1 Gradable Adjectives

This chapter provides an overview of the semantic properties of gradable adjectives–the core set of facts that any theory must explain–and surveys the primary approaches to the semantic analysis of gradable adjectives that have been developed in the literature, focusing on two approaches. The first, which I refer to as the "vague predicate analysis", builds on the hypothesis that gradable adjectives denote partial functions from individuals to truth values. I survey the basic claims of this type of analysis, then discuss several sets of facts which are highly problematic for it, concluding that the analysis, in its basic form, cannot be maintained. I then discuss a second account, which I refer to as the "scalar analysis", in which gradable adjectives are analyzed as expressions that denote relations between objects and abstract measures, or *degrees*, and degree constructions are analyzed as expressions which quantify over degrees. I show that this type of approach contains the machinery necessary for an explanation of the data which is problematic for the vague predicate analysis. I conclude by laying out some additional facts which are problematic for a traditional scalar analysis, focusing on the scopal properties of comparatives.

1.1 The Semantic Characteristics of Gradable Adjectives

A defining characteristic of gradable adjectives is that there is some gradient property associated with their meaning with respect to which the objects in their domains can be ordered. For example, any set of objects that have some positive linear dimension can be ordered according to how *long* the objects are or how *short* they are, and any set of objects that move can be ordered according to how *fast* or *slow* they are.[1] Some connection between gradable adjectives and ordering relations is incorporated into all approaches to their semantics; what distinguishes the two analyses that I will discuss in the sections 1.2 and 1.3

[1]The distinction between so-called "positive" adjectives like *long* and *fast* and "negative" adjectives like *short* and *slow* is discussed below in section 1.1.4, and forms the focus of chapter 3.

of this chapter is the way in which the ordering on the domain is determined, in particular, whether the ordering on the domain is presupposed and the adjective is analyzed as a function from objects in an ordered set to truth values, or whether the ordering on the domain is actually determined by the meaning of the adjective. In order to appreciate the differences between the two approaches, however, it is necessary to first review some of the crucial facts that any analysis must explain. One important set of facts was discussed in the introduction: the presence of gradable adjectives in comparatives and other degree constructions. The goal of this section is to introduce several additional empirical domains that provide important insight into the semantic characteristics of gradable adjectives.

1.1.1 Vagueness

Sentences containing adjectives are inherently vague; (1), for example, may be judged true in one context and false in another.

(1) The Mars Pathfinder mission is expensive.

In a context in which the discussion includes all objects that have some cost associated with them, (1)would most likely be judged true, since the cost of sending a spacecraft to Mars is far greater than the cost of most things (e.g., nails, dog food, a used Volvo, etc.). If the context is such that only missions involving interplanetary exploration are salient, however, then (1) would be judged false, since a unique characteristic of the Mars Pathfinder mission was its low cost compared to other projects involving the exploration of outer space.

This discussion brings into focus an important aspect of the vagueness of gradable adjectives: determining the truth of a sentence of the form *x is* φ (where φ is a gradable adjective in its absolute form) involves a judgment of whether x "counts as" φ in the context of utterance. The problem of resolving the vagueness of a gradable adjective, then, can be viewed as the problem of answering the question *does x count as* φ *in context c?* Although there may be many different ways to construct an algorithm for answering this question, two approaches have predominated in research on the semantics of gradable adjectives. In the following paragraphs, I will present an informal outline of these two approaches, returning to a more formal discussion of the same issues in Sections 1.2 and 1.3.

The first approach, which I will refer to as the "vague predicate analysis" (see McConnell-Ginet 1973, Kamp 1975, Klein 1980, 1982, 1991, van Benthem 1983, Larson 1988a, and Sánchez-Valencia 1995), starts from the assumption that gradable adjectives are of the same semantic type as non-gradable adjectives and other predicates: they denote functions from objects to truth values. What distinguishes gradable adjectives from other predicative expressions is that the domains of the former are partially ordered with respect to some property that permits gradation, such as *cost, temperature, height*, or *brightness*. On this view, the observation that objects can be ordered according to the amount to which they possess some property is interpreted as basic principle (see Sapir 1941 for relevant discussion), and the meaning of a gradable adjective is built on top of it. Specifically, a gradable adjective φ is analyzed as a function that induces a partitioning on a partially ordered set into objects ordered above some point and objects below that point: for objects ordered towards the upper end of the set, *x is* φ is true, and for objects ordered towards the lower end, *x is* φ is false.[2]

In this type of approach, the problem of vagueness can be characterized as the problem of determining how the domain of a gradable adjective should be partitioned in a particular context. One way to go about solving this problem is to assume a very general algorithm whereby a gradable adjective partitions any partially ordered set according to some "norm value", and to allow for the possibility that in different contexts, instead of applying the adjective to its entire domain, only a subset of the domain is considered.[3] Specifically,

[2]Klein 1980, 1982 argues that gradable adjectives should actually be analyzed as *partial* functions, allowing for the possibility that for some objects in the domain of φ, *x is* φ is undefined (see also Kamp 1975), resulting in a three-way partitioning of the domain. I will return to this issue in section 1.2.

[3]How exactly the norm value is determined in this type of analysis is not a question that I will attempt to answer here, though I will return to this question in the context of a different analysis in chapter 2. See Siegel 1979 for a general survey of different approaches to this question; see also Bierwisch 1989 for related discussion of the notion of "norm".

when evaluating a sentence of the form *x is* φ in a context *c,* attention is restricted to a subset of the domain of φ that contains only objects that are deemed to be "like *x*" in some relevant sense in *c* (assuming that the relation "is like *x*" is reflexive, this subset will always include *x*), and then checking to see whether the partitioning of the subset by φ is such that *x is* φ is true.

Following Klein, I will refer to this contextually relevant subset as a *comparison class*. Intuitively, a comparison class is a subset of the domain of a gradable adjective that contains just those objects that are determined to be relevant in a particular context of utterance, in particular, those objects that are similar to *x* in some appropriate respect. The intuition underlying this type of approach is that in order to make a precise judgment about whether an object "counts as" φ, it is first necessary to focus attention on a subset of the domain that contains objects that are in some way similar to *x*, and then check to see whether *x* falls "at one end of the other" of the ordered subset. The basic idea can be illustrated by considering example (1). Assume that the domain of the adjective *expensive* is the set of entities that can have some cost value. Among this set are the objects in (2), which are ordered according to increasing cost.

(2) $D_{expensive}$ = {... a nail ... a bag of dog food ... a Hank Mobley album ... a copy of *Stricture in Feature Geometry* ... a dinner at Chez Panisse ... a new BMW ... a house in San Francisco ... the Mars Pathfinder mission ... a manned mission to Mars ...}

If all of the objects in the domain of *expensive* must be considered when evaluating the truth of (1), then it is clear that (1) should be true, since the cost of the Mars Pathfinder mission is greater than the cost of most things. If, however, the context is such that only projects in the space program are relevant, then the comparison class would consist of the subset of $D_{expensive}$ illustrated in (3).

(3) {the Mars Pathfinder mission ... a 15 day space shuttle mission ... a mission to the

moon ... the international space station ... a manned mission to Mars}

In this context, (1) would be false, because the Mars Pathfinder mission falls at the low end of the ordering. Other contexts might give rise to comparison classes in which the Mars Pathfinder mission falls at the upper end of the ordered set (e.g., contexts in which the comparison class consists of expeditions involving 6-wheeled vehicles), in which case (1) would again be true.

The initial assumption that the domain of a gradable adjective has an inherent ordering imposed upon it is crucial to the vague predicate analysis, since the truth or falsity of a sentence of the form *x is φ* is determined by the position of *x* in the ordered set (whether it is ordered at the upper end or whether it is ordered at the lower end). Moreover, the inherent ordering on the domain plays an important role in the analysis of vagueness outlined here, as well, since it is necessary that any comparison class constructed from an ordered set *S* preserves the ordering on *S*. If the ordering on the domain was not inherent, but could change from context to context, then a subset of the domain of *expensive* as presented in (2) with the ordering indicated in (1) would be a possible comparison class for (1), with the result that (5) would be false and (6) true *in the same context*.

(4) {a manned mission to Mars ... the international space station ... a mission to the moon ... a 15 day space shuttle mission ... the Mars Pathfinder project}

(5) The Mars Pathfinder mission is expensive.

(6) A manned mission to Mars is expensive.

This would be an unacceptable result: there is a clear intuition that if the basic ordering on the domain of *expensive* is as in (2), then any context in which (6) is true should also be one in which (5) is true. In order to avoid this problem, Klein (1982:126) stipulates that the ordering on a comparison class must preserve the initial ordering on the domain of the adjective, pointing out that this is not an unjustified assumption; rather, it is "fundamental

to the expression of ordering relations in natural language." This claim raises the following question, however: should a principle like this be made to follow more directly from the meaning of a gradable adjective itself? More generally, should the ordering on the domain of a gradable adjective be viewed as a primitive, or should it be determined in some way by the meaning of the adjective itself? The analysis that I have outlined here takes the former position; in the following paragraphs, I will sketch an alternative approach that makes the latter assumption.

The second approach to the problem of vagueness, first articulated in Cresswell 1976 (see also Seuren 1973) but since incorporated into many analyses of the semantics of gradable adjectives (see e.g., Hellan 1981, Hoeksema 1983, von Stechow 1984a, Heim 1985, Lerner & Pinkal 1992, 1995, Moltmann 1992a, Gawron 1995, Rullmann 1995, Hendriks 1995), provides a means of answering the question *does x count as φ in c?* by constructing an abstract representation of measurement and defining the interpretation of a gradable adjective in terms of this representation.[4] This abstract representation, or *scale*, can be construed as a set of points ordered by a relation ≤, where each point represents a measure or *degree* of "*φ*-ness". The introduction of scales and degrees into the ontology makes it possible to analyze gradable adjectives as relational expressions; specifically, as expressions whose semantic function is to establish a relation between objects in its domain and degrees on the scale. A more general consequence of defining the interpretation of an adjective in terms of a scale is that the ordering on the domain of a gradable adjective is determined by a semantic property of the adjective itself: by establishing a relation between objects and points in a totally ordered set, the adjective imposes a partial order on its domain.

For illustration, consider the domain of the adjective *expensive*, repeated below as (7).

[4]In Cresswell's analysis, scales are actually constructed out of equivalence classes of objects partially ordered according to some gradable property; in this respect, the analysis starts from the same assumption as the approach discussed above: that objects can be ordered according to the degree to which they possess some property, independently of the meaning of particular adjectives. In the discussion that follows, however, I will treat scales as abstract objects with respect to which the ordering on a set is determined (see Bierwisch 1989). I will go into the formalization of scales in more detail in section 1.3, and, in particular, in chapter 3.

(7) D = {a nail ... a bag of dog food ... a Hank Mobley album ... a copy of *Stricture in Feature Geometry* ... dinner at Chez Panisse ... a new BMW ... a house in San Francisco ... the Mars Pathfinder project ... a 15 day space shuttle mission ... a mission to the moon ... the international space station ... sending people to Mars}

In the vague predicate analysis outlined above, the ordering represented in (7) is assumed to be an inherent property of the domain of the adjective. In the alternative "scalar" analysis, however, the domain of the adjective is unordered, but an ordering corresponding to the one illustrated in (7) can be derived as a consequence of the fact that the adjective *expensive* establishes relations between the objects in $D_{expensive}$ and elements in a totally ordered set of points, i.e., degrees on a scale of *expensiveness*.

The characterization of gradable adjectives as relational expressions supports an alternative approach to the interpretation of vague sentences like (1). Specifically, a sentence of the form *x is φ* is taken to mean *x is at least as φ as d*, where d is a degree on the scale associated with φ that identifies a "standard" of φ-ness. Intuitively, a standard-denoting degree is a value that provides a means of separating those objects for which the statement *x is φ* is true from those objects for which *x is φ* is false, in some context. The structure of scales–specifically, the fact that they are defined as totally ordered sets–ensures that the relative ordering of a standard-denoting degree and a degree which corresponds to the measure of an object's "adjectiveness" can always be determined.

For example, a sentence like (1), on this view, is assigned an interpretation that can be paraphrased as (8), which is true just in case the degree that indicates the expensiveness of the Mars Pathfinder mission is at least as great as the standard value (I will return to a more formal discussion of this approach in section 1.3).

(8) The Mars Pathfinder mission is at least as expensive as a standard of expensiveness.

Within this type of analysis, the problem of vagueness can be cast as the problem of

determining the actual value of the standard in the context of utterance. The standard assumption is that the standard value is set indexically, and that its value may be determined by the a contextually relevant comparison class (see Cresswell 1976, von Stechow 1984a, and, in particular, Bierwisch 1989 for discussion).[5] For example, assume that in a context in which the comparison class is determined to be projects in the space program, as in (3) above, the relation between the projections of the objects in the comparison class onto the scale of expensiveness may (i.e., their "degrees of expensiveness") stand in relation to the standard degree d_{stnd} as shown in (9).

(9) *expensive*: —— $d_{Pathfinder}$ —— $d_{Shuttle}$ —— d_{stnd} —— d_{Moon} —— $d_{Station}$ — $d_{People\ to\ Mars}$ ——→

In this context, (8) is true, because $d_{Pathfinder}$–the degree to which the Pathfinder mission is expensive–is ordered below d_{stnd}. In an alternative context, however, in which the comparison class were such that the standard value were to shift to a point below $d_{Pathfinder}$, (8) would be false. What the scalar analysis "gets for free" is the preservation of the ordering on the domain, because a change in comparison class, with a concomitant change in the value of the standard, does not affect the overall ordering of the degrees on the scale. Since the scale determines the ordering on the domain of the adjective, this ordering remains constant, regardless of a shift in comparison class.

An important similarity between the two approaches to vagueness discussed here is that the context-dependence of vague sentences like (1) is ultimately explained in the same way: in terms of comparison classes. In order to know whether a sentence of the form *x isφ* is true in a context *c*–whether *x* "counts as" *φ* in *c*–it is first necessary to determine what subset of the domain of the gradable adjective that is taken to be relevant in the context. This subset–the comparison class–is then used as the basis for evaluating the truth of the

[5] In chapter 2, I will undertake a detailed discussion of how the makeup of the comparison class affects the actual value of the standard. For the moment, we may make the simplifying assumption (as above) that the standard value represents a "norm" for a particular comparison class (cf. Bartsch and Vennemann 1972, Cresswell 1976, von Stechow 1984a, Bierwisch 1989).

sentence. In the first account, the comparison class introduces the set that is partitioned by the adjective; in the second account, the comparison class is used as the basis for fixing the value of the standard. In both cases, when the comparison class is changed, the truth of the original sentence may be affected: either the partitioning induced by the adjective may change, or the standard value may be shifted accordingly.

Despite this similarity, the two analysis outlined here differ in a fundamental way. Specifically, they make very different claims about the relation between the meaning of a gradable adjective and the ordering on its domain. In the vague predicate analysis, the ordering on the domain is assumed to be inherent. This assumption not only permits a straightforward semantic analysis of gradable adjectives as predicative expressions, it also provides justification for the assumption the construction of a comparison class always preserves the ordering on the domain. In contrast, the scalar analysis derives the ordering on the objects in the domain of a gradable adjective from the meaning of the adjective itself, which establishes a relation between domain objects and degrees on a scale (i.e., points in a totally ordered set). This result does not come without a cost, however. Although the scalar approach derives the ordering on the domain, it gives up the analysis of gradable adjectives as simple predicates, treating them instead as relational expressions. In addition, it requires the introduction of abstract objects into the ontology, namely scales and degrees.

The latter difference is of primary importance, as it introduces a potential basis for making an empirical distinction between the two analyses sketched here. If scales and degrees do play a role in the interpretation of gradable adjectives, then it should be possible to show that there are facts which can be explained only if scales and degrees are part of the ontology; such facts would then constitute an argument for a scalar approach. One of the goals of this thesis is to make exactly this argument. In section 1.2, I will introduce several sets of facts which are problematic for the analysis of gradable adjectives as simple predicates, and in section 1.3, I will show that these facts can be straightforwardly explained if scales and degrees are part of the ontology.[6] Before moving on to this discussion, though, some additional semantic characteristics of gradable adjectives that will play crucial roles in the argument will be introduced.

