, but I think it can be made clearer with one of their examples, one which involves the basic antonym pair *wet* and *dry*

¹⁶Pairs in which both adjectives are linked to antonyms (e.g., *parched/soggy*, in which *parched* is linked to *dry* and *soggy* is linked to *wet*) are also considered indirect antonyms in this model.

and the several adjectives which are related to these antonyms. These include *damp*, *moist*, and *waterlogged*, which they consider to be synonyms of *wet*, and *baked*, *arid*, and *parched*, which they consider to be synonyms of *dry*. Gross, Fischer, and Miller say that the synonyms of *wet* form a conceptual "cluster" in semantic memory, and likewise the synonyms of *dry*. These two semantic clusters are held together by the fact that they share a bipolar attribute. Although these two conceptual clusters contrast in meaning, they argue that antonymy is a relationship between particular words, so that it cannot be said that the clusters themselves are antonyms. Instead, *wet* and *dry* are somehow selected as the "labels" for the two poles of the attribute, and the fact that they are chosen as labels makes them direct antonyms.

As further evidence that antonymy is a relationship between words, not concepts, Gross, Fischer, and Miller note that most antonyms in English are formed by prefixes such as *un-* and *in-*, saying, "When an antonym is formed by affixation, the direct bond between the two adjectives is apparent from their shared root."¹⁷

1.4.2.2 A test of the model

Gross, Fischer, and Miller say, "One of the implications of the lexical organization that has been proposed here is that direct antonyms should be easier to recognize than indirect antonyms." In order to test this prediction, Gross, Fischer, and Miller used native speakers' ratings of semantic distance to compile lists of direct and indirect antonyms. For example, they took antonym pairs such as *alive* and *dead* and for each adjective found a "near" synonym (a

¹⁷They do not mention, however, that these prefixes do not always form antonyms (consider *easy* and *uneasy*) or that the most frequently occurring antonyms are not formed by affixation (for example, none of the 39 pairs on Deese's list in Figure 1 are formed this way).

word judged to be close in meaning to the original) and a "far" synonym (a word similar in meaning to one of the less common senses or to a figurative sense of the original). For example, *living* was chosen as a near synonym of *alive* and *bouncy* as a far synonym of *alive*, while *lifeless* was chosen as a near synonym of *dead* and *asleep* as a far synonym of *dead*. Next they made lists of different types of word pairs, including pairs of direct antonyms (e.g. *near/far*), pairs of indirect antonyms using near synonyms (*far/proximate*), pairs of indirect antonyms using far synonyms (*far/juxtaposed*), pairs using one of the adjectives from the antonym list and one color term (*far/violet*) and pairs using unrelated adjectives or color terms (*regretful/clumsy* and *chartreuse/brown*).

Next Gross, Fischer, and Miller presented these various pairs to a new set of subjects, telling them it was a "color study." First, without using the words *antonym* or *opposite*, they trained the subjects to recognize an unnamed "relation" by showing them two lists. One list had pairs of opposites--including direct pairs and both types of indirect pairs--and the other list had the pairs of unrelated adjectives and color terms. They asked the subjects if they could identify the relation that held between all the things on the one list but not the other. When the subjects said that they recognized the relation, they were presented with pairs of words on a computer screen. As each pair appeared, the subject was asked to press one key if the pair showed the relation they had been trained to recognize, and another key if it did not. The reaction time was measured.

This experiment found that subjects were able to identify pairs of direct antonyms faster than pairs of indirect antonyms, and that they identified the near indirect antonyms faster than the far ones. Gross, Fischer, and Miller conclude

that these results are compatible with their model--if direct antonyms are immediately linked in the mind, it would make sense that they can be identified faster than indirect antonyms which are not so linked.

Gross, Fischer, and Miller then conducted a second experiment to see if subjects could distinguish between direct and indirect antonyms. It was much like the first one, except that this time, they were asked to identify the relation that held among the pairs of direct antonyms only. Indirect antonyms in this case were grouped with the non-antonyms. They found that subjects were able to identify the direct antonyms quickly and to distinguish them from the indirect antonyms, although they had some difficulty in making decisions about the indirect antonym pairs that contained near synonyms (e.g, *alive/lifeless*). Gross, Fischer, and Miller admit that the results of these two experiments do not prove the validity of their model, but they say, "the conservative statement is that the results of experiments 1 and 2 do not refute the hypothesis that they were designed to test" (Gross, Fischer and Miller 1988, 12).

Gross, Fischer, and Miller consider another possible explanation of their findings, a hypothesis that they call the "co-occurrence hypothesis." They note that some words are used together more often than expected by mere chance, and that this is probably especially true in the case of antonyms since antonyms label the two poles of a single attribute. In other words, the co-occurrence hypothesis is that an adjective such as *cold* is likely to be used together (in the same sentence, paragraph, or stretch of speech) with *hot*, and it is more likely to occur with *hot* than with any of the synonyms of *hot* such as *scorching* or *ardent* (the near and far synonyms of *hot* in this experiment). They go on to say that the co-occurrence hypothesis, if true, could probably explain why subjects identified

direct antonyms faster than indirect antonyms, but they add, "We would object, however, that the co-occurrence hypothesis does not go far enough to explain all that we know about antonymous English adjectives" (Gross, Fischer and Miller 1988, 13). In particular, they feel that the large number of morphologically marked antonyms in English cannot be explained by co-occurrence alone; instead, the prefixes show a particularly tight bond between two antonyms that they feel is better explained through a model which directly links antonyms than through the co-occurrence hypothesis.

1.4.2.3 Summary of the WordNet model

The model of antonymy developed in Gross, Fischer, and Miller (1988) seems to provide answers to several of the questions raised in section 1.1. It offers a description of where semantic dimensions come from--they reflect bipolar attributes that exist in the world (or at least in the way we experience the world). It has an explanation for what makes two words antonyms--antonyms are adjectives which have (somehow) been chosen to label the two opposing poles of an attribute. It also explains why some pairs of opposites are better examples of antonyms than others--direct antonyms are linked directly in the mind which causes them to be more easily recognized as antonyms. The indirect antonyms (those which are not directly linked) are not so quickly recognized and so seem less prototypical.