1.1.2 Indeterminacy and the Dimensional Parameter

In most cases, the resolution of vagueness–the judgment of whether an object x "counts as" φ–can be accomplished as described above: either by restricting attention to a particular comparison class, or by determining an appropriate standard. Both of these operations presuppose that the ordering associated with the adjective (either on the domain or with respect to the scale) is determinate, however, since it is with respect to this ordering that the ultimate judgment is made. For many adjectives, however, this presupposition is not met. Consider, for example, the following sentences:

(10) Richard is smart.

(11) *The Devils* is a slow book.

(12) William is liberal.

The truth of a sentence like (10) is *indeterminate* in a way that is different from that of a typical vague sentence such as *Richard is tall*. A particular individual might be considered smart in the role of, for example, a political advisor, but decidedly *not* smart when it comes to social behavior and discreetness. As a result, the truth or falsity of a general statement like (10) is unclear, raising the following question: *smart in what sense*? (11) and (12) are similarly indeterminate. A book might be exciting and engaging, but nevertheless be slow to read due to the complexity of its characters and language. Similarly, an individual might be judged liberal with respect to some issues (e.g., health care, affirmative action); but with respect to other issues (e.g., welfare, immigration), the same individual might not be.

One way to approach the problem of indeterminacy would be to assume that it is a

[6] In chapter 3, I will address the question of the actual structure of scales and degrees, arguing that degrees should be formulated as intervals on a scale or "extents", rather than points, as standardly assumed.

kind of vagueness, arising from a difficulty in some contexts of determining an appropriate comparison class. Although this might be true of (12), examples like (10) and (11) call this characterization of indeterminacy into question. What is at issue in these sentences is not the content of the comparison class, but rather the actual ordering on the domain of the adjective. Adjectives like *smart*, *slow* and *liberal* have a wider range of interpretations than an adjective like *tall*, in that they permit different orderings on their domains in different contexts of use. For example, *smart* may involve an ordering according to political or strategic skill, or it may be associated with an ordering according to more general notions of social behavior and personal conduct. In the former case, (10) might be judged true; in the latter case (10) might be judged false. What is important to note is that even if the comparison class remains constant–the set of political consultants, for example–the truth value of a sentence like (10) can still vary depending on which of these two interpretations of *smart* is chosen.[7]

Indeterminacy is a characteristic of a large number of gradable adjectives in English, which McConnell-Ginet (1973) and Kamp (1975) refer to as the NON-LINEAR adjectives (see also Klein 1980). A defining characteristic of non-linear adjectives is that comparative constructions in which they appear do not have definite truth values, in contrast to comparative constructions in which otherwise vague adjectives appear; indeed, this characteristic explicitly distinguishes indeterminacy from the type of vagueness discussed in section 1.1.1. For example, (13) has the same status as (10)–we cannot evaluate the truth of this sentence without first knowing the sense in which *smart* is used–i.e., what the criteria for "smartness" are. In contrast, (14) can be evaluated simply by determining the costs of the different missions.

(13) Richard is smarter than George.

(14) The Mars Pathfinder mission was less expensive than the Viking missions.

[7]The interpretation of *slow*, particularly in attributive contexts, is similar (cf. *a slow book*, *a slow car*, *a slow student*). See Pustejovsky and Boguraev 1993 for relevant discussion.

What the facts discussed here indicate is that the relation between an adjective and a particular ordering relation is not one to one: some gradable adjectives may be associated with more than one ordering on a domain. For example, consider the adjective *large* in the context of cities. Cities can be ordered according to different aspects of largeness, such as volume, population, or even size of the bureaucracy (see Klein 1991:686 and Cresswell 1976:270-271 for discussion); not surprisingly, the NP *a large city* is ambiguous in at least these three ways. In the discussion that follows, I will refer to the aspect according to which objects in a set are ordered as a DIMENSION (cf. Bierwisch's (1989) notion of ASPECT). The underlying idea is that a dimension corresponds to a property that permits grading, i.e., a property such as *temperature, monetary cost, physical size, social grace, skill at manipulating people,* etc., that can be used as a basis for imposing an ordering on a set of objects. Dimensions play a fundamental role in the analysis of gradable adjectives, as they determine the actual ordering of the objects in a gradable adjective's domain.

To make things concrete, I will assume that every gradable adjective is associated with a *dimensional parameter*, which identifies the dimension that a particular adjective is associated with, and so determines how the objects in an adjective's domain should be ordered. How this is accomplished differs depending on which of the two approaches discussed in the previous section is adopted, as the implementation of dimensions differs in the two accounts. In the vague predicate analysis, the dimensional parameter specifies the dimension according to which the initial ordering on the domain of the adjective is constructed; in effect, it identifies which of many possible orderings on the set of objects that satisfy the selectional restrictions of the adjective should be used to build the (partially ordered) domain of the function denoted by the adjective. Crucially, since different dimensions may determine different orderings, the domains of adjectives with different dimensional parameters may have distinct orderings, even if they contain the same objects.

In the scalar analysis, the dimensional parameter plays a slightly different role, but ultimately has the same effect. Specifically, the dimensional parameter identifies the scale onto which the adjective maps the objects in its domain. The underlying idea is that dimensions distinguish one scale from another: a scale along a dimension of *temperature* and a scale along a dimension of *brightness* are different objects. The consequence of this distinction is that adjectives with different dimensional parameters may impose different orderings even if their domains are the same, since the mapping between a set of objects and one scale need not be the same as the mapping between the same set of objects and a different scale.[8]

Given these basic assumptions, the fact that non-linear adjectives support more than one ordering on their domains can be taken as evidence that they are underspecified for their dimensional parameter. In other words, the difference between adjectives like *smart* and *tall* is that the dimensional parameter of the former may take on different values in different contexts, while the dimensional parameter of the latter is fixed. If this is correct, then indeterminacy is a type of ambiguity, rather than a type of vagueness: it is the problem of determining in some context of use what the actual value of the dimensional parameter of a non-linear adjective is. Once the dimensional parameter of a non-linear adjective is fixed, however, sentences like those discussed here can be evaluated in the same way as sentences with other gradable adjectives. The fact that a sentence like (10) might be true on one interpretation of *smart* and false on the other, even with respect to the same comparison class, follows from the fact that the two interpretations of *smart* have different dimensional parameters, which in turn introduce different orderings on the adjective's domain.

1.1.3 Incommensurability

It is often the case that the object of which a non-linear gradable adjective is predicated indicates which of several dimensions associated with the adjective is appropriate for a

[8]A second consequence of this distinction is that degrees on different scales are elements of different ordered sets. This aspect of the scalar analysis will play a crucial role in the discussion of *incommensurability*, a phenomenon that will be introduced in the next section and used as a basis for distinguishing between the vague predicate and scalar analyses in sections 1.2 and 1.3.

particular sentence. For example, the adjective *long* is associated with at least two dimensions: one that supports an ordering according to *linear extent*, and one that supports an ordering by *temporal duration*. Because of the sortal restrictions associated with the corresponding interpretations of the adjective, neither (15) nor (16) is indeterminate: the former has the *temporal duration* interpretation; the latter the *linear extent* reading.

(15) The class was long.

(16) The table is long.

An important fact is that when an adjective like *long* appears in a comparative construction, the compared objects must be ordered along the same dimension. If they are not, as in (17) and (18), the comparative is anomalous (see Hale 1970).

(17) #The class was longer than this table is.

(18) #*The Devils* isn't as slow as the people in this class.

The anomaly of (17) and (18) is actually a specific instance of a more general phenomenon that will play a crucial role in the discussion to follow: *incommensurability* (see Klein 1991:686 for discussion). Incommensurability is illustrated by (19) and (20), which show that adjectives that are associated with different dimensions are anomalous in comparative constructions.

(19) #Larry is more tired than Michael is clever.

(20) #My copy of *The Brothers Karamazov* is heavier than my copy of *The Idiot* is old.

That the anomaly of these sentences is due to a mismatch of dimension, rather than a general constraint prohibiting different adjectives from appearing in comparative constructions, is shown by (21)-(22).

(21) Most boats are longer than they are wide.

(22) Our Norfolk Island Pine is almost as tall as the bedroom ceiling is high.

Although the comparatives in these examples are constructed out of different adjectives, the pairs of adjectives arguably have the same or very similar dimensional parameters, as they all introduce orderings according to different aspects of the same basic property: some notion of "linear extent".[9]

The contrast between (19)-(20) and (17)-(18) on the one hand, and (21)-(22) on the other, suggests the descriptive generalization stated in (23).

(23) A comparative construction is semantically well-formed only if the compared adjectives have the same dimensional parameter.

A plausible explanation for this generalization is that a necessary condition for comparison is that the compared objects be ordered along the same dimension. Ideally, this condition should follow as a general consequence of the semantic analysis of gradable adjectives and comparative constructions; it should not have to be stipulated. This requirement suggests a potential point of difference between the vague predicate and scalar analyses of gradable adjectives. Since the role of the dimensional parameter of a gradable adjective differs in the two accounts–in the former, it identifies how the objects in the domain of the adjective should be ordered; in the latter, it identifies the scale onto which the objects in the domain of the adjective should be mapped–the two analysis should differ in their approaches to the

[9]It should be observed that while sentences like (17)-(20) are anomalous in a "basic" interpretation, as shown by the clear contrast between these examples and (21)-(22), sentences of this type are often used with a sense of irony or humor. In these contexts, a sentence of the form *"x is more/as φ than/as y is ψ"*, where φ and ψ are otherwise incommensurable adjectives, takes on a "metalinguistic" interpretation: roughly, "it's more/as appropriate to say that '*x* is φ' than/as it is to say that '*y* is ψ'" ((18), for example, clearly has this kind of interpretation; see McCawley 1988 and Klein 1991 for some discussion of this phenomenon). The possibility of a metalinguistic interpretation does not eliminate the need for an account of incommensurability, however, since the contrast between sentences like (17)-(20) and (21)-(22) must still be explained.

problem of incommensurability. In sections 1.2 and 1.3, I will take a closer look at these issues, arguing that only an analysis that introduces scales and degrees into the ontology provides an explanatorily adequate account of this phenomenon.

1.1.4 Polarity

The final semantic characteristic of gradable adjectives that I will consider here is *polarity* (see Seuren 1978, Rusiecki 1985, Bierwisch 1989, Sánchez-Valencia 1994, and H. Klein 1996 for relevant discussion). Many (though not all) gradable adjectives come in antonymous pairs, such as *tall/short*, *safe/dangerous*, *sharp/dull*, *liberal/conservative*, and so forth. At some fundamental level, the members of such pairs provide the same kind of information about an object: they provide a characterization of an object according to some property that permits grading. What distinguishes the "positive" and "negative" members of these pairs is that they do this from different (and in some sense *complementary*) perspectives. For example, both *tall* and *short* are used to make claims about the height an object has, but the information conveyed by a sentence like *Carmen is tall* is qualitatively different from that conveyed by *Mike is short*. This difference in perspective has several important empirical consequences that are relevant to the overall semantic analysis of gradable adjectives.

1.1.4.1 Montonicity Properties

As shown by (24)-(29), antonymous adjectives differ with respect to the licensing of negative polarity items (NPIs): the negative member of the pair licenses NPIs in a clausal complement of the adjective; the positive member does not (Seuren 1978, Ladusaw 1979, Linebarger 1980, Sanchez-Valencia 1996).

(24) It's difficult for Tim to admit that he has *ever* been wrong.

(25) *It's easy for Tim to admit that he has *ever* been wrong.

(26) It's sad that you have to talk to *any* of these people *at all*.

(27) *It's great that you have to talk to *any* of these people *at all*.

(28) It would be foolish of her to *even bother to lift a finger* to help.

(29) *It would be clever of her to *even bother to lift a finger* to help.

Similarly, (30)-(33) show that negative adjectives license downward entailments in clausal complements, while positives license upward entailments.

(30) It's dangerous to drive in Rome. ⇒ It's dangerous to drive fast in Rome.

(31) It's safe to drive in Des Moines. ⇐ It's safe to drive fast in Des Moines.

(32) It's strange to see Frances playing electric guitar. ⇒ It's strange to see Frances playing electric guitar poorly.

(33) It's common to see Frances playing electric guitar. ⇐ It's common to see Frances playing electric guitar poorly.

The conclusion to be drawn from these facts is that negative adjectives generate monotone decreasing contexts, while positives generate monotone increasing contexts. A plausible hypothesis, then, is that positive and negative adjectives are associated with inverse ordering relations. The validity of statements like (34)-(35) provides initial support for this idea.

(34) Venus is brighter than Mars if and only if Mars is dimmer than Venus.

(35) Driving in Rome is more dangerous than driving in Los Angeles if and only if driving in Los Angeles is safer than driving in Rome.

More generally, the facts discussed here suggest that gradable adjectives have logical properties that are connected to the way in which the ordering on their domains is introduced. As a result, adjectival polarity provides an empirical domain for exploring the

relation between gradable adjectives and ordering relations. The important questions are: how is adjectival polarity represented in the lexical semantics of pairs of adjectives like *tall* and *short*, *safe* and *dangerous*, and so on, and how does the representation of polarity explain the logical properties of the adjectives discussed above: licensing of negative polarity items, entailments, and the validity of statements like (34) and (35)? These questions will be addressed in chapter 3.

1.1.4.2 Cross-Polar Anomaly

Sentences such as (36)-(39) show that comparatives constructed out of positive and negative pairs of adjectives are anomalous, a phenomenon that I will refer to as *cross-polar anomaly* (see Hale 1970 and Bierwisch 1989 for discussion of similar facts).

(36) #Mike is shorter than Carmen is tall.

(37) #*The Brothers Karamazov* is longer than *The Idiot* is short.

(38) #The Tenderloin is dirtier than Pacific Heights is clean.

(39) #A Volvo is safer than a Fiat is dangerous.

(36)-(39) contrast with structurally similar sentences in which the polarity of the adjectives is the same, such as (40) and (41), which are perfectly interpretable (cf. (21)-(22)).

(40) The desk is longer than the table is wide.

(41) Luckily, the ficus is shorter than the ceiling is low, so it'll fit in the room.

An initially promising hypothesis is that cross-polar anomaly is a type of incommensurability; i.e., (36)-(39) are anomalous for the same reason as sentences like (17)-(20), discussed in section 1.1.3: positive and negative adjectives are associated with different dimensional parameters, and so are incomparable. An analysis of this sort faces two important difficulties, however. The first comes from tautologies like (34) and (35). As the discussion of indeterminacy in section 1.1.2 showed, different dimensions may introduce different orderings on the same domain. As a result, it is not possible to make "cross-dimensional" inferences: given two dimensions d_1 and d_2 that define partial orderings on a set A, the fact that two objects a and b in A stand in a particular ordering relation with respect to d_1 does not tell us anything about the relative ordering of a and b with respect to d_2. The importance of examples like (34) and (35) is that they show that there is a non-arbitrary relation between positive and negative pairs of adjectives: the ordering relation associated with the latter is the inverse of the ordering associated with the former. Without additional stipulations, an analysis of cross-polar anomaly that asserts that positive and negative adjectives have different dimensional parameters would lose this crucial relation.

The second difficulty facing this account of cross-polar anomaly that it would conflict with the basic characterization of a dimension as an ordering with respect to a property that permits grading. As noted in the introduction to this section, there is a strong intuition that antonymous adjectives provide complementary perspectives on how an object is characterized with respect to the same gradable property, e.g. a dimension of *height* for the adjectives *tall* and *short*. One way to account for this intuition would be to assume that antonymous adjectives introduce inverse orderings along *the same* dimension; indeed, it is this assumption that provides the basis for an explanation of the validity of examples like (34) and (35). If positive and negative adjectives have different dimensional parameters, however–an assumption required by an explanation in terms of incommensurability–then this explanation would be unavailable, and the underlying intuition would remain unexplained.

Despite these difficulties, the intuition that cross-polar anomaly is a kind of incommensurability remains. The challenge facing an analysis that seeks to explain the anomaly of sentences like (36)-(39) *and* sentences like (17)-(20) in terms of the same underlying principles is to do so in a way that maintains the assumption that antonymous adjectives have the same dimensional parameter. In section 1.3 and in more detail in chapter 3, I will show that a positive result of a scalar analysis of gradable adjectives is that,

given an appropriate formalization of degrees, both cross-polar anomaly and incommensurability can be explained in terms of a failure of the comparison relation to be defined for the compared degrees.

1.1.4.3 Comparison of Deviation

A set of facts that bears directly on the analysis of cross-polar anomaly consists of sentences which are superficially similar to examples like (36)-(39), in that they involve comparative (and equative) constructions constructed out of positive and negative pairs of adjectives but are not anomalous. These sentences, which I will refer to as *comparison of deviation* constructions, are exemplified by (42)-(45).

(42) Robert is as short as William is tall.
(43) Alex is as slim now as he was obese before.
(44) Frances is as reticent as Hilary is long-winded.
(45) It's more difficult to surf Maverick's than it is easy to surf Steamer Lane.

Comparison of deviation constructions differ semantically from standard comparatives and equatives in two important ways.[10] First, they differ in terms of basic interpretation. Roughly speaking, sentences like (42)-(45) compare the relative extents to which a pair of objects differ from some "standard" associated with the relevant adjectives. This is clearly illustrated by an equative construction like (42), which can only mean that the extent to which Robert exceeds some standard of shortness is (relatively) the same as the

[10]A third difference between comparison of deviation constructions and standard comparatives is that they do not license morphological incorporation of the adjective and the comparative morpheme. This is clearly illustrated by the minimal pair (i)-(ii). (i) has a comparison of deviation interpretation, but (ii) is an example of cross-polar anomaly.

(i) San Francisco Bay is more shallow than Monterey Bay is deep.
(ii) #San Francisco Bay is shallower than Monterey Bay is deep.

extent to which William exceeds some standard of tallness. (42) *cannot* mean that Robert and William are equal in height.[11]

Second, unlike standard comparatives, comparison of deviation constructions entail that the property predicated of the compared objects is true in the absolute sense. For example, (45) entails that surfing Maverick's is difficult and surfing Steamer Lane is easy, as shown by (46), which is contradictory.

(46) It's more difficult to surf Maverick's than it is easy to surf Steamer Lane, though Maverick's is quite easy.

In contrast, (47) does not entail that surfing either location is difficult, although this information is conveyed as a cancelable implicature, as shown by (48).

(47) It's more difficult to surf Maverick's than it is to surf Steamer Lane.

(48) It's more difficult to surf Maverick's than it is to surf Steamer Lane, though they're both quite easy.

The challenge presented by comparison of deviation is to construct an analysis that both accounts for the inferences associated with sentences like (42)-(45) *and* supports an explanation of cross-polar anomaly. In sections 1.2 and 1.3, I will show that only a scalar analysis achieves this result.

1.1.5 Summary

This section provided an overview of the empirical domain that a theory of the semantics of gradable adjectives and degree constructions must explain, and it gave an informal

[11]In this way, the interpretation of (42) differs from that of a more typical equative, such as (i), which asserts that the length of Robert's feet and the width of William's feet are the same.

(i) Robert's feet are as long as William's feet are wide.

introduction to two approaches to these issues: the vague predicate analysis and the scalar analysis. The vague predicate analysis assumes a partial ordering on the domain of a gradable adjective and constructs a semantic analysis of gradable adjectives as partial functions from objects to truth values; the scalar analysis extends the ontology to include abstract representations of measurement, or "scales", and characterizes the interpretation of gradable adjectives in terms of such objects. In the next two sections, I will focus in more detail on these two analyses. Although the two are very similar in terms of their empirical coverage, I will show that three of the phenomena discussed in this section–incommensurability, cross-polar anomaly, and comparison of deviation–provide evidence that the interpretation of gradable adjectives should be formalized in terms of scales and degrees.

1.2 The Vague Predicate Analysis

The first approach to the problem of vagueness discussed in section 1.1.1, which can be found in the work of McConnell-Ginet 1973, Kamp 1975, Klein 1980, Klein 1982, van Benthem 1983, Larson 1988a, and Sanchez-Valencia 1994 (see also Lewis 1973), analyzes gradable adjectives as predicates whose domains are partially ordered according to some dimension. In the following sections, I will examine the basic assumptions of this type of approach in more detail, focusing on the versions articulated in Klein 1980, 1982. I will show first how this type of analysis explains the interpretation of gradable adjectives in the absolute form, then how it explains the interpretation of more complex degree constructions, focusing on the analysis of comparatives. Finally, I will discuss several problematic sets of data, concluding that the analysis in its basic form cannot be maintained.

1.2.1 Overview

As observed in section 1.1.1, the vague predicate analysis starts from the assumption that gradable adjectives are of the same semantic type as non-gradable adjectives: they denote functions from objects to truth values. Gradable adjectives are distinguished from non-gradable adjectives (and other predicative expressions) in that their domains are partially ordered according to some gradable property, such as *cost, temperature, height*, or *brightness*; in section 1.1.2, I assumed that the ordering on a gradable adjective's domain is determined by its dimensional parameter. Klein (1980, 1982) (building on Kamp 1975) makes a second distinction between gradable and non-gradable adjectives: the latter always denote complete functions from individuals to truth values, but the former can denote *partial* functions from individuals to truth values. In other words, non-gradable adjectives like *hexagonal* and *Croatian* always denote functions that return a value in {0,1} when applied to objects in their domains, but gradable adjectives like *dense, bright*, and *shallow* can denote functions that return 0, 1 or no value at all when applied to objects in their domains.

The interpretation of a proposition with a gradable adjective as the main predication can be stated as follows. First, assume as above that the domain of a gradable adjective is partially ordered according to some dimensional. A gradable adjective φ in a context c can then be analyzed as a function that induces a tripartite partitioning of its (ordered) domain into: (i) a *positive extension* ($pos_c(\varphi)$), which contains objects above some point in the ordering (objects that are definitely φ in c), (ii) a *negative extension* ($neg_c(\varphi)$), which contains objects below some point the ordering (objects that are definitely not φ in c), and (iii) an *extension gap* ($gap_c(\varphi)$), which contains objects that fall within an "indeterminate middle", i.e., objects for which it is unclear whether they are or are not φ in c. The net effect of these assumptions is that the truth conditions of a sentence of the form *x is* φ in a context c can be defined as in (49) (where $\varphi(x)$ is the logical representation of "*x is* φ").[12]

(49) i. $||\varphi(x)||^c = 1$ iff x is in the positive extension of φ at c,

ii. $||\varphi(x)||^c = 0$ iff x is in the negative extension of φ at c, and

[12]As noted in section 1.1.1, the point that identifies the "lower edge" of the positive extension corresponds to a "norm" value for a given adjective, which Klein (1982) suggests can be identified with the extension gap.

iii. $||\varphi(x)||^c$ is undefined otherwise.

As observed in section 1.1.1, the partitioning of the domain into a positive and negative extension and extension gap is context-dependent, determined by the choice of *comparison class*. Roughly speaking, a comparison class is a subset of the domain of discourse that is determined to be somehow relevant in the context of utterance, and it is this subset that is supplied as the domain of the function denoted by the adjective. The role of the comparison class can be illustrated by considering an example like (50).

(50) Bill is tall

If the entire domain of discourse were taken into consideration when evaluating the truth of (50), then it would turn out to be either false or undefined, since relative to mountains, redwoods, and skyscrapers, humans fall at the lower end of an ordering along a dimension of *height*. As a result, the individual denoted by *Bill* would be at the lower end of the ordered domain of the adjective, and so would fall within the negative extension of *tall* (or possibly in the extension gap). When attention is restricted to humans, however, then a comparison class consisting only of humans is used as the basis for the partitioning of the domain of *tall*, and the truth or falsity of (50) depends only on the position of *Bill* in this smaller set.

An important constraint on the construction of a comparison class is that it must preserve the original ordering on the domain, in order to avoid undesirable entailments. For example, consider a context in which the ordering on the domain of *tall* is as in (51). If no restrictions were placed on the construction of a comparison class from D_{tall}, then the ordered set K_{tall} in (52) would be a possible comparison class, allowing for a partitioning of the domain as shown in (53).

(51) D_{tall} = {..., *Nadine, Bill, Aisha, Chris, Tim, Frances, Polly, Erik*, ...}

(52) K_{tall} = {..., *Aisha, Frances, Polly, Nadine*, ...}

(53) $pos_c(tall)= \{Polly, Nadine\}$
$neg_c(tall) = \{Aisha\}$
$gap_c(tall) = \{Frances\}$

In this context, (54) would be true while (55) would be false, a result which is inconsistent with our intuitions if the actual ordering on the domain of *tall* is as in (51).

(54) Nadine is tall.
(55) Frances is tall.

This undesirable result is avoided by invoking the "Consistency Postulate" informally defined in (56) (see Klein 1980, 1982, van Benthem 1983, Sanchez-Valencia 1994, and (63) below), which requires any partitioning of a subset of the domain of a gradable adjective to preserve the original ordering on the entire domain.

(56) *Consistency Postulate* (*informal*)
For any context in which *a is* φ is true and $b \geq a$ with respect to the ordering on the domain of φ, then *b is* φ is also true, and for any context in which *a is* φ is false, and $a \geq b$, then *b is* φ is also false.

A consequence of this condition is that (53) is not a possible partitioning: since the partitioning indicated in (53) makes (54) true, and $Frances \geq Nadine$ with respect to the original ordering in (51), the Consistency Postulate requires it to also be the case that (55) is true.

The Consistency Postulate has another important consequence: it entails that any objects in the domain of a gradable adjective that are ordered above objects in the positive extension of some comparison class in context *c* fall in the positive extension of a corresponding partitioning of the entire domain in *c*. For example, assume that the domain

of *tall* is as shown above in (51), and in context *c*, the partitioning on the comparison class (57) is as shown in (58).

(57) $K_{tall} = \{..., Nadine, Aisha, Frances, Polly, ...\}$
(58) $pos_c(tall)= \{Frances, Polly\}$
$neg_c(tall) = \{Aisha\}$
$gap_c(tall) = \{Nadine\}$

In this context (59) is true, and according to the Consistency Postulate, (60) is true as well, since $Erik \geq Polly$ with respect to the ordering in (51).

(59) Polly is tall.
(60) Erik is tall.

This discussion brings into focus the fundamental ideas underlying the vague predicate analysis. Given any set of objects partially ordered along a dimension δ, it is possible to define a family of (possibly partial) functions that induce a partitioning on the set in accord with the consistency postulate informally stated in (56). In effect, the vague predicate analysis claims that the interpretation of a gradable adjective with dimensional parameter δ is a value selected from this family of functions, though it may vary from context to context. On this view, the vagueness of sentences constructed out of gradable adjectives is due to the fact that it is necessary to choose a value from this family of functions for the adjective. The importance of the comparison class is that it provides a means of narrowing down the set of possible choices.

1.2.2 Comparatives

The analysis of comparative constructions within the vague predicate analysis builds on the intuitions underlying the Consistency Postulate (see McConnell-Ginet 1973, Kamp 1975, Klein 1980, 1982, van Benthem 1983, Larson 1988a, Sánchez-Valencia 1994).[13] Consider, for example, (61) and (62).

(61) Jupiter is larger than Saturn (is).

(62) The earth is as large as Venus (is).

Given the conditions imposed by the Consistency Postulate, it follows that if there is a context that makes the proposition expressed by *Jupiter is large* true but makes *Saturn is large* false, then it must be the case that the object denoted by *Jupiter* is ordered above the individual denoted by *Saturn* with respect to the ordering on the domain of *large*, i.e., it must be the case that Jupiter is larger than Saturn. Similarly, if every context in which the proposition expressed by *Venus is large* is true is such that *the earth is large* is true as well, then it must be the case that the object denoted by *the earth* is ordered at least as high as the object denoted by *Venus* in the domain of *large*, i.e., that the earth is as large as Venus.

This analysis can be made precise by building on the observation made at the end of the previous section that the interpretation of a gradable adjective in a context *c* is a member of a family of functions that partition a partially ordered set in accord with the Consistency Postulate. Specifically, we can introduce a set of *degree functions* that apply to a gradable adjective and return some member of this family; in particular, following Klein 1980, we can assume that the result of applying a degree function to a gradable adjective is always a complete function.[14] The underlying idea is that a degree function performs the role normally played by context: it fixes the denotation of the adjective, ultimately determining how the domain is to be partitioned. The difference is that all of the partitionings induced by a degree function are bipartite: none contain an extension gap.

---

[13] I will focus here on the analysis of comparative and equative constructions; see Klein 1980, 1982 for discussion of *how adj* questions and anaphoric *that adj* constructions.

[14] In Klein's analysis, the denotations of *very*, *fairly*, and other degree modifiers are taken from the set of degree functions.

Once we have degree functions, the Consistency Postulate can be restated more formally, as in (63) (where *GrAdj* is the set of gradable adjective meanings, *D* is the domain of discourse, and *Deg* is the set of degree functions; cf. Klein 1982:126).[15]

(63) *Consistency Postulate*

For all $\varphi \in GrAdj$, $a,b \in D$, $c \in C$, and $d \in Deg$:

$$\left[[\![(d(\varphi))(a)]\!]^c = 1 \wedge b \geq_\varphi a\right] \rightarrow \left[[\![(d(\varphi))(b)]\!]^c = 1\right] \text{ and}$$

$$\left[[\![(d(\varphi))(a)]\!]^c = 0 \wedge a \geq_\varphi b\right] \rightarrow \left[[\![(d(\varphi))(b)]\!]^c = 0\right]$$

The net effect of the Consistency Postulate is to ensure that the only admissible degree functions are those that induce partitionings of the domain of a gradable adjective in a way that is consistent with the inferences discussed in the previous section. Given this constraint, the interpretation of comparatives and equatives can be straightforwardly formalized in terms of quantification over degree functions, as in (64) and (65) the formalism adopted here is most similar to that in Klein 1982).[16]

(64) *a is more φ /φ-er than b:* $\exists d[(d(\varphi))(a)\ \&\quad (d(\varphi))(b)\ ]$

(65) *a is as φ as b:* $\forall d[(d(\varphi))(b) \rightarrow (d(\varphi))(a)]$

Consider the analysis of (61), which has the logical representation in (66).

(61) Jupiter is larger than Saturn (is).

(66) $\exists d[(d(large))(Jupiter)\ \&\quad (d(large))(Saturn)\ ]$

---

[15] For two objects a,b in the domain of a gradable adjective φ, $a \geq_\varphi b$ iff a is at least as great as b with respect to the ordering determined by the dimension identified by φ's dimensional parameter.

[16] I focus here on comparatives with *more* for perspicuity; see Klein 1980 and Larson 1988a for some discussion of comparatives with *less*. In addition, see Larson 1988a for some refinements of the basic analysis developed to handle the interpretation of quantificational expressions in the comparative clause (the complement of *than* or *as*).

According to (66), (61) is true just in case there is a function that, when applied to *large*, induces a partitioning of the domain of *large* so that the positive extension includes *Jupiter*, while the negative extension contains *Saturn*. Assuming the domain of *large* to be as in (67) (limiting the domain to the planets in the solar system), (61) is true, because there is a partitioning of the domain of *large* such that *Jupiter* is in the positive extension and *Saturn* is in the negative extension, namely the one shown in (68) (where the notation *a/b* indicates that *a* and *b* are nondistinct with respect to the ordering on the domain; I'm assuming for the sake of argument that Venus and the earth are the same size). [17]

(67) D_{large} = {*Pluto, Mercury, Mars, Venus/Earth, Neptune, Uranus, Saturn, Jupiter*}

(68) pos_d(*large*) = {*Jupiter*}

neg_d(*large*) = {*Pluto, Mercury, Mars, Venus/Earth, Neptune, Uranus, Saturn*}

Since the possible values of the function *d* must satisfy Consistency Postulate, partitionings such as (69) are impossible, and we derive the desired result that (61) entails that for any context, *Jupiter* is ordered above *Saturn* in the domain of *large*; i.e., that Jupiter is larger than Saturn is.