This model does not answer all the questions, however. How exactly is one adjective (out of a set of synonyms) chosen to represent the pole of an attribute? And what is the explanation for the fact that some words have more than one antonym, for example, *happy*, which has two antonyms *unhappy* and

sad?¹⁸ Are both *unhappy* and *sad* somehow chosen as labels for the same attribute? Or is *happy* chosen as the label for two different (though related) attributes? Is the distinction between direct and indirect links enough to account for the clang phenomenon, or is that better accounted for simply through frequency of association? What about the fact that subjects had a lot of trouble categorizing the indirect pairs that used near synonyms (e.g, *alive/lifeless* and *living/dead*)? The fact that the subjects had difficulty deciding whether such pairs were direct or indirect antonyms argues against a structural difference between these two types of antonyms. And most importantly, even though the model does explain a great deal, is there any evidence that the mind is actually organized in this way? Or could other models explain the antonymy equally well or better? Gross, Fischer, and Miller themselves acknowledge that the results of their experiments did not provide strong support for their model, although they did not contradict it either. However, their model did lead to further study by researchers seeking to answer these questions. Some of this research is discussed in the rest of this chapter.

1.5 Responses to the WordNet model of antonymy

1.5.1 A study of the contexts of antonymous adjectives

Perhaps the most interesting question raised by the WordNet model is the question of how two particular adjectives come to be associated as antonyms. Since in this model, antonymy is a relationship between two "word forms" rather than between two concepts, the meanings of the words cannot be

¹⁸Gross, Fischer, and Miller note the existence of words that have more than one antonym, and they claim (without any justification that I can see) that this is additional evidence that antonymy is a relationship between particular words, not between concepts. However, they do not explain how their model accounts for one word with two antonyms.

used as an explanation. In other words, Egan's kind of explanation, that two words are antonyms because they are exactly opposite in meaning and implications, is not considered satisfactory in this model. So the question remains, how do people learn which words are antonyms?

Charles and Miller (1989) attempt to answer this question. They propose two possible answers to this question, one of which they go on to study in greater detail. Their two hypotheses are called the co-occurrence hypothesis and the substitutability hypothesis, and they are quoted below:

Co-occurrence: Two adjectives are learned as direct antonyms because they are heard together and are spoken together in the same sentences more often than chance would allow.

Substitutability: Two adjectives are learned as direct antonyms because they are interchangeable in most contexts, i.e., because any noun phrase that can be modified by one member of the pair can also be modified by the other (Charles and Miller 1989, 360).

Charles and Miller say that the co-occurrence hypothesis is simpler, but that there is an independent need for the substitutability hypothesis, for example, to explain other facts of word association and to explain how people learn which part of speech a word belongs to, so they decided to test the substitutability hypothesis first to see if it can explain how people learn antonyms.

This was their basic procedure. First, they chose a set of four adjectives to study. The set was made up of a pair of antonyms, *strong* and *weak*, and two synonyms of these antonyms, namely *powerful* and *faint*. The synonyms were used so that they could compare the results found with the antonym pair *strong/weak* to the results found with the indirect antonym (near-opposite) pairs *strong/faint* and *powerful/weak*. For each adjective, they found 25 sentences containing the adjective in the Brown corpus. Then they typed each sentence

on a separate card, leaving a blank in place of the target word, as in this example: "Third, there were those, notably Patrice Lumumba, who favored a Unified Congo with a very _____ central government." Subjects were asked to sort the cards, placing all the cards which were missing the same word together; then they were asked to fill in the missing word. Charles and Miller predicted that if the substitutability hypothesis were correct, subjects would confuse the contexts for the direct antonyms, but not for the indirect pairs. That is, they would be likely to infer that the missing word was *weak* when in fact it was *strong* and vice versa, but they would not infer that the missing word was *faint* when in fact it was *strong* or *weak* when it was *powerful*.

The results seemed to argue against the substitutability hypothesis: Subjects rarely confused the contexts of the antonyms. For example, when subjects were asked to sort out the sentences which contained *strong* and *weak*, *strong* was guessed correctly in 179 responses out of 200 (200 responses = 25 sentences filled in by eight different subjects), and *weak* was substituted for *strong* only 4 times. (The other 17 responses were other adjectives, e.g., *ardent*, *great*, etc.) Similarly, *weak* was guessed correctly in 164 responses out of 200, and *strong* was substituted for *weak* in only 13 out of 200 responses. Subjects did not substitute the indirect antonyms for each other either. When the contexts for *strong* and *faint* were sorted, for example, the subjects correctly identified *strong* in 113 cases out of 200 and only substituted *weak* for *strong* in one case; they correctly identified *faint* only 46 out of 200 times, and substituted *faint* for *strong* only 23 out of 200 times. Overall, subjects had a harder time identifying *faint* and *powerful* from the context than they did *strong* and *weak*.

Charles and Miller repeated the experiment with four other adjectives, the antonyms *public* and *private* and their synonyms *open* and *secret*. The results were similar to those in the first experiment. *Private* and *public* were easier to identify from the context than *secret* and *open*, but they were rarely substituted for each. Charles and Miller say that these results argue against the substitutability hypothesis, since they show that the sentential contexts of direct antonyms are not interchangeable. However, they say, "It may still be true, however, that a noun phrase that can incorporate one adjective can also incorporate its antonym. Disambiguation may be based on the rest of the sentence." Therefore, they did another experiment, this time cutting down the amount of context to just the noun phrase. They gave these examples of the types of noun phrase contexts that were used:

- (3) a very _____ central government
 _____, sickly, little things
 the _____ eye
 many _____ sources

In this experiment, the noun contexts for four adjectives were compared: *strong*, *weak*, *public* and *private* (that is, this experiment had no indirect antonyms), and the subjects were told what the adjectives were. Again, they were told to sort the cards so that the cards in each stack could be filled in with the same missing word. This time, subjects substituted antonyms for each other more often than in the previous experiment, but they substituted the unrelated adjectives too. For example, the context of *strong* was correctly identified in 132 out of 200 cases, but in 29 examples, *weak* was chosen instead, in 27 cases *public* was chosen, and 12 cases *private* was chosen. Therefore, it seems that the contexts of antonyms are not interchangeable.

Charles and Miller say that even within a noun phrase, there is enough disambiguating evidence that the correct adjective can often be identified. In the example " _____, sickly, little things" the other adjectives *sickly* and *little* help the subject to choose *weak*. They admit that some of the noun phrases form collocations that are "almost compound nouns." But, they say:

The point, however, is that it is simply not true that any noun phrase that can take one adjective can equally well take its direct antonym. Hence, even though substitutability might provide an important cue for learning which contrasting adjectives to pair as direct antonyms, it is still not an attractive hypothesis. (Charles and Miller 1988, 372)

Since Charles and Miller believe that the results of this study show that antonyms are not learned through substitution, they turn to the co-occurrence hypothesis as a more likely explanation. That is, they suggest that speakers learn that *strong* and *weak* are antonyms (but *strong* and *faint* are not) because they encounter *strong* and *weak* (but not *strong* and *faint*) used together in the same sentence fairly often. The Brown corpus is too small to accurately measure co-occurrence, but they note that they found five examples in which *strong* and *weak* occurred in the same sentence, but no sentences with *strong* and *faint*, *weak* and *powerful*, or *faint* and *powerful*.