(69) pos_d(*large*) = {*Uranus, Saturn*}

neg_d(*large*) = {*Pluto, Mercury, Mars, Venus/Earth, Neptune, Jupiter*}

The analysis of the equative construction is very similar. The logical representation of (62) is (70).

(62) The earth is as large as Venus (is).

[17]To distinguish the partitioning introduced by a degree function from the context-dependent partitioning associated with the absolute construction, I will represent the positive and negative extensions of a gradable adjective φ with respect to a particular degree function d as $pos_d(\varphi)$ and $pos_d(\varphi)$, respectively.

(70) $\forall d[(d(large))(Venus) \rightarrow (d(large))(the\ earth)]$

(62) is true just in case every value of *d* that results in a partitioning of the domain of *large* in which the object denoted by *Venus* is in the positive extension is also a partitioning in which the object denoted by *the earth* is in the positive extension. Since all values for *d* must obey the Consistency Postulate, this will be the case in a context in which the ordering on the domain of *large* is as in (67).[18]

Finally, it should be observed that this analysis does not entail that the two conjuncts in the logical representation of the comparative (or equative) are true in the context of utterance. Consider, for example, the analysis of (71).

(71) My Volvo is faster than Jason's Honda.

(72) $\exists d[(d(fast))(my\ Volvo)\ \&\ \ (d(fast))(Jason's\ Honda)\ ]$

All that is necessary to satisfy the truth conditions of (71) is that there be some partitioning of the domain of *fast* that makes *my Volvo is fast* true and *Jason's Honda is fast* false; for example, the one in (73).

(73) pos_d(*fast*) = {*my Volvo, my old Dodge, ..., Ken's Jaguar, Jorge's Morgan, Kari's Dodge*}

neg_d(*fast*) = {*..., Rachel's scooter, Jason's Honda, ...*}

It does not follow, however, that this partitioning is the one derived contextually (relative to an appropriate comparison class) in the context of utterance. For example, it might be the case that the actual partitioning of the domain of *fast* in the context of utterance of (71) is as in (74), in which case neither (75) nor (76) would be true, according to the analysis of the

[18]A positive result of this analysis is that it accounts for the fact that (i) and (ii) are logically equivalent (see Klein 1980).

(i) Mars is not as large as Jupiter.
(ii) Jupiter is larger than Mars.

absolute form outlined above.

(74) pos_c(*fast*) = {..., *Ken's Jaguar, Jorge's Morgan, Kari's Dodge*}
neg_c(*fast*) = {..., *Rachel's scooter, Jason's Honda, my Volvo, ...*}
gap_c(*fast*) = {..., *my old Dodge, ...*}

(75) My Volvo is fast.

(76) Jason's Honda is fast.

1.2.3 Problems with the Vague Predicate Analysis

The following sections discuss a number of problems for the analysis of comparative constructions outlined here, and, by extension, for the general analysis of gradable adjectives within the vague predicate approach. I will focus on four problems specifically, which involve facts from two of the domains discussed in section 1.1: polarity and incommensurability.

1.2.3.1 Cross-Polar Anomaly

Although Klein (1980) does not explicitly discuss the differences between antonymous pairs of positive and negative adjectives such as *tall/short, clever/stupid*, and *safe/dangerous*, a natural approach to adjectival polarity within a vague predicate analysis is to assume, building on the observations about the logical properties of gradable adjectives discussed in section 1.1.4.1, that the domains of antonymous pairs are distinguished by their orderings: one is the inverse of the other.[19] A positive result of this assumption is that it explains why sentences like (77) are valid.

(77) Jason's Honda is more dangerous than my Volvo if and only if my Volvo is safer than Jason's Honda.

[19]This idea is implicit in Klein's (1980:35) discussion of examples like *Mona is more happy than Jude is sad* (see the discussion of comparison of deviation in section 1.2.3.2 below). Sánchez-Valencia (1994) shows how this assumption can be used to build an explanation of the monotonicity properties of polar adjectives.

If the domains of *safe* and *dangerous* are identical except for the ordering on the objects they contain, and if the ordering of one is the inverse of the other, then any partitioning of the domain of *dangerous* that satisfies the truth conditions of the first conjunct in (77)–i.e., any partitioning that makes *Jason's Honda is dangerous* true and *my Volvo is dangerous* false–will have the opposite effect on the domain of *safe*, since the two sets, in effect, stand in the dual relation to each other. For example, any function that partitions the domain of *dangerous* as in (78) must induce a corresponding partitioning on the domain of *safe* as shown in (79), with the result that both conjuncts of (77), shown in (80) and (81), are true.

(78) $D_{dangerous}$ = {..., *c, b, my Volvo, a, ..., x, Jason's Fiat, y, x, ...*}
pos_d(*dangerous*) = {..., *x, y, Jason's Fiat, z, ...*}
neg_d(*dangerous*) = {*a, my Volvo, b, c, ...*}

(79) D_{safe} = {..., *x, y, Jason's Fiat, z, ..., a, my Volvo, b, c, ...*}
pos_d(*safe*) = {*a, my Volvo, b, c, ...*}
neg_d(*safe*) = {..., *x, y, Jason's Fiat, z, ...*}

(80) $\exists d[(d(dangerous))(Jason's\ Fiat)\ \&\ (d(dangerous))(my\ Volvo)]$

(81) $\exists d[(d(safe))(my\ Volvo)\ \&\ (d(safe))(Jason's\ Fiat)]$

This analysis runs into problems when confronted with examples of cross-polar anomaly, however, which is illustrated by (82) and (83).

(82) #Mona is happier than Jude is sad.

(83) #Bill is older than Chelsea is young.

Consider, for example, the case of (82), which has the logical representation in (84).

(84) $\exists d[(d(happy))(Mona)\ \&\ \ (d(sad))(Jude)]$

According to (84), (82) is true just in case there is a function that effects a partitioning of the domains of *happy* and *sad* in such a way that *Mona is happy* is true and *Jude is sad* is false; e.g., if Mona is very happy and Jude is not very sad. Given the assumption that the domains of the antonymous pair *happy* and *sad* have opposite ordering relations, in a context in which the domain of *happy* is (85), the domain of *sad* is (86).

(85) $D_{happy} = \{x, y, Jude, z, Mona\}$

(86) $D_{sad} = \{Mona, z, Jude, y, x\}$

In such a context, there is a function that satisfies the truth conditions associated with (84), for example, the one that induces the partitioning of the domains of *happy* and *sad* shown in (87).

(87) $pos_d(happy) = \{Jude, z, Mona\}$
$neg_d(happy) = \{x, y\}$
$pos_d(sad) = \{y, x\}$
$neg_d(sad) = \{Mona, z, Jude\}$

As a result, (82) should be true. More generally, (82) should be perfectly interpretable: nothing about the architecture of the analysis predicts that comparatives constructed out of antonymous pairs of adjectives should be anomalous. The basic problem is that the assumption that the domains of positive and negative adjectives contain the same objects under inverse ordering relations–an assumption that is necessary to account for the validity of sentences like (77)–predicts that it should be possible to interpret sentences like (82) in

the way I have outlined here. One could stipulate that comparison between positive and negative pairs of adjectives is impossible, but there is no aspect of the analysis of comparatives within the vague predicate approach that derives this constraint. Moreover, such a stipulation would be empirically unmotivated, since comparison of deviation constructions, discussed in the next section, show that comparison between positive and negative adjectives is possible in certain circumstances.

1.2.3.2 Comparison of Deviation

The problems presented by cross-polar anomaly for a vague predicate analysis extend to comparison of deviation constructions such as (88) and (89), though their effects are somewhat different.

(88) Robert is as short as William is tall.

(89) Mona is more happy than Jude is sad.

In the discussion of these constructions in section 1.1.4.3, I observed that they have two important semantic characteristics. First, they compare the relative extents to which two objects deviate from some contextually-determined "standard" value associated with the adjective, and second, they entail that the properties predicated of the compared objects are true in the absolute in the context of utterance. The problem that comparison of deviation presents for the vague predicate analysis is that the semantic analysis of sentences like (88) and (89) does *not* derive the fact that comparison of deviation constructions, unlike standard comparatives, entail that the corresponding absolute sentences are true in the context of utterance.

Consider the case of (89), which has the logical representation in (90) (see Klein 1980:35-36 for discussion of the same example).

(90) $\exists d[(d(happy))(Mona)\ \&\ \ (d(sad))(Jude)]$

The logical representation in (90) is exactly the same as the logical representation assigned to the example of cross-polar anomaly (82) discussed above. The problem presented by cross-polar anomaly was that there was no way to explain why such examples are anomalous; the problem of comparison of deviation is not the case that the analysis is inconsistent with the interpretations of these structures–as noted in the previous section, the logical representation in (90) would be true if e.g. Mona were very happy and Jude were not very sad, which is a rough paraphrase of what (89) means–but rather that it is too weak: it does not entail that Mona is happy and Jude is sad. To see why, assume that the domains of *happy* and *sad* are as in (91) and (92), and assume also that the contextual partitionings of the domains–the partitionings relevant to the interpretation of the absolute construction (e.g., *Jude is happy*)–are as in (93) and (94).

(91) $D_{happy} = \{x, Mona, y, z, Jude\}$

(92) $D_{sad} = \{Jude, z, y, Mona, x\}$

(93) $pos_c(happy) = \{z, Jude\}$
$neg_c(happy) = \{x, Mona, y\}$

(94) $pos_c(sad) = \{y, Mona, x\}$
$neg_c(sad) = \{Jude, z\}$

In this context, (95) is false, because *Mona* falls in the negative extension of the adjective, while (96) is true, since *Jude* appears in the positive extension of *happy*.

(95) Mona is happy.

(96) Jude is happy.

Similarly, (97) is false, because *Jude* falls in the negative extension of *sad*, and (98) is true.

(97) Jude is sad.

(98) Mona is sad.

Unfortunately, the comparison of deviation construction (89) is also true in this context, since there is a function *d* that introduces alternative partitionings of the domains of *happy* and *sad*–those shown in (99) and (100)–which satisfies the truth condition of (90).

(99) $pos_d(happy) = \{Mona, y, z, Jude\}$
$neg_d(happy) = \{x\}$

(100) $pos_d(sad) = \{x\}$
$neg_d(sad) = \{Jude, z, y, Mona\}$

With respect to the partitionings in (99) and (100), *Mona is happy* is true and *Jude is sad* is false, therefore (89) should be true. In other words, since the analysis requires only that there is a possible partitioning of the domain of *happy* and *sad* in which *Mona is happy* is true and *Jude is sad* is false, it allows for the possibility that (89) is true while (95) and (97) are false. This result is inconsistent with the facts of comparison of deviation, however: (89) entails that *Mona is happy* is true and *Jude is sad* is false in the context of utterance.

In fact, the analysis of (89) is even more problematic. If the ordering on the domains of *happy* and *sad* are as specified in (91) and (92), (101) is also true, since there is a function that partitions the domain so that *Jude is happy* true and *Mona is happy* false, namely the one that generates partitionings equivalent to those shown in (93) and (94).

(101) Jude is happier than Mona.

The analysis thus allows for the possibility that (89) and (101) can be true in the same context. The fact that (102) is a contradiction, however, shows that this result is incorrect.[20]

(102) Mona is more happy than Jude is sad, but Jude is happier than Mona.

1.2.3.3 Incommensurability

A third problem for the vague predicate analysis comes from the phenomenon of incommensurability, illustrated by sentences like (103) and (104).

(103) #Morton is as tall as Richard is clever.
(104) #My copy of *The Brothers Karamazov* is heavier than my copy of *The Idiot* is old.

(103) and (104) contrast quite clearly with sentences like (105).

(105) Our Norfolk Island Pine is as tall as its branches are long.

In section 1.1.3, I used the contrast between examples like (103)-(104) and (105) as the basis for the descriptive generalization in (106).

(106) A comparative construction is semantically well-formed only if the compared adjectives have the same dimensional parameter.

(106) reflects the general hypothesis that comparison of two objects is possible only if they

[20]Comparison of deviation constructions involving equatives are similarly problematic. Consider an example like (i), which has the interpretation in (ii).

(i) Mona is as happy as Jude is sad.
(ii) $\forall d[(d(sad))(Jude) \rightarrow (d(happy))(Mona)]$

The logical representation in (ii) is *consistent* with the interpretation we want to derive–it is true, for example, in a context in which Mona is quite happy and Jude is quite sad–but it has the same problem as the comparative: it does not entail that the propositions *Mona is happy* and *Jude is sad* are true in the context of utterance.

are ordered according to the same dimension. For example, in the case of a sentence like (105), the dimension common to both *tall* and *long* could be described as "linear measurement": both adjectives order the objects in their domain according to some general notion of length, a vertical one in the case of *tall*, and a horizontal one in the case of *long* (cf. Cresswell 1976). The problem with an example like (104), on this view, is that *long* and *tragic* have different dimensional parameters, therefore the objects in the domains of the adjectives are not ordered according to the same dimension, and comparison is impossible.

Although this approach to incommensurability seems reasonable, it should nevertheless be the case that the actual constraint underlying the descriptive generalization in (106) is sensitive to the *linguistic* properties of gradable adjectives and the comparative construction. That is, the explanation of incommensurability should be stated in terms of the semantic properties of linguistic expressions (i.e., gradable adjectives and comparative morphology), rather than in terms of a general, conceptual constraint on comparison. This is shown by a sentence like (107).

(107) My copy of *The Brothers Karamazov* is higher on a scale of heaviness than my copy of *The Idiot* is on a scale of age.

(107) represents a coherent thought, and moreover is a reasonable paraphrase of the intended meaning of a sentence like (104). Most importantly, the fact that (107) is a perfectly acceptable *linguistic* expression indicates that the anomaly of examples like (103) and (104) must be due to the interaction of the meanings of the "incomparable" adjectives *in the context of* the comparative construction. In other words, it must be some semantic property of the adjectives in these examples which, when they are inserted into the comparative, triggers incommensurability.

The problem for the vague predicate analysis is that it does not provide a means of explaining incommensurability in these terms. To see why, first consider the analysis of a typical example of comparative subdeletion such as (108).[21]

(108) The tree is taller than the ceiling is high.

(109) $\exists d[(d(tall))(tree) \;\&\; (d(high))(ceiling)]$

According to (109), (108) is true just in case there is a function *d* that introduces a partitioning on the domain of *tall* so that *the tree is tall* is true, and also introduces a partitioning on the domain of *high* so that *the ceiling is high* is false (e.g., if the tree is very tall and the ceiling is not very high). Note that in any context, the domains of *tall* and *high* may contain the same objects, but they need not be ordered in the same way. The fact that (109) is interpretable indicates that it must be the case that *d* can apply to sets with unrelated orderings.

Now consider the analysis of a context in which the domain of discourse is restricted to include only my copies of Dostoevski's novels, and the ordering on the domain of *heavy* is (110), while the ordering on the domain of *old* is (111).

(110) D_{heavy} = {*Crime and Punishment, The Devils, The Idiot, The Brothers Karamazov*}

(111) D_{old} = {*The Idiot, The Devils, The Brothers Karamazov, Crime and Punishment*}

According to the analysis of comparatives outlined in section 1.2.2, the interpretation of (104) is (112).

(112) $\exists d[(d(heavy))(The\ Brothers\ K) \;\&\; (d(old))(The\ Idiot)]$

In order to determine whether the vague predicate analysis supports an explanation of incommensurability, we first need to ask the following question: is there a function that

[21]"Subdeletion" structures are those of the form *x is more A_1 than A_2*, where A_1 and A_2 are lexically distinct (cf. Bresnan 1973, 1975, Grimshaw 1987, Corver 1990, 1993, Izvorski 1995).

partitions the domain so that *The Brothers Karamazov* is in the positive extension of *heavy* and *The Idiot* is in the negative extension of *old*? Such a function would have to apply to sets with unrelated orderings, but as we saw above with the analysis of (108), such functions must be available. Consider then a function *d*, which partitions the domains of *heavy* and *old* so that the highest-ranked object in each set falls in the positive extension and the rest are in the negative extension. The partitionings induced by *d* on the domains of *heavy* and *old* in the context specified above are shown in (113) and (114).