Other adjectives occur more frequently in the Brown corpus and so provide might a better test of the co-occurrence hypothesis; this is the case with *big*, *little*, *large* and *small*,¹⁹ a set which Charles and Miller find particularly interesting because *big* and *large* are similar in meaning and therefore, according to the WordNet model, associated with the same concept; *small* and *little* are likewise associated with the same concept. Considering the fact that *large* is usually paired with *small* rather than *little* and *big* is usually paired with

¹⁷*Big* occurs 316 times, *little* 273 times, *large* 347 times and *small* 504 times.

little rather than *small*, Charles and Miller say, "If co-occurrence in the Brown Corpus could be shown for these adjective pairs [*large/small* and *big/little*], it should serve to discount any proposal that antonymous pairing reflects some kind of conceptual or semantic generalizations" (Charles and Miller 1988).

Counting all the co-occurrences in the corpus, they find that *big/little* and *large/small* co-occur more often in the same sentence than chance would predict and three times more often than either *big/small* or *large/little*. Charles and Miller say that given these findings,

[I]t seems reasonable to suppose that the antonymous pairings--the "clang" association between direct antonyms--is a consequence of frequently perceiving and using these words together in the same syntactic structures. (Charles and Miller 1988, 374)²⁰

Thus according to Charles and Miller, the results of this study help to answer one of the biggest questions raised by Deese's research, the question of how antonyms come to be associated in the mind. They say that the substitution hypothesis does not seem plausible and that co-occurrence can probably explain how antonyms become associated. Charles and Miller believe that this in turn supports the characterization of antonymy as a lexical rather than conceptual relation--if antonyms are linked as a result of co-occurrence, there appears to be no need to rely on word meanings to explain their association. However, there are many researchers who do not accept this characterization of antonymy, and as will be shown in the next sections, they

²⁰*Big* and *little* occur together in 12 sentences, but would be expected to occur together in 1.6 if they were occurring independently. *Large* and *small* occur together in 26 sentences (3.2 expected by chance), *large* and *little* occur together in only 3 sentences (1.7 expected by chance), and *big* and *small* occur together in only 4 sentences (2.9 expected by chance). Charles and Miller do not discuss the fact that some English speakers do in fact pair *big* with *small* as antonyms, and that *big/small* sounds much better than *large/little*, even though neither of these pairs co-occur very often in the Brown corpus.

interpret Charles and Miller's findings in very different ways.

1.5.2 A critique of the view of antonymy as a relation between word forms

Murphy and Andrew (1992) are two linguists who are somewhat skeptical of the claim that there is no role for conceptual meaning in antonymy because antonymy is only a relationship between word forms. In their 1992 research report, they offer some apt criticisms of the WordNet model and describe their own studies which suggest that conceptual meaning plays a role in antonymy just as it does in synonymy.

Murphy and Andrew start by discussing Charles and Miller's suggestion that co-occurrence alone can explain how two words come to be associated. They begin by pointing out, "There are a number of other words that co-occur but are not even near-antonyms: *dentist/teeth*, *food/eat*, *sick/tired*, and a variety of idioms and phrasal expressions" (Murphy and Andrew 1992, 3). How do word learners know that some kinds of co-occurring words are antonyms, but others are not?²¹

Their second criticism is even more to the point. They say:

[A]n explanation is needed for why antonyms co-occur. If it is because they are associated in semantic memory, then we have formed a completely circular explanation in which co-occurrence is caused by the relation, and the relation is caused by co-occurrence. (Murphy and Andrew 1992, 3-4)

This also occurred to me the first time I read Charles and Miller (1989). It seems that the only way out of this circularity is to take the meanings of the words into

²¹In sections 1.5.3 and 1.5.4 below, a partial explanation will be discussed: when antonyms occur together in the same sentence, they usually are found in one of a set of particular syntactic frames that emphasize the contrast in meaning. Noun-verb pairs such as *food/eat* and noun-noun pairs such as *dentist/teeth* do not occur in these frames.

account. After all, when people speak, they choose words to express their ideas; if people are choosing to use antonyms together in the same sentence, it is because they intend to convey a contrast in meaning and they are choosing words which effectively convey that contrast. As Murphy and Andrew put it, "If antonymy is just a kind of lexical association, then the semantic component would be superfluous, whereas in fact it seems to be the crucial element" (Murphy and Andrew 1992, 4).

Murphy and Andrew's final criticism is that an increasing amount of the research done on word meanings suggests that "[S]emantic relations should be definable in terms of conceptual relations," but in the WordNet model, antonymy is a lexical association that is not based on conceptual relations.

Murphy and Andrew believe that word meanings are mentally represented as concepts and that semantic relations, including antonymy, can be "computed" by a language user as necessary. To test whether this in fact is true of antonymy, they conducted a few simple experiments making use of research on conceptual combination which shows that the meaning of an adjective depends to some extent on the noun it is modifying. As an example, they point out that the meaning of *fresh* is somewhat different in the phrases *fresh shirt*, *fresh idea* and *fresh fish*. If, as Gross, Fischer and Miller claim, antonymy is a relationship between word forms and not concepts, Murphy and Andrew say that *fresh* should have a single antonym, and this antonym should be the same regardless of the context in which *fresh* occurs. If, on the other hand, antonymy is a relation between concepts, then an adjective would be expected to have different antonyms in different contexts because the same adjective evokes different concepts (each with its own antonym) in different

contexts.

Murphy and Andrew test this prediction by giving subjects two lists, one of adjectives in isolation and one of the same adjectives modifying nouns. They then asked the subjects to give the antonym of the adjective. For example, they were asked in the first part of the test to produce antonyms for *fresh*, *cold*, *dry* and so on, and in the second part, they were asked to give the antonym of the adjective in phrases such as *fresh fish*, *fresh shirt*, *cold water*, *cold facts*, etc. In the second part, each adjective occurred in two noun phrases, one which Murphy and Andrew felt would have the same antonym as the adjective in isolation and one which would probably be different. For example, they expected *hot* to be given as the antonym of *cold* in isolation, and they also expected *hot* to be given as the antonym of *cold* in the phrase *cold water* but not in the phrase *cold facts*.

The results matched their expectations quite well. They found that subjects generally did not give the same antonym for the same adjective in all three cases. They report, "Only half the time did subjects give the same opposite to an adjective when it was alone as when it was in a phrase" (Murphy and Andrew 1992, 12). They feel that this is lower than would be expected by the account of antonymy as a relation between word forms; since the same word form is used in all cases, that account would predict a match rate of close to 100%, much higher than the actual rate found by Murphy and Andrew.