(113) $pos_d(heavy)$ = {*The Brothers Karamazov*}
$neg_d(heavy)$ = {*Crime and Punishment, The Devils, The Idiot*}

(114) $pos_d(old)$ = {*Crime and Punishment*}
$neg_d(old)$ = {*...The Idiot, The Devils, The Brothers Karamazov*}

Since *d* satisfies the truth conditions specified in (112), (104) should be true. But if the vague predicate analysis supports an explanation of the anomaly of a sentence like (104), it must be the case that the computation outlined in the previous paragraph is impossible. Note that it is not enough to show that a function such as *d* is actually unavailable–this would have the result that (104) is false, rather than the desired result that the sentence is anomalous. It must be shown that it is impossible to even ask the question "is there a function that partitions the domains of *long* and *tragic* in the relevant way".

Indeed, one way to derive this result would be to build on the initial assumption that the objects in the domains of *heavy* and *old* are ordered according to different dimensions, and to somehow encode a "dimension identity" requirement into the semantics of the comparative construction. The problem for a vague predicate analysis is that there is no way to achieve this result, since the ordering relations associated with the adjectives are presumed, and so do not play a direct role in the compositional interpretation of a construction like (104). If the truth conditions for the comparative require only that there be a function which partitions the domains of the adjectives as in (113) and (114), then, at least in a context in which the domains of the adjectives are the same, as in (110) and (111), comparison should be possible. The only way to rule out these sentences is by stipulating the dimension identity requirement. (107), however, shows that this stipulation must be localized to comparative constructions; it cannot be a general conceptual constraint on comparison. Ideally, then, the anomaly of examples like (104) should follow from the semantics of gradable adjectives and the comparative construction, and not from a general stipulation.

1.2.3.4 Negative Adjectives and Measure Phrases

Klein (1980:27-28) discusses the analysis of sentences such as (115), which involve gradable adjectives and measure phrases.

(115) Mona is three feet tall.

Klein suggests that measure phrases such as *three feet* denote equivalence classes of objects that are three feet in height (see also Cresswell 1976, von Stechow 1984a, Klein 1991), and that *Mona is three feet tall* is true just in case there is some object in the denotation of *three feet* (the equivalence class of objects which are three feet in height) and Mona is nondistinct from that object with respect to the ordering imposed by *tall*. On this view, (115) has the truth conditions in (116).

(116) $(\mathit{three\ feet}(\mathit{tall}))(\mathit{Mona}) = 1$ iff $\exists y \in$ *three feet:* $\mathit{Mona} =_{tall} y$

This analysis fails to make a distinction between positive and negative adjectives, however. As a result, it does not support an explanation of the well-known fact that negative adjectives such as *short* do not permit measure phrases:

(117) #Mona is three feet short.

Assuming that *tall* and *short* are extensionally equivalent, but associated with the opposite ordering relations, (117) should not only be non-anomalous, it should be logically equivalent to (115):

(118) $(\mathit{three\ feet}(\mathit{short}))(\mathit{Mona}) = 1$ iff $\exists y \in$ *three feet:* $\mathit{Mona} =_{short} y$

1.2.4 Summary

The primary claim of the vague predicate analysis is that gradable adjectives are of the same semantic type as other predicative expressions–they denote (possibly partial) functions from objects to truth values–but their domains are partially ordered with respect to some dimension. An important aspect of this type of analysis–one that distinguishes it from the type of analysis I will discuss in section 1.3–is that gradable adjectives are of the same semantic type as non-gradable adjectives. The difference between gradable and non-gradable adjectives is, in effect, a *sortal* one: they denote different sorts of predicates (partial vs. complete functions; ordered vs. unordered domains). This difference can be used as the basis for an explanation of the distributional characteristics I discussed in the introduction–the fact that gradable adjectives, but not non-gradable ones, can appear in degree constructions (and with degree modifiers). If we assume that the set of degree functions contains functions from partial functions from individuals to truth values to complete functions from individuals to truth values, then only gradable adjectives provide appropriate arguments for degree functions. The anomaly of sentences like (119) and (120) follows from the analysis of comparatives (and other degree constructions) in terms of quantification over degree functions, and the assumption that degree modifiers (such as *extremely, very, quite*) denote degree functions (Klein 1980, McConnell-Ginet 1973).

(119) ??Nixon is extremely dead.

(120) ??Nixon is more dead than Reagan.

The problem with both examples is that the non-gradable adjective *dead* denotes a complete function from objects to truth values. When the denotation of *dead* is supplied as an argument to the degree functions in these sentences–*extremely* in (119), and *d* (the variable quantified over by the comparative) in (120)–the result is sortal anomaly, since a degree functions expect a partial function from objects to truth values as its argument.

This discussion suggests that the vague predicate analysis supports an explanation of the distributional characteristics of gradable adjectives; the facts discussed in the previous section, however, show that it fails to provide an explanation of several other important sets of facts: cross-polar anomaly, comparison of deviation, incommensurability, and the unacceptability of measure phrases with negative adjectives. In the next section, I will outline an alternative semantic analysis of gradable adjectives, which differs from the vague predicate analysis in that it expands the ontology to include abstract representations of measurement, and defines the interpretation of gradable adjectives in terms of such abstract objects.

1.3 The Scalar Analysis

There are two primary differences between the vague predicate analysis of gradable adjectives and the analysis that I referred to in section 1.1.1 as the "scalar analysis". The first difference concerns the semantic type of a gradable adjective. Whereas the vague predicate analysis assumes that gradable adjectives have the same semantic type as other adjectives (and other predicative expressions in general)–they denote functions from individuals to truth values–the scalar analysis reanalyzes gradable adjectives as relational expressions, specifically, relations between individuals and abstract representations of measurement, or "degrees". The second difference concerns the nature of the ordering on the domain of the adjective. Both analyses claim that a partial ordering can be imposed on the domain of the adjective, but they differ in their assumptions about how the ordering is derived. In the vague predicate analysis, the ordering on the domain is presumed; in the scalar analysis, however, the adjective imposes an ordering on its domain by relating objects to degrees on a scale.

In the following sections, I will go over the basic assumptions of the scalar analysis in more detail. As in the discussion of the vague predicate analysis, I will focus on the interpretation of comparatives. In particular, I will show that the introduction of scales and degrees into the ontology provides the basis for an explanation of the facts that were problematic for a vague predicate analysis: cross-polar anomaly, comparison of deviation, incommensurability, and the distribution of measure phrases. In section 1.4, however, I will introduce a set of facts involving comparatives and scope that are problematic for the traditional scalar analysis of gradable adjectives and comparatives. These facts will provide the empirical basis for the analysis of gradable adjectives and degree constructions that I will develop in chapter 2, which falls within the general category of scalar analyses, but differs it its claims about the core meaning of gradable adjectives and the compositional semantics of comparatives and other degree constructions.

1.3.1 Degree Arguments

Cresswell (1976:266) suggests that "[w]hen we make comparisons we have in mind points on a scale". Building on this intuition, Cresswell develops a theory in which gradable adjectives are analyzed as expressions whose semantic function is to define a mapping between objects and points on a scale. Intuitively, a scale is an abstract representation of measurement: an infinitely long measuring stick, which provides a representation of the amount to which an object possesses some gradable property. To make things precise, I will define a scale as a dense, linearly ordered set of points, or "degrees", where the ordering is relativized to a DIMENSION. As noted in section 1.1.2, a dimension corresponds to a gradable property such as *height*, *length*, *speed*, *density*, *beauty*, etc., and provides a means of differentiating one scale from another.[22]

Once scales and degrees are introduced into the ontology, it becomes possible to analyze gradable adjectives as relational expressions, specifically, as expressions that relate objects in their domains to degrees on a scale, where the particular scale is specified by the dimensional parameter the adjective (see e.g. Seuren 1973, Cresswell 1976, Hellan 1981, Hoeksema 1983, von Stechow 1984a, Heim 1985, Bierwisch 1989, Lerner and Pinkal 1992, 1995, Moltmann 1992a, Rullmann 1995, Gawron 1995 and others for approaches along these lines).[23] A consequence of defining the interpretation of a gradable adjective in this way is that the ordering that can be imposed on its domain is derived from a semantic property of the adjective itself: by relating objects in a set to degrees on a scale (a totally ordered set of points), a gradable adjective determines a partial ordering on that set.[24] This property of a gradable adjective's meaning represents the fundamental difference between the scalar analysis and the vague predicate analysis discussed in section 1.2, since in the latter the ordering of the domain of a gradable adjective is presumed.

Within a framework in which gradable adjectives are analyzed as relational expressions, the logical representation of a sentence of the form *x is* φ can be stated as in (121), which has the truth conditions in (122), where δ_φ is a function that maps objects to the scale associated with φ.

[22]This aspect of the scalar analysis is crucial to the analysis of incommensurability that I will discuss below in section 1.3.3.1; in chapter 3, I will show that a similar distinction provides the basis for an explanation of cross-polar anomaly.

[23]A different type of analysis, in which adjectives denote *functions* from objects to scalar values, is developed in Wunderlich 1970 and Bartch and Vennemann 1972, and forms the basis for the analysis I will develop in chapter 2. I will have more to say about the general characteristics of this alternative approach there; for now, it should simply be observed that these approaches fall into the general category of scalar analyses, since they assume that scales and degrees are part of the ontology.

[24]Although a scale is defined as a totally ordered set of points, nothing prohibits a gradable adjective from relating different objects in its domain to the same degree on the scale; as a result, the ordering on the domain is partial.

(121) $\varphi(x,d)$

(122) $||\varphi(x,d)|| = 1$ iff $\delta_\varphi(x) \geq d$

Stated informally, *x is* φ is true just in case the projection of *x* on the scale associated with φ (i.e., the degree to which *x* is φ) is at least as great as *d*. The first question raised by this analysis is the following: what is the value of *d* in (121)? For a sentence of the form *x is* φ (an absolute construction), the answer is that *d* represents a "standard". Intuitively, a standard-denoting degree is a degree that identifies the point on a scale that can be used to separate those elements for which the statement *x is* φ is true from those elements for which *x is* φ is false in some context.

For illustration of this idea, consider an example like (123), which has the logical representation in (124), where $d_{s(long)}$ is the degree argument of *long* and denotes a contextually determined standard of "longness".

(123) *The Brothers Karamazov* is long.

(124) $long(BK,d_{s(long)})$

According to the truth conditions in (122), (123) is true if and only if $\delta_{long}(BK) \geq d_{s(long)}$ holds, i.e., just in case the projection of *The Brothers Karamazov* on a scale of length is at least as great as the standard of longness in the context of utterance. Note that the structure of the scale–specifically, the fact that scales are defined as totally ordered sets of points along some dimension–ensures not only that the relative ordering of the standard-denoting degree and the degree that represents the measure of *The Brothers Karamazov*'s length can be determined, but also that the standard value must be a degree of length. If it were a degree along some other dimension, then it would not be an element of the same scale as $\delta_{long}(BK)$. As a result, the partial ordering relation associated with the absolute would be undefined, rendering the sentence uninterpretable. I will return to this point in section 1.3.3.

For the moment, I will leave open the question of how exactly the standard value is determined, and assume following Bierwisch 1989 that the standard value is determined contextually, relative to a particular comparison class.[25] As observed in section 1.1.1, this hypothesis forms the basic explanation of vagueness in a scalar analysis of gradable adjectives. The context-dependency of a sentence of the form *x is* φ in the scalar analysis is parallel to the context-dependency of this type of sentence in the vague predicate analysis: the actual value of the standard in a context is determined by some subset of the domain of the gradable adjective that is taken to be relevant in that context, i.e., the comparison class. If the comparison class is changed, then the standard value may be shifted accordingly. Changing the identity of the standard does not affect the overall ordering of the degrees on the scale, however, so the scalar analysis derives the result that the ordering of objects in a comparison class preserves the ordering on the entire domain, as observed in section 1.1.1.[26]

1.3.2 Comparatives

In a traditional scalar analysis, in which gradable adjectives denote relations between objects and degrees, comparatives are typically analyzed as quantificational expressions, specifically, as expressions that quantify over degrees (see e.g., Seuren 1973, Hellan 1981, von Stechow 1984a,b, Heim 1985, Lerner & Pinkal 1992, 1995, Gawron 1995, Hazout 1995, Rullmann 1995). For example, in Heim 1985, comparatives are analyzed as *indefinite degree descriptions*,

[25]In chapter 2, section 2.2.2, I will go into the identification of the standard value in much greater detail.

[26]It should also be observed that nothing requires e.g. the standard of *tallness* to be the same as the standard of *shortness*. As a result, we allow for the possibility that an object may be neither tall nor short in a particular context. For example, if the standard of *tallness* for humans is 5'11'' and the standard of shortness is 5'5'', then objects whose projections on the scale of height fall between these two degrees are neither tall nor short. (In the terminology of the vague predicate analysis, such objects fall in the "extension gap".) Of course, it must be the case that the standard of shortness cannot *exceed* the standard of tallness. That is, we must assume a general constraint that prohibits a context in which the standard of *tallness* is, for example, 5'9'' and the standard of shortness is 5'11'', otherwise some objects could be both tall and short. See Bierwisch 1989 for relevant discussion.

which restrict the possible values of the degree argument of a gradable adjective.[27] In this analysis, the comparative morpheme denotes a relation between two degrees: one bound by an existential quantifier, and one introduced by the comparative clause (the complement of *than* or *as*). The logical representation of a typical comparative like (125) is (126).

(125) *x is more* φ *than* d_c

(126) $\exists d[d > d_c][\varphi(x,d)]$

Given the truth conditions for the absolute construction presented in section 1.3.1, a statement of the form *x is more* φ *than* d_c, where d_c is the degree denoted by the comparative clause, is true just in case there is a degree d such that d exceeds d_c and x is at least as φ as d. Throughout this thesis, I will adopt von Stechow's (1984a) position that the comparative clause is a type of definite description that denotes a maximal degree (see also Rullmann 1995 for extensive discussion of this issue).[28] According to von Stechow, the

[27]Two remarks are in order here. First, although many scalar analyses of comparatives are stated in terms of existential quantification over degrees, some analyzes comparatives as universal quantification structures (Cresswell 1976; see also Postal 1974, Williams 1977) or as generalized quantifiers (Moltmann 1992, Hendriks 1995). I focus here on the existential analysis for perspicuity, but I will examine the other accounts in more detail below.

Second, in the subsequent discussion, I will ignore the distinction between *phrasal* comparatives–comparatives in which the complement of *than* is a non-clausal constituent, and *clausal* comparatives–comparatives in which the complement of *than* is a (possibly reduced) clause. Although this distinction is an important one, it is not crucial to the general discussion here. It will be examined in detail in chapter 2, however.

[28]A number of researchers, including Cresswell 1976, Lerner & Pinkal 1992, Moltmann 1992a, and Gawron 1995, have proposed that the comparative clause is a universal quantification structure rather than a definite description. On this view, a sentence like (i) has an interpretation that corresponds roughly to the paraphrase in (ii).

(i) Some star is brighter than Venus has ever been.

(ii) Some star is brighter than every degree d such that Venus has ever been d-bright.

Although I will not attempt to resolve this debate here, I will point out two facts that argue in favor of a definite description analysis. First, the comparative clause supports discourse anaphora, as shown by (iii), in which the anaphor *that* picks up its reference from the comparative clause in the first conjunct: it denotes the degree to which Venus is bright.

complement of *than* denotes the set of degrees that satisfy the restriction derived by abstracting over the degree variable in the comparative clause (in (128), the set of degrees that are at least as great as the degree to which *The Idiot* is long; the mapping from the syntactic structure to this interpretation is straightforward if, as argued in Chomsky 1977, the complement of *than* is a *wh*-construction). This set is then supplied as the argument of a covert maximality operator, which I have represented as MAX. The interpretation of MAX is given in (127), where D is a totally ordered set of degrees (cf. Rullmann 1995).