Murphy and Andrew found similar rates of matching in another experiment in which the subjects were asked to give synonyms of adjectives alone and in context. They argue that these results show that antonymy and synonymy are similar types of relations and are not structurally different as in

the WordNet model.

In their discussion, Murphy and Andrew say that although some synonyms and antonyms may be "pre-stored" in the mental lexicon as a result of associations that arise from frequent co-occurrence, the relations of synonymy and antonymy are basically conceptual in nature and so they can also be figured out when necessary. They say:

Synonyms can be computed via conceptual similarity--they do not have to be pre-stored. Antonyms for a given word can likewise be produced by searching the space of semantically related words...they can be generated by choosing words with similar meanings and searching for one that differs only in one dimension...(Murphy and Andrew 1992, 18)

The significance of Murphy and Andrew's study for my research is that it offers good reasons to reject the claim of Miller and his colleagues that antonymy is simply an association between particular word forms, and it encourages us to look for a semantically based explanation for the question of why speakers choose to use two particular words together so often.

1.5.3 Antonymy and co-occurrence from a computational approach

In their paper "Redefining Antonymy: The Textual Structure of a Semantic Relation"²² Justeson and Katz (1992) set out to systematically test the co-occurrence hypothesis proposed in Charles and Miller (1989). They find that, although Charles and Miller found that substitutability per se does not apply to most antonyms (that is, in most ordinary contexts, antonyms cannot be substituted for one another), the correct generalization actually involves a combination of co-occurrence and substitution. Through their study of

²²This paper summarizes and extends earlier work described in Justeson and and Katz (1991).

antonyms in two corpora, they find that "all lexically antonymous adjectives that have a relatively high frequency of occurrence for either member of the antonym pair have a greater-than-chance occurrence rate," a finding which supports the co-occurrence hypothesis, but at the same time, they find that "these co-occurrences are constrained syntactically so that they occur in parallel and often essentially identical phrases" (Justeson and Katz 1992).

Although Justeson and Katz conclude that antonymy involves a link between two specific words, just as Miller and his colleagues claimed, they also say that a semantic component is necessary to antonymy:

We conclude that antonyms are *those semantically opposed words* that are conjoined and often opposed to one another at relatively high rates in sentences by substitution for one another in otherwise essentially identical (or parallel) phrases (italics mine). (Justeson and Katz 1992).

Justeson and Katz reached these conclusions after investigating the co-occurrence patterns of many pairs of antonyms, first in the Brown corpus (useful because it was tagged) and then in a larger, untagged corpus of 25,000,000 words. The larger corpus was used to check patterns seen in the smaller corpus but which were not statistically significant because some of the adjectives did not occur very frequently.

They began by looking at the occurrences of the antonyms studied by Deese (listed in (4) above). Of the 35 pairs they looked at,²³ 30 pairs occurred together in the Brown corpus in the same sentence at least once; if the co-occurrence was simply random, Justeson and Katz argue, they would have

²³Four of the original pairs on Deese's list were eliminated from the study because one or both of the members were not tagged as adjectives in the Brown corpus. The four pairs are; *alone/together*, *few/many*, *first/last* and *married/single*.

expected only about 9 pairs to co-occur.²⁴ They say that a more stable and relevant measure of the strength of the association between the antonyms is in terms of rate, defined as "the conditional probability that the more frequent of the antonyms occurs in a sentence given that the less frequent does so."²⁵ On average, the antonym pairs from Deese's list occurred together at a rate of once in every 14.7 opportunities.

Next, Justeson and Katz measured the co-occurrences of antonyms on a new list of 22 frequently occurring adjectives, which included pairs such as *absent/present*, *ancient/modern* and *beautiful/ugly*.²⁶ Again, they found that as a group, these antonyms co-occurred much more frequently than expected by chance. Out of the 22 pairs, 14 occurred at least once in the Brown corpus, and the overall rate of co-occurrence was once every 18.2 opportunities.

Even though most of the adjective pairs Justeson and Katz examined were found to co-occur in the Brown corpus, there were several that did not, presumably because of the small size of the corpus. For example, the pair *deep/shallow* did not co-occur, but *deep* occurred 84 times and *shallow* only 14

²⁴Justeson and Katz give actual and predicted co-occurrence rates for each individual pair, but for some of the pairs, the actual and predicted frequencies are so low that no conclusions can be drawn. For example, the pair *heavy/light* would be predicted by chance to co-occur in the Brown corpus 0.11929 times, but it actually occurs once. Since it is impossible for something to occur 0.1 times, it is hard to interpret this figure. However, for other individual pairs, the figures are more meaningful. For example, *new* and *old* were predicted to occur 10.40936 times by chance and actually co-occurred 28 times.

²⁵For example, *large* occurs 347 times in the Brown corpus and *small* occurs 504 times, so there are potentially 347 opportunities for *large* and *small* to occur in the same sentence (assuming that *large* does not occur in any sentence twice.) They actually occur together in 26 sentences, so the rate is one out of every 13.3 opportunities.

²⁶In the previous study, Justeson and Katz (1991) described how they came up with this list. First, they made a list of all the adjectives that occurred frequently in the Brown corpus but were not on Deese's list. Then they asked a lexicographer, a linguist, and an elementary school teacher to form pairs of antonyms from these adjectives.

times, so there were in fact only 14 possible opportunities for the two to co-occur. Therefore Justeson and Katz checked for co-occurrence in a larger untagged corpus of 25 million words. In the larger corpus, they found that all of the antonym pairs on both lists co-occurred at relatively high rates, e.g. *shallow* and *deep* co-occurred at a rate of one out of every 26.5 sentences.

Using the large corpus, Justeson and Katz also examined the co-occurrence patterns of 346 antonym pairs in which one of the words was morphologically derived from the other, e.g., *rational/irrational* and *contributory* and *non-contributory*. They say that overall, the adjectives in these pairs occurred so infrequently that it was not possible to compare the actual and predicted co-occurrence rates of individual pairs, but they were able to measure the co-occurrence rates for the class as a whole. They report, "[T]he seventy-four observed co-occurrences represent a highly significant excess over the 2.1 expected by chance, yielding an overall co-occurrence rate of once every 18.2 opportunities."

Justeson and Katz say that the co-occurrence rates of the frequent antonyms are high enough to create a "training effect" which causes the antonyms to be associated in the mind; in other words, they believe that the clang effect, the strong association described by Deese and later by Miller and his associates, comes about because people learn from experience which antonyms are used together. This training effect is strengthened by the fact that when the antonyms co-occur, they usually do so in a syntactic frame that highlights the contrast in meaning. For example, of the 237 cases of co-occurrences of antonyms from Deese's list, 58 cases are of the form *adjective-conjunction-adjective*, as in this example: *That was one reason she did not look*

forward to Cathy's visit, short or long. In another 49 cases, the antonyms occur in noun phrases that are identical except for the adjective, e.g., *Under normal circumstances, he had a certain bright-eyed all-American charm, with great appeal for young ladies, old ladies and dogs.*

Justeson and Katz also believe that the training effect of co-occurrence can explain why speakers have a hard time deciding whether a particular pair of words are antonyms or not. They say:

[I]t is our experience that all uncertain cases involve a rather clear semantic contrast; accordingly, the variation in strength of antonymy must be ascribed mainly to the lexical criterion [that is, to the association of word forms]. We suggest that such cases typically involve too little co-occurrence via phrasal substitution, at least in the experience of the person who is unsure about antonymy judgments. (Justeson and Katz 1992)

In other words, speakers have strong intuitions about antonym pairs that they frequently encounter, but when they are asked to judge whether two less frequently occurring words are antonyms or not, they do not have such strong intuitions. Justeson and Katz give the example of *sick/well*. These two words show a clear contrast in meaning, but *sick* also contrasts with *healthy* and *well* also contrasts with *ill*, so semantic criteria alone are not enough to determine an antonym. If a speaker has enough exposure to *sick* occurring in the same sentence with *well*, she will quickly identify them as antonyms, but if not, she may not be able to make a judgment about them. If Justeson and Katz are right in this regard, then it seems that co-occurrence goes a long way toward explaining why antonymy seems to be a "gradient phenomenon" (their words), but unfortunately, no one seems to have yet tested this hypothesis by checking the co-occurrence rates of pairs such as *sick/well* to see if they are really significantly lower than the co-occurrence rates for more easily recognized

antonym pairs.

In an earlier paper describing some of the same research, Justeson and Katz (1991) say that the lexical association learned through co-occurrence is enough to define antonymy, but in this paper, they step back from that claim and admit that some kind of semantic criteria is required too. They point out (as Murphy and Andrew 1992 also did) that antonyms are not the only kinds of words which regularly co-occur. They give several examples of other sets of co-occurring words, including the pair *mental* and *physical* and the triplet *social*, *political* and *economic*. Justeson and Katz suggest that the difference between these types of co-occurring words and antonyms is that, "Whereas antonyms oppose extreme values of a single dimension, these adjectives designate dimensions, and it is the dimensions themselves that stand in opposition." A semantic criteria must therefore be involved in antonymy in order for language users to be able to distinguish antonymic opposition from other kinds of semantic contrasts.

To summarize, like Charles and Miller, Justeson and Katz believe that co-occurrence in the same sentence leads language learners to associate antonyms in semantic memory. Unlike Charles and Miller, however, they believe that co-occurrence alone is not enough to explain the association; in particular, they pay attention to the syntactic frames in which the antonyms often occur, speculating that these frames draw attention to the antonyms and thus aid the association. Finally, they say that semantics must also be involved to some extent because speakers can distinguish antonyms from other kinds of contrasting words which also co-occur frequently. However, as always, there are some unanswered questions. As was shown earlier, Charles and Miller do

not explain why speakers chose to use antonyms together so often in the same sentences. Justeson and Katz provide a clue in their observation of the commonly used syntactic frames, but they do not follow it up, so I am left wondering what is so special about these syntactic frames and why speakers use them. Fellbaum (1995) considers this question and comes up with some reasonable explanations which are discussed in the next section.

1.5.4 Co-occurrence of antonyms across syntactic class boundaries

The studies of co-occurrence phenomena discussed so far have focused on adjectives, but in her (1995) study, Fellbaum widens the scope to other categories of words. She points out that although antonymy is widespread in the adjectival lexicon, it is also found among nouns, verbs, and adverbs, and that many times, a single concept is expressed by words of different syntactic classes, e.g., *death, die, dead, deadly* and the contrasting *life, live, alive*. She accepts the WordNet model's premise that associations between antonyms organize the adjectival lexicon, but she asks:

[D]oes the associative effect of antonymy transcend the borders of syntactic categories? If so, then there is nothing special about antonymous adjectives, other than that antonymy is more pervasive among adjectives; rather, there is something special about antonymous concepts, no matter in what form these concepts are lexicalized.
(Fellbaum 1995)

In other words, she suggests, co-occurrence may not be limited to adjectives and antonymy may in fact reflect an association between antonymous concepts rather than an association between words forms.

Fellbaum began to investigate this by testing co-occurrences among verbs and nouns in the Brown corpus. (Like many others, she finds although

the Brown corpus is smaller than she would like, the fact that it is tagged makes it convenient). She chose three pairs of concepts to study, *lose/gain*, *start/finish*, and *increase/decrease*. Each of these concepts can be expressed as either verbs or nouns.²⁷ Counting up the actual co-occurrences she found that although the overall number was small, it was still higher than would be expected by chance.²⁸

Fellbaum then expanded the study to include several more concept pairs; this time, she focused on concepts that were expressed by adjectives and/or adverbs in addition to nouns and verbs, e.g., the concept pair *dark* and *light*, expressed by *dark*, *darken*, *darkness*, *light* (adj), *light* (v) and *light* (n). She found many examples of co-occurrences across syntactic class boundaries. For example, the adjective *dark* co-occurred not only with the adjective *light*, but also with the verb and the noun forms.

Fellbaum says that the results of her study argue strongly against the view of antonym learning proposed by Justeson and Katz, that antonyms are learned through a "training effect" which results from exposure to frequent co-occurrence of words in a particular syntactic frame. Fellbaum shows that co-occurrence is not limited to adjectives, and that words of one syntactic class often co-occur with semantically contrasting words from other syntactic classes. For example, she found that the adjective *wet* occurred with the verb *dry* as well

²⁷In the case of *lose*, the noun takes a slightly different shape, *loss*, but in the other cases the morphological shape of noun and verb are the same.

²⁸For example, *lose* and *gain*(v) occurred together in 5 sentences when only 0.384 would be expected by chance, and *loss* and *gain* (n) occurred 4 times, with only 0.14 expected.

as with the adjective *dry*.²⁹ When the co-occurrences cross word class boundaries in this way, they do not occur in the kinds of syntactic frames that Justeson and Katz described. Therefore, she says, the question remains, "In the face of the demonstrated absence of syntactic clues, how does the child recognize semantically opposed words?" Fellbaum does not have an answer, but she believes that however children learn semantic relations such as antonymy, co-occurrence is not enough to explain it.

Fellbaum then goes on to consider the question of why speakers frequently refer to two opposed concepts in the same sentence, making some suggestions for future study. She notes that many of the semantic frames in which antonyms occur can accommodate nouns and verbs as well as adjectives; these include the ones listed below in (4).