(127) $\text{MAX}(D) = \iota d \in D[\forall\, d' \in D\text{: } d \geq d']$

For an illustration of the basic approach consider the analysis of (128), which has the logical representation in (129).

(128) *The Brothers Karamazov* is longer than *The Idiot*.

(129) $\exists d[d > \text{MAX}(\lambda d'.long(Idiot,d'))][long(BK,d)]$

According to (129), (128) is true iff there is a degree such that d exceeds the maximal degree to which *The Idiot* is long, and *The Brothers K.* is at least as long as d. In a context such as (130), then, where d_I denotes the degree of *The Idiot's* length and d_{BK} denotes the degree of

(iii) Mars isn't brighter than Venus is, but the earth is brighter than that.

Second, the comparative clause does not show the kind of scope ambiguities we expect of a universal quantification structure. If the comparative clause involved universal quantification, then (i) should permit a reading in which the comparative clause has scope over the indefinite subject, as in (iv).

(iv) For every degree d such that Venus has been d-bright, some star is brighter than d.

But (i) does not have such a reading: this sentence can only be interpreted as an existential claim about stars. Note that the unavailability of this reading is not due to a constraint on the complement of *than*: the ambiguity of e.g. *Some star is larger than every planet* shows that the complement of *than* can take scope over an indefinite subject.

The Brothers Karamazov's length, (128) is true.

(130) LENGTH: 0 ———————— d_I ———————— d_{BK} ————————— ∞

Equatives and comparatives with *less* are analyzed in essentially the same way, the only difference being the ordering relation introduced by the degree morpheme: $\geq$ for equatives; $<$ for *less*.

An important aspect of this analysis is that although the comparative construction establishes a relation between the projections of two objects on a scale, it does not make reference to a standard value. Since the truth conditions for the absolute construction are stated in term of a relation between the projection of an object on a scale and a standard value, the comparative does not support inferences to the absolute. That is, a sentence like (128) supports neither the inference in (131) nor the inference in (132), since the truth of both of these expressions can be determined only by evaluating the relation between the degrees identifying the subjects' lengths and the degree which denotes a standard of length (in some context).

(131) *The Brothers Karamazov* is long.

(132) *The Idiot* is not long.

1.3.3 Solutions to the Problems for the Vague Predicate Analysis

The primary difference between the scalar analysis and the vague predicate analysis is that the former introduces scales and degrees into the ontology. It is precisely this difference that provides a basis for empirically distinguishing between these two approaches. Specifically, the introduction of scales, and the characterization of scales as sets of degrees ordered along a dimension, provides a basis for explaining the facts that were problematic for the vague predicate analysis: incommensurability, cross-polar anomaly, the distribution of measure phrases, and comparison of deviation. In the following sections, I will reexamine the phenomena, showing how the hypothesis that gradable adjectives define mappings between objects and scales supports explanations of the facts.

1.3.3.1 Incommensurability

The general problem of incommensurability, illustrated by examples like (133), can be stated as follows: how do we explain the fact that subdeletion constructions involving adjectives that are in some intuitive sense "incomparable" are anomalous?

(133) #My copy of *The Brothers Karamazov* is heavier than my copy of *The Idiot* is old.

As observed in section 1.2.3.3, the anomaly of (133) must be explained in terms of the linguistic properties of the adjectives and the comparative construction, rather than in terms of a general conceptual restriction on comparison of objects ordered according to different dimensions. The problem for the vague predicate analysis was that this requirement could not be derived from the semantic properties of gradable adjectives and the comparative construction. In contrast, in the scalar analysis, this requirement follows directly from basic assumptions about scales.

In section 1.3.1, degrees were defined as elements of a scale, and a scale was defined as a totally ordered set of points along some dimension. The importance of the dimension is that it distinguishes one scale from another. For example, a scale S_ϕ of degrees ordered along dimension ϕ and a scale S_ψ of degrees ordered along dimension ψ are different sets. As a result, for any $d_\phi \in S_\phi$ and any $d_\psi \in S_\psi$, the expression $d_\phi \mathbf{R} d_\psi$ is undefined, where $\mathbf{R}$ is an ordering on degrees. Put another way, for any two degrees d_1 and d_2, $d_1\mathbf{R}d_2$ is defined only if d_1 and d_2 are degrees on the same scale. These general constraints on ordering relations interact with the semantics of comparatives to explain incommensurability.

According to the analysis of comparatives as quantificational structures that restrict the possible value of the degree argument of a gradable adjective to be a degree that satisfies the conditions imposed by the comparative, the logical representation of a sentence like (133)

is (134).

(134) $\exists d[d > \textsc{max}(\lambda d'.old(my\ copy\ of\ The\ Idiot,d'))][heavy(my\ copy\ of\ The\ Brothers\ K,d)]$

The formula in (134) is true just in case for some degree *d* such that *d* exceeds the maximal degree to which my copy of *The Idiot* is old, the degree to which my copy of *The Brothers K* is heavy is at least as great as *d*. That is, there must be some degree that satisfies both the restriction in (134) and the formula in (135), which represent the truth conditions of the absolute form (see section 1.3.1), where δ_{heavy} is a function from objects to the scale associated with *heavy*.

(135) $\delta_{heavy}(my\ copy\ of\ The\ Brothers\ K) \geq d$

Since the ordering relation introduced by the comparative morpheme is defined only for degrees on the same scale, the only objects that satisfy the restriction imposed by the comparative are degrees on a scale of age; as a result, the comparative restricts the possible value of the degree argument of *heavy* to be a degree of age. The adjectives *old* and *heavy* define mappings from objects to scales with different dimensional parameters, however, with the result that the partial ordering relation introduced by the adjective in the nuclear scope–the relation shown in (135)–is undefined for the degrees introduced by the comparative. The anomaly of (133) is a result of this failure to provide an appropriate value for the degree argument of *heavy*.[29]

[29] A possible criticism of this analysis is that it predicts that sentences like (133) should be contradictory, rather than anomalous, because the constraints on ordering relations described above entail that there is no degree which satisfies both the restriction and the nuclear scope in (134). Although I will assume for now that the interpretation of restricted quantification structures can be formulated in a way that avoid this criticism, I should note that the analysis of comparatives that I will develop in chapter 2, in which comparatives are *not* analyzed as quantificational expressions, avoids this criticism, because it has the consequence that the relation introduced by the comparative morpheme in examples like (133) is undefined. The most important point of the discussion in this section is that in order to construct the type of explanation outlined here in the first place, it is necessary to introduce scales and degrees into the ontology.

In general terms, incommensurability is predicted to arise whenever a comparative construction restricts the degree argument of a gradable adjective to be a degree on a scale that is distinct from the scale associated with the adjective. In such a context, the partial ordering relation introduced by the adjective in the nuclear scope is undefined, triggering anomaly. In order to avoid this anomaly, then, it must be the case that the comparative introduces a degree on the same scale as the one associated with the adjective in the nuclear scope; this will be the case only if both adjectives in the subdeletion construction have the same dimensional parameters. The scalar analysis thus derives the descriptive generalization adduced in section 1.1.3 and repeated below in (136) from basic assumptions about ordering relations and the analysis of gradable adjectives as expressions that map objects to degrees on a scale.

(136) A comparative construction is semantically well-formed only if the compared adjectives have the same dimensional parameter.

Most importantly, the anomaly of sentences like (133) is explained in terms of the *linguistic* properties of gradable adjectives and comparatives, rather than a general conceptual constraint on comparison.

This analysis also explains the fact that a comparative such as (137) must evaluate the compared objects with respect to the same dimension (see the discussion of this point in section 1.1.2).

(137) Richard is more clever than George is.

According to the analysis of incommensurability outlined above, the comparative must introduce a degree on the same scale as the one associated with the adjective in the main clause. Since scales are differentiated by their dimensional parameters, evaluating the compared objects in (137) with respect to different dimensions would amount to a violation

of this constraint. Since such a violation triggers anomaly, the interpretation of (137) that would trigger it is unavailable.

Finally, it should be noted that the analysis of incommensurability outlined here has important consequences for the analysis of non-anomalous comparatives constructed out of different adjectives, such as (138), which has the logical representation in (139).

(138) Our Norfolk Island Pine is as tall as its branches are long.

(139) $\exists d[d \geq \textsc{max}(\lambda d'.long(our\ NI\ pine's\ branches, d'))][tall(our\ NI\ pine, d)]$

The comparative in (139) restricts the possible values of the degree argument of *tall* to be some degree which is at least as great as the (maximal) degree to which our Norfolk Island pine's branches are long; in other words, the degree argument of *tall* must also be a degree of *length*. The fact that (138) is not anomalous indicates that degrees of length and degrees of tallness must be objects on the same scale. That is, it must be the case that the adjectives *long* and *tall* have the same dimensional parameters, i.e., they must map objects to the same scale. Intuitively, this seems right: both adjectives order the objects in their domain according to different aspects of a general concept of "linear extent": a vertical one in the case of *tall*, and a horizontal one in the case of *long* (cf. Cresswell 1976). The fact that the actual orderings imposed on the domains are different is simply a result of the fact that *long* and *tall* define different relations between objects and degrees. Since the relations are distinct, there is no entailment that the orderings they support should be the same.

1.3.3.2 Cross-Polar Anomaly

Like the explanation of incommensurability, the explanation of cross-polar anomaly relies crucially on the assumption that the meaning of a gradable adjective is defined in relation to a scale. Recall that cross-polar anomaly is exemplified by sentences like (140)-(141), which show that comparatives formed out of positive and negative pairs of adjectives are anomalous.

(140) #Venus is brighter than Mars is dim.

(141) #*The Dream of a Ridiculous Man* is shorter than *The Brothers Karamazov* is long.

In order to develop an explanation for this phenomenon, it is first necessary to introduce a theory of adjectival polarity into the scalar analysis. This project is undertaken chapter 3; for now, I will limit the discussion to an overview of the aspects of the analysis that provide the basis for an explanation of cross-polar anomaly.

Antonymous pairs of adjectives such as *bright/dim* and *tall/short* provide fundamentally the same kind of information about the degree to which an object possesses some gradable property (for, example, both *tall* and *short* provide information about an object's height), but they do so from complementary perspectives. Intuitively, *tall* is used either neutrally or to highlight the height an object has, while *short* is used to highlight the height an object does *not* have. In chapter 3, I use this difference in perspective to develop a theory of adjectival polarity in which positive and negative degrees are treated as distinct objects on the same scale. Specifically, I show that if degrees are analyzed as intervals on a scale, as in Seuren 1978, von Stechow 1984a, and Löbner 1990 (cf. Bierwisch 1989), rather than as points on a scale, as traditionally assumed, the facts of cross-polar anomaly can be explained. At the same time, important inferences associated with antonymous pairs of adjectives are also captured (see section 1.1.4.1). I will not attempt to go into the details of the analysis here; instead, I will outline the basic claims of the analysis and point the reader to chapter 3 for detailed argumentation in support of these claims (see also Kennedy 1997b).[30]

First, assume that antonymous pairs of positive and negative adjectives define the same mapping of objects in their domains to a shared scale, but that positive and negative

[30]The analysis of cross-polar anomaly that I will outline in this section is essentially the same as the one developed in Kennedy 1997b. The analysis to be presented in chapter 3 is the same in its basic assumptions, but differs slightly in implementation, keeping in line with the analysis of comparatives that I will present in chapter 2.

degrees are distinct objects. If this is correct, adjectival polarity can be characterized as a sortal distinction: positive adjectives denote relations between individuals and positive degrees; negative adjectives denote relations between individuals and negative degrees. These assumptions, combined with the analysis of comparatives as restricted quantification structures, provide the basis for an explanation of cross-polar anomaly as a type of sortal anomaly: the comparative restricts the degree argument of the adjective to be a degree of the wrong sort.[31]

For illustration, consider (140), repeated below with its logical representation.

(140) #Venus is brighter than Mars is dim.

(142) $\exists d[d > \text{MAX}(\lambda d'.dim(Mars,d'))][bright(Venus,d)]$

According to the analysis of the comparative outlined in section 1.3.2, (140) is true just in case for some degree *d* such that *d* exceeds the degree to which Mars is dim, and Venus is at least as bright as *d*. The distinction between positive and negative degrees developed in chapter 3 is such that in a logical representation like (142), in which the adjective in the comparative clause is negative, the only objects which satisfy the restriction imposed by the comparative are negative degrees.[32] As a result, the comparative restricts the degree argument of the positive adjective *bright* to be a negative degree. Since *bright* is requires a positive degree as its argument, the comparative introduces an argument of the wrong sort, triggering a sortal anomaly. Examples in which the polarity of the adjectives are reversed, such as (141), can be explained in exactly the same way.

Although I have postponed detailed argumentation for the claims made here until chapter 3, the current discussion makes an important point. If an explanation of cross-polar

[31]In section 3.3 of chapter 3, I show that these assumptions derive the order-reversing properties of negative adjectives discussed above, and also provide a basis for an explanation of the monotonicity properties of gradable adjectives (see the discussion in section 1.1.3).

[32]The reverse is true of examples such as (141), in which the adjective in the comparative clause is positive.

anomaly along the lines of the one I have outlined here is correct, then it provides additional support for the general hypothesis that the interpretation of gradable adjectives should be characterized in terms of scales and degrees. In order to distinguish positive and negative degrees and use this distinction as the basis for a sortal characterization of adjectival polarity, it must be the case that scales and degrees are part of the ontology, and that the interpretation of gradable adjectives is formalized in terms of such objects.[33]

1.3.3.3 Negative Adjectives and Measure Phrases

The analysis of cross-polar anomaly outlined in the previous section has the additional positive result of providing an explanation for the distribution of measure phrases. A property of the distinction between positive and negative degrees that I will motivate and develop in chapter 3 is that measure phrases such as *3 feet* can only denote positive degrees, they cannot denote negative degrees. If polarity is represented as a sortal distinction between positive and negative adjectives, then the contrast between e.g. (143) and (144) can be explained in the same way as cross-polar anomaly.

(143) Benny is 4 feet tall.

(144) #Benny is 4 feet short.

Assuming that the role of the measure phrase in an example like (143) is to denote the degree argument of the adjective, then (144) is well-formed because *4 feet* denotes a positive degree, which is of a degree of the appropriate sort for the positive adjective *tall*. In contrast, since the negative adjective *short* requires a negative degree as argument, the positive

[33] In its basic respects, the explanation of cross-polar anomaly is of the same type as the explanation of incommensurability: the comparative construction restricts the value of the degree argument of the adjective to be a degree that is incompatible with its semantic requirements. Although this incompatibility stems from a conflict in polarity, rather than a conflict in dimensional parameter, the underlying problem is the same. In both constructions, the partial ordering relation introduced by the adjective in the nuclear scope is undefined for the degree introduced by the comparative. I will return to this point in more detail in chapter 3.

degree introduced by the measure phrase triggers a sortal anomaly (see von Stechow 1984b for a similar explanation of these facts).

1.3.3.4 Comparison of Deviation

Comparison of deviation constructions, which are exemplified by sentences like (145)-(146), differ from examples of cross-polar anomaly in that they are constructed out of positive and negative pairs of adjectives, but they are not anomalous.

(145) Robert is as short as William is tall.

(146) It's more difficult to surf Maverick's than it is easy to surf Steamer Lane.

The challenge faced by an analysis of comparatives and gradable adjectives is to develop an account of these facts that both explains the unique semantic characteristics of comparison of deviation–in particular, the fact that comparison of deviation constructions entail the truth of the corresponding absolutes (see the discussion of this point in section 1.1.4.3)–and, at the same time, maintains an analysis of cross-polar anomaly. The vague predicate analysis was unable to achieve either of these goals: although it does succeed in constructing interpretations for comparison of deviation constructions, it does not account for their entailments, nor does it succeed in explaining cross-polar anomaly. In the following paragraphs, I will present an overview of how the scalar analysis succeeds in meeting this challenge. As in the earlier discussion of cross-polar anomaly, I postpone detailed argumentation until chapter 3.