(4) Some common syntactic frames of co-occurring antonyms

(Both) <i>x</i> and <i>y</i>	<i>x</i> as well as <i>y</i>
<i>x</i> and <i>y</i> alike	Neither <i>x</i> or <i>y</i>
(Either) <i>x</i> or <i>y</i>	now <i>x</i> , now <i>y</i>
from <i>x</i> , to <i>y</i>	

Fellbaum points out,

When antonyms occur in these frames, they usually constitute the most salient points at either end of a continuous scale, which expresses an attribute. Referring to the salient antonymic values of that attribute can have the effect of denoting the entire range of values, even if the antonyms may not be endpoints of the scale. (Fellbaum 1995, 295)

Thus Fellbaum is arguing that speakers use antonyms together in a sentence in order to achieve a particular rhetorical effect, that of efficiently indicating a range

²⁹She says that "the adjective *wet* occurs as often with the verb *dry* as with its antonymous adjective *dry*," but since there were only two examples of each in the Brown corpus, "as often" seems a little strong to me. Here is one of the two examples: *She leaned unconcerned against the broken porch fence, brushing and drying her wet, gilded hair in the sun.*

of possible values.

Fellbaum mentions other rhetorical effects that can be achieved by using antonyms together. One is emphasis through redundancy, using phrases such as "not x but y", "instead of x, y" and "y, not x". Humor can also be effectively achieved through using antonyms together. As she says, "Strong effects can be achieved by juxtaposing complementaries, which are incompatible," as in this example from the Brown Corpus: *How easily he could hate the lovable Irish* (Fellbaum 1995, 296).

Finally, Fellbaum notes that it is natural to use antonyms together when describing actions and events that involve a change of state. These kinds of actions and events are very common in people's experiences and so are also common topics of communication. She says,

This is reflected in the lexicon, which contains many concepts denoting physical, temporal, emotional, and other changes.³⁰ Explicit reference to one state sometimes implies the opposite or reverse state; for example, *opening* something presupposes that it was *closed*; *drying* something implies that it is *wet*. Such presuppositions are frequently overtly expressed, resulting in the co-occurrence of semantically opposed words. Other sentences explicitly refer to a change from state *A* to state *B* (Fellbaum 1995, 297).

Fellbaum gives several examples of this from the Brown corpus, including the two below.

- (5) After being closed for seven months, the Garden of the Gods club will have its gala summer opening Saturday, June 3.

Pass the iron rations please, and light another candle because it's getting dark down here...

³⁰She has a reference here to: Miller, G.A. and C. Fellbaum. 1991. Semantic networks of English. *Cognition* 41: 187-229.

By looking at the meanings of the syntactic frames in which antonyms frequently occur together, I think Fellbaum provides a very likely explanation for the co-occurrence of opposites. People are not using opposites together as a kind of reflex simply because they have heard other people use them together--they use them together because they are a very effective way of conveying a meaning or creating a rhetorical effect.

Mettinger (1994) provides further support. He identifies several frequently occurring syntactic frames and discusses their function in texts, providing many examples from a corpus of modern novels.³¹ For example, he says that the syntactic frame *X and Y* has two main functions; one is to indicate "the simultaneous validity of *X, Y*" in phrases such as *Johnson was always well up in arrivals and departures* and *...life and death are the affair of the good God*. A second function is to indicate "confrontation" in phrases such as *the old wife and the new wife* (Mettinger 1994, 47).

1.6 Directions for new research

1.6.1 Unanswered questions

The work on antonymy described so far has provided several clues to help answer my initial questions, but it has also raised new ones. This is a good point at which to take another look at those questions and the answers that have been proposed. The first questions were "What makes two words antonyms?" and "What exactly are the semantic dimensions which antonyms are said to share?" It seems that people generally agree that semantic

³¹Mettinger's study of antonymy is quite interesting, but he does not deal much with the issues that concern me in this dissertation. His main purpose is to examine the various types of dimensions (e.g. *scalar*, the type of dimension found with gradable antonyms, and *digital*, the type found with complementary opposites) using a large number of pairs of opposites.

dimensions somehow reflect attributes of things in the world, e.g., the attribute of SIZE or the attribute of SPEED, and that these attributes naturally have two main contrasting values. Antonyms are the words that label these two contrasting values.

People do not agree, however, on how particular adjectives come to be chosen as the labels for particular dimensions. In other words, they do not agree on the answer to the questions of why some words have antonyms while others have none and why some words have more than one antonym. As was discussed above, Cruse (1986), Lehrer and Lehrer (1982) and Egan (1968) focus on the meanings of the individual words involved, suggesting that subtle semantic differences, such as differences in connotative meaning or in range of application, may explain why two words which seem to contrast in meaning are not accepted as antonyms by most speakers. In contrast, the psychologists and computational linguists, Gross, Fischer and Miller (1988), Charles and Miller (1989) and Justeson and Katz (1992), view antonymy as a kind of lexical association which is to some extent independent of meaning.³² They say that contrasting concepts can be expressed by many different words (e.g., the concept of wetness expressed by *wet*, *damp*, *soaked*, *soggy*, etc.), but that antonyms have a special direct association developed through hearing the antonyms frequently used together. The reason why some words have no clear antonym is because they simply do not co-occur often enough with any semantically contrasting word for this lexical association to develop, and presumably the explanation for why some words have more than one antonymy

³²Gross, Fischer and Miller and Charles and Miller say the antonymic association is completely independent of meaning, while Justeson and Katz say that a semantic component is needed in addition to the lexical component.

must be that these words frequently occur with more than one semantically contrasting word. However, I have not found any research that directly tests this by looking at the co-occurrence patterns of words that do not have antonyms or words that have two antonyms.

The co-occurrence account of antonymy also seems to be somewhat useful in answering the questions about the clang phenomenon. The set of 39 strongly associated opposite pairs described by Deese all "sound" like very good examples of antonyms, and Justeson and Katz have shown that these pairs do in fact co-occur frequently. As was discussed above, they believe that co-occurrence can also explain why people can easily identify some contrasting pairs as antonyms but have a harder time making judgments about others.

In any case, whether or not one believes that co-occurrence can explain everything about antonymy, the co-occurrence of antonyms is a real phenomenon, and the question still remains about why speakers choose to use two particular words together so frequently in the first place. Of all the still unanswered questions about antonymy, this is the one I find most interesting and pursue in the rest of the dissertation. Charles and Miller (1989) seem to suggest that the reason is that people have a kind of conditioned response-- they often hear two words used together, and they come to associate them in their minds. But as Murphy and Andrew (1992) point out, this view of antonymy leaves out its most salient characteristic, that it is a relationship between words with contrasting meanings. It seems to me that speakers choose to use antonyms together not because it has become a habit, but because it is effective in conveying meaning or making a strong rhetorical impact of the kind described by Fellbaum (1995) and Mettinger (1994). The question then

becomes what makes two good antonyms, such as *hot /cold* and *young/old*, more effective for these uses than pairs of near-opposites, such as *hot/ freezing* and *youthful/elderly*?

As I briefly described in section 1.3 above, I think Egan's suggestions about range of scope and implications are extremely promising. Although I do not think her characterization is exactly right--I do not think there are any two words, either synonyms or antonyms, which are exactly the same in terms of range of application and connotations--I think it is likely that antonyms are very similar in terms of range of application, as well as in connotations and register. Egan picks out some of the same aspects of meaning as the linguists in section 1.2.3, such as the differences in non-propositional meaning mentioned by Cruse (1986) and the differences in distribution noted by Lehrer and Lehrer (1982). Taken together, the linguistic and lexical accounts suggest that antonyms are words which have a similarity in meaning, and I think that it is this similarity which provides opportunities for speakers to use the antonyms together in a wide range of different sentences. This in turn may explain why antonyms co-occur so frequently. The rest of this thesis is an exploration of this idea.

1.6.2 A hypothesis: Antonymy results from shared semantic range

The hypothesis that I examine here is that a word's semantic range, which reflects its range of application (in Egan's terms) as well as connotations and stylistic factors, determine whether the word will have an antonym and what that antonym is. My idea is that if two words contrast in meaning across a wide semantic range, there will be many different situations in which it will be

appropriate for a speaker to use the two words together. On the other hand, if two words only contrast in a small part of their overall range of meaning, that is, if they contrast only when they are used in a particular context, they won't be considered antonyms outside of that context.

Consider the pairs *dry/sweet* and *wet/dry*, for example. *Dry* and *sweet* contrast only when they are being used to describe wine, and this is only a small part of the overall ranges of these two adjectives. *Wet* and *dry*, in contrast, share a great deal of semantic range: they can both be used to describe weather conditions, soil conditions, the state of people and their clothing, and so on, in phrases such as *wet/dry summer*, *wet/dry sand* and *wet/dry laundry*. In fact, it is probably the case that *wet* and *dry* share so much semantic range that for most uses of *wet*, the opposite state would be best described by *dry*, and vice-versa. My claim is that *dry* and *wet* are considered antonyms even in the absence of any particular context, as in the word association tests done by Deese (1965) and the antonym tests done by Murphy and Andrew (1992), because speakers have a knowledge of the overall semantic range shared by these two words, and perhaps also because this shared semantic range allows speakers to use the antonyms together frequently, thus leading to the kind of association described by Justeson and Katz (1992). This hypothesis, that antonyms are words which share a great deal of semantic range, is pursued in the three case studies that make up the rest of this thesis.

The first case study, which makes up Chapter Two, looks at *big*, *little*, *large* and *small*. The question that arises with these four adjectives, of course, is why *large* can be paired with *small* but not *little*, while *big* can be paired with both *small* and *little* (most people seem to prefer *little*, but *small* is often listed as

an antonym for *big* in many dictionaries and lists of antonyms).³³ Charles and Miller believe that semantics cannot explain this and so they regard this case as crucial evidence that antonymy is a relationship between word forms rather than concepts. My goal is to show, through a detailed look at the four adjectives' meanings, that semantics can explain this. My study shows that *large* and *little* have almost no semantic range in common and have pronounced differences in register, while *big* shares a significant amount of semantic range with both *little* and *small*.

The second case study, Chapter Three, focuses on the distinction between antonyms and near-opposites (direct and indirect antonyms, in the terms of Gross, Fischer and Miller 1988), looking at the antonyms *wet* and *dry* and several adjectives similar in meaning, namely *damp*, *moist*, *dank*, *humid*, *arid*, and *parched*.³⁴ I show that in this case too, semantic range can explain, for example, why *humid* and *damp* are not considered antonyms of *dry*.

The final case study, Chapter Four, looks at a word with two antonyms, *happy* and its opposites *sad* and *unhappy*. Many cases of a word with two different antonyms involve a word which has clearly distinct senses (e.g., *right/wrong*, *right/left* and *old/young*, *old/new*), but in this case, *sad* and *unhappy* are similar in meaning so it is not clear that distinct senses of *happy* are involved. A careful look at the meanings of the three adjectives shows that subtle differences in *sad* and *unhappy* cause them to have somewhat different semantic ranges. The wider range of *happy* overlaps the ranges of both *sad*

³³Deese even lists *small* as the antonym of *big* in the Associational Dictionary in the appendix to his book (Deese 1965), even though in the chapter on adjectives and antonymy, he lists *big/little* and *large/small* as the antonym pairs.

³⁴As I will explain in Chapter 3, I chose these adjectives from a larger set of synonyms of *wet* and *dry* listed in WordNet.

and *unhappy* so that *happy* can contrast with both *sad* and *unhappy*.

1.7 Sources of data

In each of the case studies, my goal is to compare the meanings of the words in order to see how similar or different they really are. Since all of the case studies focus on adjectives, and since the function of adjectives is to modify nouns, a very good way to characterize adjectives' meanings is by looking at the kinds of nouns they typically modify.

One source of information I make use of in this study is definitions and usage notes in learner's dictionaries. Dictionaries designed for people learning English as a second or foreign language are more useful than regular dictionaries in that they contain more information about how to distinguish near synonyms, especially synonyms which are very commonly used, such as *big* and *large*. This information often includes specific comments about common collocational patterns. In this study, I used four recent learner's dictionaries: *Collins Cobuild English Dictionary* (abbreviated hereafter as *CCED*), *Longman Dictionary of Contemporary English, Third Edition* (abbreviated as *LDOCE*), *Oxford Advanced Learner's Dictionary, Fifth Edition* (abbreviated as *OALD*) and *Longman Language Activator* (abbreviated as *LLA*).³⁵ The entries in these dictionaries are useful as a starting point in investigating how particular adjectives compare in terms of semantic range, but they do not provide enough examples--due to space limitations, a dictionary simply cannot list all the nouns an adjective typically occurs with. It is therefore necessary to also look at data

³⁵The *Longman Language Activator* is different from the other dictionaries in that it groups words relating to the same concept together. Under the entry under *big* for example, there are groupings for "words for describing an object, building, animal, or organization etc. that is big" and words that mean "to make things bigger." There is also an alphabetical cross-listing.

from a large corpus.