One of the primary differences between the scalar analysis and the vague predicate analysis is that the former introduces an abstract representation of measurement, i.e. a scale. If scales are part of the ontology, then it should not only be possible to compare objects with respect to their measurements (i.e., to establish ordering relations between degrees, as in comparative constructions), it should also be possible to talk about *differences* between measurements. Comparatives such as (147) and (148), discussed extensively in

Hellan 1981 and von Stechow 1984a, appear to verify this prediction.

(147) Red stars are typically 5000 degrees cooler than blue stars.
(148) The space telescope's orbit is now 100 kilometers higher than it used to be.

Following von Stechow (1984a), I will refer to examples like (147) and (148) as "differential comparatives". The interesting aspect of differential comparatives is the interpretation of the measure phrases: *5000 degrees* in (147) and *100 kilometers* in (148) denote the *difference* between the projections of the compared objects on the relevant scale. For example, (148) is accurately paraphrased by (149).

(149) The current height of the space telescope's orbit exceeds its former height by 100 kilometers.

The importance of differential comparatives is that they show that it is possible to make explicit reference to "differential degrees"–degrees that measure the difference between the projections of two objects on a scale.[34] Building on the initial observations about the basic meaning of comparison of deviation constructions, we can construct an analysis of this phenomenon in terms of quantification over such differential degrees: whereas standard comparatives quantify over degrees that represent the (positive or negative) projection of an object on a scale, comparison of deviation constructions are comparatives that quantify over degrees that denotes the *difference* between the projection of an object on a scale and an appropriate standard-denoting degree. Without going into the details of how this interpretation is derived (this will be the focus of section 3.2 in chapter 3), it can be shown that this analysis satisfies the requirements outlined above.

[34]Von Stechow (1984a) claims that differential comparatives provide another argument against the vague predicate analysis, since there is no way to represent such differences without a scale. See Klein 1991 for some suggestions as to how differential comparatives could be explained within a vague predicate analysis, however.

First, if this analysis is correct, it derives the entailment properties of comparison of deviation constructions. If the comparative and equative constructions in examples like (145) and (146) compare the degrees to which two objects exceed relevant standard values, then the truth conditions for the absolute construction are satisfied whenever the truth conditions for the comparison of deviation construction are satisfied. Second, this analysis explains why comparison of deviation constructions are not anomalous. An important characteristic of differential degrees is that they are "polarity neutral", as indicated by the fact that such differential degrees can be introduced by measure phrases regardless of the polarity of the adjective that heads the comparative construction, as shown by (150) and (151) (cf. the unacceptability of (144)).

(150) Galileo was 6 inches taller than Copernicus.
(151) Copernicus was 6 inches shorter than Galileo.

According to the analysis outlined in section 1.3.3.2, cross-polar anomaly is triggered by a mismatch in the polarity of the compared degrees (where polarity is characterized in terms of a sortal distinction). If differential degrees are neutral in their polarity, and if comparison of deviation involves quantification over differential degrees, then no such mismatch is predicted to arise in the context of these constructions.

1.3.4 Summary

We are now in position to summarize the important aspects of the scalar analysis. The crucial difference between gradable adjectives and non-gradable adjectives in this approach is one of semantic type: non-gradable adjectives denote functions from individuals to truth values; gradable adjectives denote relations between objects and degrees. In other words, gradable adjectives actually have an extra "degree argument". This explains the unacceptability of non-gradable adjectives in comparatives and other degree constructions: if e.g. comparatives quantify over degrees, then the fact that non-gradable adjectives do not appear in these constructions can be explained in terms of vacuous quantification: non-gradable adjectives do not introduces a degree variable for the comparative to bind.

The fundamental differences between the scalar analysis and the vague predicate analysis are the introduction of scales and degrees into the ontology and the analysis of gradable adjectives as relational expressions. The empirical advantage of introducing scales and degrees into the ontology is that it supports explanations of the facts that were problematic for the vague predicate analysis: incommensurability, cross-polar anomaly, the distribution of measure phrases, and comparison of deviation. These facts thus provide compelling evidence in favor of the hypothesis that the semantics of gradable adjectives should be characterized in terms of scales and degrees. A separate question is whether the specific approach to adjective meanings discussed so far, in which adjectives are analyzed as relational expressions, is the correct one. An important component of this analysis is the claim that comparatives and other degree constructions are expressions that quantify over the degree argument of an adjective. Indeed, this analysis of degree constructions is closely tied to the general claim that gradable adjectives have a degree argument: if degrees are introduced by the adjective, then, like other arguments, they should support quantification. The claim that degree constructions are quantificational raises other expectations, however. Focusing on comparatives in particular, if these constructions quantify over degrees, then we would expect them to participate in scope ambiguities in contexts in which other quantificational expressions (e.g., quantified NPs) show similar ambiguities. In the next section, I will show that this expectation is not met.

1.4 Comparatives and Scope

The interpretation of quantified nominals indicates that the syntactic or semantic components of natural language contain mechanisms whose function is to associate surface strings with multiple logical representations, differing in the relative scope of quantificational (and intensional) expressions. If comparatives quantify over an argument of the adjective, then the expectation is that they should show the same sorts of interpretive

possibilities as other quantificational expressions. That is, like other quantificational expressions, we expect them to interact with other operators (negation, universal quantification, intensional operators, etc.) to trigger scope ambiguities. In the following sections, I will show that this predicted parallelism is in fact *not* observed; instead, the facts indicate that if comparatives are quantificational, then the quantificational force they introduce always has narrow scope with respect to other operators in the sentence.

In order to demonstrate this, it is necessary to make a distinction is between the *scope of the comparative* and the *scope of the comparative clause*. The "scope of the comparative" is the scope of the quantificational expression that denotes the degree argument of a gradable adjective, which is the expression that, given general assumptions about quantification over argument expressions, we expect to participate in scope ambiguities. the "scope of the comparative clause", on the other hand, is the scope of the complement of *than*–the expression that introduces the degree that provides the basis for comparison. What will emerge from the discussion is that the comparative clause *does* show scope ambiguities; in particular, it has scopal characteristics that are similar to those of a definite description, an observation that is well-established in the literature on comparatives (see in particular von Stechow 1984a for an overview of the scopal properties of the comparative clause; see also Russell 1905, Hasegawa 1972, Postal 1974, Horn 1981, Heim 1985, Larson 1988a, Kennedy 1995, 1996, and Rullmann 1995).

A final note: in the discussion that follows, I will focus on the approach to comparatives outlined in section 1.3.2, in which comparatives are analyzed as restricted existential quantification structures (see e.g., Hellan 1981, von Stechow 1984a, Heim 1985, Lerner & Pinkal 1992, 1995, Gawron 1995, Hazout 1995, Rullmann 1995). I do this for perspicuity; I will, however, extend the discussion to include semantic analyses of comparatives as generalized quantifiers in 1.4 (see Moltmann 1992a, Hendriks 1995; cf. Postal 1974, Cresswell 1976, Williams 1977).

1.4.1 Negation

In general, existentially quantified nominals are ambiguous in the scope of negation, as shown by (152).

(152) Max didn't see a recent supernova in the Hercules Cluster.

The indefinite in (152) *a recent supernova in the Hercules Cluster* can be interpreted either inside or outside the scope of negation: on the former reading, Max didn't see any supernovas in the Hercules Cluster, on the latter reading, there is a supernova in the Hercules Cluster such that Max didn't see it. Although the favored interpretation of (152) may be one in which the indefinite has narrow scope, a wide scope reading is clearly available possible, and is in fact the only possible reading when (152) is followed by a pronoun that refers to the object introduced by the indefinite, as in (153).

(153) Max didn't see a recent supernova in the Hercules Cluster because it was obscured by cosmic dust.

Now consider a comparative construction such as (154) in the context of negation.[35]

(154) Max isn't taller than his brother is.

(154) does not have an interpretation in which the existential quantification associated with the comparative takes wide scope with respect to negation. That is, only (155), not (156), represents a possible interpretation of this sentence.

[35]I focus here on negation, but the following discussion applies to examples involving negative quantifiers such as (i) and (ii) as well, as can be easily verified.

(i) No one is taller than Max is.
(ii) Few people are taller than Max is.

(155) $\exists d[d > \textsc{max}(\lambda d'.tall(Max's\ brother,d'))][tall(Max,d)]$

(156) $\exists d[d > \textsc{max}(\lambda d'.tall(Max's\ brother,d'))]\ [tall(Max,d)]$

On the reading represented by (155), (154) is true just in case it is not true that for some degree *d* which exceeds the degree of Max's brothers' tallness, Max is *d* tall; this is an accurate characterization of the interpretation of (154). On the wide scope interpretation in (156), (154) would be true if there were a degree *d* which exceeds the degree to which Max's brother is tall, and it is not the case that Max is *d* tall. If (156) were a possible interpretation of (154), then this sentence could be true in a situation in which Max's height actually *exceeds* Max's brother's height, as in the context illustrated by (157), because there is a degree–namely *d*–which exceeds the degree to which Max's brother is tall, and it is not the case that Max is at least as tall as that.

(157) HEIGHT: $0 \text{———} d_{M's\ brother} \text{———} d_{Max} \text{———} d \longrightarrow \infty$

One response to this problem is that the wide scope reading of the comparative is a tautology: assuming there is no maximal degree of height (cf. von Stechow 1984b), it is always true that some degree which satisfies the restriction–*d is greater than Max's brother's height*–will also satisfy the nuclear scope–*Max is not at least as tall as d*. This response allows for the possibility that (154) is ambiguous, but claims that since the tautologous interpretation does not provide useful information about the actual relation between Max's height and his brother's height, it is ignored.

There are at least two problems with this explanation. The first is that even in contexts in which tautologous interpretations of sentences are strange or even anomalous, they are nevertheless *detectable*. Consider, for example, a sentence like (158) (see Lakoff 1970, Huddleston 1971, Carden 1977, Manaster-Ramer 1978, and Horn 1981 for discussion of sentences of this type).

(158) We were amazed that the Hale-Bopp comet was as bright as it was.

(158) allows a "sensible" reading, in which the actual brightness of the Hale-Bopp comet, in contrast to what was expected, is responsible for our amazement. It also has a "strange" interpretation in which we were amazed that the brightness of the comet was equal to (or at least as great as) the brightness of the comet. The strangeness of this interpretation clearly stems from the fact that the embedded equative is interpreted tautologously. What is crucial to the current discussion is that this reading is detectable; in contrast, the predicted tautologous reading of (154) is *not* detectable.

The second problem with this sort of explanation is that there is reason to believe that it is too weak. Ladusaw (1986) observes that two types of semantic filtering can be defined. A sentence can be semantically ill-formed either because its interpretation fails to meet conditions of informativity or because it has no interpretation at all. The explanation of the unavailability of wide-scope interpretations of comparatives with respect to negation outlined above is an example of the first type of filtering: the reading of (154) represented by (156) is unavailable because it is uninformative. A general characteristic of this type of filtering, however, is that it allows for the possibility that there are contexts in which the generally unavailable interpretation becomes available–indeed, the fact that tautologous equatives embedded under factive predicates permit a "sensible" interpretation, as described above, is a case in point.[36]' The problem is that there is no context that supports a reading of (154) with the truth conditions in (156): this sentence simply cannot be used to describe the situation in (157), which is what such a reading should allow. The conclusion, then, should be that a wide-scope readings of a comparative under negation is ruled out by the second type of filtering mentioned by Ladusaw: the grammar of comparatives simply does not permit such interpretations to arise. If comparatives are quantificational expressions,

[36]See also Barwise and Cooper's (1981) discussion of definiteness and the English existential construction.

however, then this result can be achieved only by stipulating that the existential quantification introduced by the comparative must take narrow scope (see e.g. von Stechow 1984a and Rullmann 1995 for proposals to this effect).

1.4.2 Distributive Quantifiers

A second context in which existentially quantified NPs typically show scope ambiguities is in sentences containing distributive quantifiers like *every, most, each,* etc. For example, (159) has the two readings characterized in (160)-(161), which are distinguished by the relative scope of the existential and universal quantifiers.

(159) Every student in Semantics 1 read a book on adjectives.

(160) $\forall x[\mathit{student\text{-}in\text{-}semantics\text{-}1}(x)][\exists y[\mathit{book\text{-}on\text{-}adjectives}(y)][\mathit{read}(x,y)]]$

(161) $\exists y[\mathit{book\text{-}on\text{-}adjectives}(y)][\forall x[\mathit{student\text{-}in\text{-}semantics\text{-}1}(x)][\mathit{read}(x,y)]]$

The crucial difference between (160) and (161) is that the former allows books to covary with students, while the latter requires that there be one book that every student read.

Comparatives in the scope of distributive quantifiers do not show a similar pattern of ambiguity. Consider (162).

(162) Every planet in the solar system is larger than Earth's moon.

If comparatives participate in scope ambiguities, then (162) should have the two logical representations in (163) and (164).

(163) $\forall x[\mathit{planet}(x)][\exists d[d > \textsc{max}(\lambda d'.\mathit{large}(\mathit{Earth's\ moon},d'))][\mathit{large}(x,d)]]$

(164) $\exists d[d > \textsc{max}(\lambda d'.\mathit{large}(\mathit{Earth's\ moon},d'))][\forall x[\mathit{planet}(x)][[\mathit{large}(x,d)]]$

(163) correctly captures the interpretation of (162): for every x such that x is a planet, for

some degree d such that d exceeds the degree to which Earth's moon is large, x is d-large. At first glance, the interpretation in which the comparative has wide-scope appears to be equivalent. On the interpretation in (164), (162) is true just in case for some degree d such that d exceeds the degree to which Earth's moon is large, every planet is d-large. Although it appears on the surface that (164) should be true in the same contexts as (163), this is contingent on an additional assumption. (163) and (164) are equivalent only if the adjective in the nuclear scope of the comparative (i.e., the absolute form) introduces a partial ordering relation, rather than a relation of equality, as specified in (165) (recall from the discussion in section 1.3.1 that δ_φ is a function that maps objects onto the scale associated with φ).

(165) $||\varphi(a,d)|| = 1$ iff $\delta_\varphi(a) = d$

If the absolute were interpreted as in (165), however, (163) and (164) would have very different truth conditions. On the interpretation in (164), (162) would be true just in case every planet were large to the same degree (see Rullmann 1995 for discussion).

The fact that (162) does not have this type of interpretation appears to provide support for the hypothesis that the absolute form of a gradable adjective is stated in terms of a partial ordering relation, a hypothesis that up to now I have assumed without justification. Additional support for this hypothesis appears to be provided by exchanges like (166).

(166) A: You have to be at least 5 feet tall to be an astronaut.
B: I'm 5 feet tall; in fact, I'm over 5 feet tall.

Although the typical interpretation of a sentence such as *Benny is 5 feet tall* is one in which Benny is assumed to be *exactly* 5 feet tall, the fact that B's utterance in (166) is not contradictory suggests that the "exactly" interpretation is not an entailment, but rather a scalar implicature.[37]

[37]See Carston 1988 and Horn 1992 for relevant discussion of this issue.

If the absolute form is interpreted in this way, however, it presents a serious problem for the analysis of comparatives with *less*. Specifically, an analysis of *less* comparatives in terms of existential quantification ends up with the wrong truth conditions. Consider, for example, (167), which has the logical representation in (168).

(167) That red giant is less dense than the black hole it's orbiting.

(168) $\exists d[d < \textsc{max}(\lambda d'.dense(\textit{the black hole}, d'))][dense(\textit{the red giant}, d)]]$

According to (168), (167) is true just in case for some degree d such that d is exceeded by the maximal degree to which the back hole is dense, the red giant is at least as dense as d. The problem is that these truth conditions actually allow for the possibility that the density of the red giant exceeds that of the black hole, a result that is clearly incompatible with the actual meaning of (167). In the unlikely situation represented in (169), in which d_b denotes the density of the black hole and d_r denotes the density of the red giant, (167) would be true, because there is a degree–d_1–that satisfies the conditions imposed by (168): d_1 is exceeded by d_b, and d_r is at least as great as d_1.