For finding a wide range of examples of the nouns which an adjective typically modifies, my main source of data was a corpus of over 50 million words from 6 months of the *New York Times*. It is not a balanced corpus, but it is large enough to be useful for finding statistical patterns in adjective-noun usage. A fellow graduate student at Northwestern, John Wickberg, first ran the corpus through a simple parser to tag the corpus for part-of-speech. Then he wrote a program which could take an adjective, find the nouns which it was modifying, and calculate the strength of the association between the adjective and each noun using a measure called the *mutual information statistic*. I then used this data in making my analysis of the semantic ranges of the adjectives.

The *mutual information statistic* is introduced in Church and Hanks (1990). They give the following formula, in which $P(x)$ is the probability of one word occurring (calculated by counting the number of times a word actually occurs in a particular corpus), $P(y)$ is the probability of a second word occurring, and $I(x,y)$ is the mutual information measure (the measure of association) for the two words: $I(x,y) = \log_2 P(x,y)/P(x)P(y)$. Or as Church and Hanks describe this formula in prose:

Informally, mutual information compares the probability of observing x and y together (the joint probability) with the probabilities of observing x and y independently (chance). If there is a genuine association between x and y , then the joint probability $P(x,y)$ will be much larger than chance $P(x)P(y)$, and consequently $I(x,y) \gg 0$. [the mutual information measure] If there is no interesting relationship between x and y , then $P(x,y) \approx P(x)P(y)$, and thus $I(x,y) \approx 0$. If x and y are in complementary distribution, then $P(x,y)$ will be much less than $P(x)P(y)$, forcing $I(x,y) \ll 0$. (Church and Hanks 1990, 23)

In other words, the mutual information measures the association between two

words by comparing the number of times they actually occur together in a particular corpus to the number of times they would occur by chance. The higher the mutual information value, the stronger the association.

With a sufficiently large database, such as the *New York Times* database, it is possible to find some useful patterns using this statistic. For example, *big* is a very common word, occurring 13,914 times in the corpus, and there are many nouns which occur with *big* with a mutual information value of much higher than zero. The highest values of all, not surprisingly, are for collocations with an idiomatic meaning such as *big bang* and *big leagues*, which both have a mutual information value of more than 9. Of the 112 occurrences of *bang* in the corpus, 31 of them occur with *big* in the same noun phrase, probably right before it. Some other nouns with which *big* has a high mutual information value (higher than 6) are *chunk*, *grin*, *jump* and *loser*.

The mutual information statistic is thus a good way of identifying strong relationships between particular adjectives and nouns, and thus a good source of data for someone who wants to characterize the typical semantic range of an adjective. But this method does have some limitations. The biggest drawback is that the measure is most meaningful only with adjectives that occur frequently in the corpus, e.g., *big*, *large* (which occurred 18143 times), *wet* (685), and *dry* (2147). It is much less useful for studying less frequent adjectives such as *damp*, which occurred 43 times in this corpus, or *parched*, which occurred only 47 times. Because these adjectives occur so rarely, the chance of them occurring with any noun except the most common ones are exceptional. Thus, the combination *parched fairways* occurs only once in the corpus, but it has a high mutual information value because *fairways* is a infrequent noun, occurring

only 104 times. Thus, for the less frequent adjectives discussed in the second case study, the mutual information value does not provide a strong measure of typicality, so in order to find patterns, I supplemented the corpus data with examples from dictionaries and from electronic versions of texts of English literature available through Project Gutenberg.³⁶

Another limitation to this method is that although the large database had been tagged for part of speech (not always accurately, but on the whole, pretty well), the phrase structure of sentences were not marked. This makes it hard to identify when nouns are being modified by predicative uses of an adjective. The program written by Wickberg simply looked for an adjective followed by a head noun within a three word window (to allow for compound nouns and multiple adjectives); this could accurately pick out the noun being modified by the adjective in phrases such as *large department stores*, but it could not pick out the noun modified by the adjective in a sentence such as *Her house was quite large*. In fact, the program might identify a completely unrelated noun if there was one in the first three words of the following clause. With frequently occurring adjectives, this would not be a problem--since such a noun would probably be picked out only once, randomly, it would have a mutual information value close to 0. In this study, I considered only the nouns with which the adjective had a mutual information value of 3 or more, a value which indicated the relationship was not accidental. With the infrequent adjectives, some of the

³⁶Project Gutenberg makes available many different kinds of texts in electronic form; the texts have been selected and scanned in by volunteers and reflect a wide variety of interests, from novels to historical documents to works of mathematics. For reasons having to do with copyright laws, however, they tend to be older works that are in the public domain. For this research, I downloaded and searched about 35 different works, including novels by Lewis Carroll, Thomas Hardy, and others. More information about Project Gutenberg and the kinds of texts that are available can be found on the Internet at <<http://www.promo.net/pg/>> or by email to <hart@pobox.com>.

nouns picked out by the program occurred only once with the adjective, so I checked on any combinations that seemed unexpected to see if they were due to parsing mistakes.³⁷ In the case of the adjectives *big*, *little*, *large*, and *small*, the fact that the program only picked out prenominal (attributive) uses of the adjectives was not a problem, since I was able to determine from other sources (dictionary definitions and examples picked out of a smaller corpus) that the predicative uses and the prenominal uses of these adjectives were basically the same.

In the case of the adjectives related to *wet* and *dry*, there were a few differences between predicative and attributive uses. For example, *wet* and *dry* are often used as predicative adjectives to describe the state of people, in sentences such as *I was all wet*, but there is no attributive use corresponding to these uses because *wet* and *dry* do not usually occur before nouns which name people. Since data from the New York Times corpus did not provide information of predicative uses, I used additional examples from the electronic texts from Project Gutenberg which showed that the differences between the attributive and predicative uses were limited in scope and easy to characterize.

In the case of *happy*, *unhappy* and *sad*, the differences between the predicative and the prenominal uses was greater, and I had to compensate for the limitations of the mutual information data by finding many examples from other sources. One of these sources was the example sentences in the four learners' dictionaries, *LDOCE*, *OALD*, *CCED*, and *LLA*. The examples from all four are taken from large corpora of British and American English. Another

³⁷For example, the list of co-occurring nouns and adjectives contained the combination *humid athletes*, but this turned out to be a parsing mistake, with *humid* occurring in the predicate of a clause, modifying *air*, while *athletes* started the following clause.

main source of data was examples from the quotations in the CD-ROM version of the Oxford English Dictionary.³⁸ Throughout the discussion, I have identified which data came from the large corpus, which examples came from the *Oxford English Dictionary*, and which came from sources available through Project Gutenberg. In the case of the later, I only identify individual novels when I quote sentences from them.

³⁸I used the *Oxford English Dictionary* rather than the electronic texts in the third case study because it contained a greater variety of sentences using *happy*, *sad*, and *unhappy*, and because the examples were more recent than those in the Project Gutenberg texts.