(169) DENSITY: $\text{———} \; d_1 \; \text{———} \; d_b \; \text{———} \; d_r \; \text{———}\!\!\rightarrow$

If, on the other hand, the absolute construction is interpreted with respect to a relation of equality, as in (165), this problem disappears (this point is made in Rullmann 1995): (167) would be false in the context shown in (169), because there is no degree that is both exceeded by d_b and is equal to d_r.

Note that the problem of an "at least as" interpretation of the absolute for comparatives with *less* cannot be solved pragmatically. The fact that (170) is a contradiction indicates that unlike the case with the absolute, the "maximality" of *less* comparatives is not due to scalar implicature (cf. (171)).

(170) #That red giant is less dense than the black hole it's orbiting.; in fact, it's denser than the black hole.

(171) Some black holes are extremely dense; in fact, all black holes are extremely dense.

We thus arrive at a paradox. In order to construct an analysis of comparatives with *less* that has the correct truth conditions, it is necessary to assume that the absolute form is interpreted with respect to a relation of equality, as in (165), rather than a partial ordering relation, as I have assumed up to now. If the ordering relation associated with the absolute construction is one of equality, however, then a quantificational analysis of comparatives predicts that a sentence like (162) should be ambiguous; in particular, it should have a reading in which all the planets in the solar system are claimed to have the same size. (162) cannot be interpreted in this way, however. The conclusion to be drawn from these facts, then, is that if the quantificational analysis outlined here is to correctly capture the truth conditions of comparatives with *less*, it must also stipulate that the existential force introduced by the comparative has narrow scope with respect to distributive quantifiers, just as we saw with negation in the previous section.[38]

1.4.3 Intensional Contexts

Although the discussion in the previous two sections illustrated that comparatives do not show scope ambiguities with respect to negation and distributive quantifiers, there are contexts in which comparatives clearly *are* ambiguous. The best known example of such contexts involves contradictory comparatives in intensional contexts, as in (172) (see Russell 1905, Hasegawa 1972, Postal 1974, Hankamer and Sag 1976, Williams 1977, Hellan 1981, Napoli 1983, von Stechow 1984a, Hoeksema 1984, Heim 1985, Larson 1988a, Kennedy 1995, 1996b and others for discussion).

(172) Max thinks the moon is larger than it is.

As originally observed by Russell (1905), a sentence like (172) is ambiguous between the reading paraphrased in (173), in which it is asserted that Max is mistaken about the size of the moon, and the one in (174), in which it is claimed that Max believes a contradiction.

(173) The size that Max thinks the moon is exceeds the size that it actually is.

(174) Max thinks that the size of the moon exceeds the size of the moon.

If comparatives are quantificational, then this contrast is expected, because (172) is assigned the two logical representations in (175) and (176), which differ in the relative scope of the comparative and the intensional verb *think*.

(175) $\exists d[d > \textsc{max}(\lambda d'.large(moon,d'))][think(Max,\hat{}large(moon,d))]$

(176) $think(Max,\hat{}\exists d[d > \textsc{max}(\lambda d'.large(moon,d'))][large(moon,d)])$

The interpretation of (172) in (175) accurately characterizes the reading in (173), while the interpretation in (176) corresponds to the reading in (174). On the surface, then, the ambiguity of (172) appears to provide support for the hypothesis that comparatives are quantificational.

Problems arise for this analysis when attention is expanded to include comparatives in other intensional contexts. Consider, for example, the interpretation of contradictory comparatives in counterfactuals, as in (177) (cf. Postal 1974, von Stechow 1984a).

[38] Carston (1988) argues that numerals can have an "exactly" reading in some contexts and an "at least as" reading in others, suggesting ultimately that their interpretations are indeterminate. Even if Carston's analysis can be carried over to the absolute form of gradable adjectives, it does not resolve the problems I have laid out here, since it allows for the possibility of an "exactly" interpretation of the absolute in examples like (162) and an "at least as" interpretation in examples like (170), predicting that both should be ambiguous. The optimal solution to this problem, then, is one in which the interpretation of the comparative forms is not stated in terms of the interpretation of the absolute. This is exactly the type of analysis that I will develop in chapter 2.

(177) If Jones had been taller than he was, he would have been decapitated by the flying saucer.

(177) is ambiguous in the same way as (172): it has a "sensible" interpretation, in which the antecedent of the conditional describes a relation between Jones' height in some alternative world and Jones' height in the actual world, and a trivial interpretation, in which the antecedent of the conditional is contradictory. The problem presented by counterfactuals like (177), as pointed out by Von Stechow (1984a), is that an analysis of this sentence along the lines of the one given for (172) above fails to get the right truth conditions.

I assume the semantics for the counterfactual is as in (178), where the symbol "⇒" indicates counterfactuality (cf. Lewis 1973a, Stalnaker 1968).

(178) $||\varphi \Rightarrow \psi||^w = 1$ iff (i) there are no possible worlds in which φ is true or (ii) there is a world w' in which φ and ψ are true and w' is closer to w than any world in which φ holds but ψ does not.

Building on the analysis of (172), the non-trivial interpretation of (177) is assigned the logical representation in (179).

(179) $\exists d[d > \textsc{max}(\lambda d'.tall(jones,d'))][tall(Jones,d) \Rightarrow decapitate(flying\ saucer, jones)]$

Now consider a context in which the set of worlds consists of the five worlds in (180) where the specified facts obtain and w_0 is the world in which the sentence is interpreted.

(180)

WORLD	JONES' HEIGHT	DECAPITATED?
w_0	5'	no
w_1	5' 1"	no
w_2	5' 2"	no
w_3	5' 3"	no
w_4	5' 4"	yes

In this context, (179) is true, because there is a degree of height–the one corresponding to 5' 4"–that exceeds Jones' actual height and also satisfies the truth conditions of the counterfactual: there is no world closer to w_0 than w_4 in which it is true both that Jones is at least 5' 4" tall and that he is decapitated by the flying saucer. The problem is that (177) makes a stronger claim that this: (177) asserts that *any* increase in height would have resulted in Jones being decapitated. As a result, the truth conditions associated with (179) are too weak.

Von Stechow concludes from examples like this that it is not the entire comparative construction that interacts with intensional operators to generate scope ambiguities, rather it is the *comparative clause* that triggers scope ambiguities in intensional contexts. According to von Stechow, the actual logical representation of the non-trivial interpretation of (177) is something like (181), in which the comparative clause ($\textsc{max}(\lambda d'.tall(jones,d'))$) has scoped out of the conditional, but the quantifier introduced by the comparative remains inside.

(181) $\lambda d_x[\exists d[d > d_x][tall(Jones,d)] \Rightarrow decapitate(flying\ saucer, jones)](\textsc{max}(\lambda d'.tall(jones,d')))$

(181) requires it to be the case that every world in which Jones' height exceeds his height in the actual world is such that he is decapitated. Therefore, unlike (179), (181) is false in the context illustrated by (180).

Von Stechow motivates this analysis by observing that the scopal properties of the comparative clause are entirely expected if it is a type of definite description, as originally claimed by Russell (1905) (and as assumed in this thesis; see the discussion of this point in section 1.3.2). What is relevant to the current discussion is that if von Stechow's analysis is correct in its basic claims, then facts like (172) do not provide evidence in favor of a quantificational analysis of the comparative, since the range of scope ambiguities of comparatives in intensional contexts can be explained if we assume that it is the comparative clause (qua definite description) that is responsible for the ambiguities.[39]

1.4.4 Summary

The starting point of this section was the observation that a basic prediction of an analysis in which comparatives involve quantification over degrees is that they should show scope ambiguities relative to other operators. The facts discussed here show that in contexts involving negation, distributive quantifiers, and counterfactual conditionals, this prediction is not borne out: there is no evidence that the quantificational force introduced by the comparative interacts with other operators to generate scope ambiguities. At the same time, the interpretation of comparatives in intensional contexts, in particular, the interpretation of comparatives in counterfactual conditionals, shows that the comparative clause *does* participate in scope ambiguities, a result that is expected if it is a type of a definite description.[40] We thus arrive at a somewhat paradoxical generalization: the comparative does not show scope ambiguities; the comparative clause does. In effect, the comparative clause is behaving as if it, rather than the comparative construction, were the argument-denoting expression. In chapter 2, I will develop an alternative analysis of gradable adjectives and comparatives that makes exactly this claim; before I conclude this chapter, however, I

[39]Indeed, this was Russell's original analysis of these facts.

[40]Note that if the comparative clause is a definite description, then the contexts in which it triggers scope ambiguities should be limited to those involving intensional expressions; it should not interact with quantificational determiners or negation to generate scope ambiguities. As observed in footnote 28, this does seem to be the case.

will take a brief look at an alternative quantificational approach, showing that it also makes incorrect predictions with respect to scope ambiguities.

1.4.5 Comparatives as Generalized Quantifiers

An alternative to the analysis of comparatives as existential quantification structures is one in which they are analyzed as generalized quantifiers (see Moltmann 1992a and Hendriks 1995; see also Postal 1974, Cresswell 1976, and Williams 1977 for very similar accounts). On this view, comparatives do not introduce a degree, but rather denote relations between sets of degrees (where one set is introduced by the comparative clause and the other is derived by abstracting over the degree argument associated with the adjective).

The basic form of the analysis is as follows. Assume that degree morphemes such as *er/more* and *less* denote determiners in the sense of Barwise and Cooper 1981; i.e., relations between sets. Unlike determiners in the nominal domain, which denote relations between sets of individuals, comparative determiners denote relations between sets of degrees. The interpretation of *more* on this view is shown in (182), where φ and ψ correspond to the comparative clause and main clause respectively, and denote sets of degrees.

(182) $||\textsc{more}(\varphi)(\psi)|| = 1$ iff $||\varphi|| \subset ||\psi||$

For an illustration of this type approach, consider the analysis of (183), which has the interpretation shown in (184) (cf. Moltmann 1992a).

(183) Jupiter's atmosphere is thicker than Titan's atmosphere.

(184) MORE($\lambda d.$*thick*(*Titan's atmosphere,d*))($\lambda d.$*thick*(*Jupiter's atmosphere,d*))

According to (184), (182) is true just in case the set of degrees which makes *Titan's atmosphere is d-thick* true is properly included by the set of degrees which makes *Jupiter's atmosphere is d-thick* true. Assuming the semantics of the absolute to be stated as in (185) (cf. the analysis of the absolute in section 1.3.1), (183) is true just in case the set of degrees which are ordered below the degree of thickness of Titan's atmosphere on a scale of thickness (inclusive) is a proper subset of the set of degrees which are ordered below the degree of thickness of Jupiter's atmosphere (inclusive).

(185) $||\varphi(a,d)|| = 1$ iff $\delta_{\varphi}(a) \geq d$, where δ is a function from objects to degrees.

Since scales are formalized as totally ordered sets of degrees, the truth conditions for the comparative in (182) entail that the degree of thickness of Titan's atmosphere exceeds the degree of thickness of Jupiter's atmosphere. An immediate positive result of this analysis is it avoids problems associated with the analysis of comparatives with *less* that were discussed in the previous section. Since the arguments of the comparative morpheme are sets which contain *all* of the degrees ordered below an object's projection on a scale, the "maximality" required to capture the correct truth conditions of *less* comparatives is derived.

There are two problems with the analysis of comparatives as generalized quantifiers, however. The first is essentially the same as the one discussed in section 1.4.1: like an analysis of comparatives as existential quantification structures, a generalized quantifier account predicts that comparatives should show scope ambiguities which do not actually exist. Whereas the unavailable readings discussed in section 1.4.1 were tautologies, the generalized quantifier approach predicts comparatives in certain scopal contexts to have *contradictory* interpretations. I will focus only on the case of negation to illustrate this point.

Given the basic assumptions outlined above, a sentence like (186) should have the two interpretations in (187) and (188).

(186) Titan's atmosphere is thicker than Jupiter's atmosphere.

(187) MORE($\lambda d.thick(Jupiter's\ atmosphere,d)$)($\lambda d.thick(Titan's\ atmosphere,d)$)

(188) MORE($\lambda d.thick(Jupiter's\ atmosphere,d)$)($\lambda d.\ \ thick(Titan's\ atmosphere,d)$)

(187) accurately captures the truth conditions for (186), but the formula in (188) is a contradiction: the denotation of the scope clause is the set of degrees which make *Titan's atmosphere is d-thick* false; this is the set of degrees which are ordered above the degree of Titan's atmosphere's thickness. Because of the ordering of the scale, it will never be the case that the set of degrees denoted by the restrictor clause–the degrees ordered below the degree of Jupiter's atmosphere's thickness (inclusive)–is properly included in the set denoted by the scope clause, so (188) is contradictory. (186) does not have a contradictory interpretation, however, any more than it has a tautological one.

The second problem with the analysis of comparatives as generalized quantifiers involves a fundamental principle of natural language determiners: *conservativity*. Conservativity, defined in (189), is a property shared by all natural language determiners (see Barwise and Cooper 1981, Keenan and Stavi 1986).

(189) A determiner D is conservative iff: $D(\varphi)(\psi) \Leftrightarrow D(\varphi)(\varphi \cap \psi)$

If the interpretation of *more* is as defined in (182), then according to(189) it is not conservative, because the equivalence in (190) does not hold (cf. Gawron 1995).

(190) $\varphi \subset \psi \Leftrightarrow \varphi \subset (\varphi \cap \psi)$

A possible response to this objection is that the semantic analysis of *more* in (182) is incorrect; its interpretation is actually that in (191), in which the main clause provides the restriction and the comparative clause the scope.

(191) $||$MORE$(\varphi)(\psi)|| = 1$ iff $||\varphi|| \supset ||\psi||$

If (191) is the correct analysis of *more*, then it is conservative, because the equivalence in (192) holds:

(192) $\varphi \supset \psi \Leftrightarrow \varphi \supset (\varphi \cap \psi)$

This analysis runs into problems with *less* however. The interpretation of *less* is the inverse of the interpretation of *more*, as shown by the validity of examples like (193).

(193) Titan's atmosphere is less thick than Jupiter's atmosphere if and only if Jupiter's atmosphere is thicker than Titan's atmosphere.

If (191) is the actual the interpretation of *more,* then the interpretation of *less* must be as in (194), in which case *less* is nonconservative.

(194) $||\text{LESS}(\varphi)(\psi)|| = 1$ iff $||\varphi|| \subset ||\psi||$

The bottom line is that if conservativity is a constraint on quantificational determiners in natural language, then degree morphemes can't be quantificational determiners, because they are nonconservative.

As the discussion here shows, an analysis of comparatives as generalized quantifiers not only leads to the same puzzles as the degree description account regarding scopal interpretation, it also forces us to abandon the assumption that all quantificational determiners in natural language are conservative. At the same time, this analysis contains a very intuitive analysis of degree morphemes as relational expressions. In chapter 2, I will develop an analysis of degree constructions that is similar to the generalized quantifier approach in this respect, but differs in that it does not analyze degree morphemes as quantificational determiners, and so avoids the problems of unattested scope ambiguities and nonconservativity.

1.5 Conclusion

The primary conclusion of this chapter is that the semantic analysis of gradable adjectives should be stated in terms of abstract representations of measurement, i.e., scales and degrees. This conclusion was arrived at by showing that a number of facts, including incommensurability, cross-polar anomaly, the distribution of measure phrases, and the interpretation of comparison of deviation constructions, receive a natural explanation only if scales and degrees are introduced into the ontology and the interpretation of gradable adjectives is characterized in terms of these abstract objects.

At the same time, a number of facts called into question the traditional scalar analysis of gradable adjectives as relational expressions and comparatives as expressions that quantify over degrees. The data discussed in sections 1.4 clearly showed that the quantification hypothesized to be introduced by the comparative construction does not interact with other quantificational expressions to generate scope ambiguities. It follows that if comparatives are analyzed as quantificational structures, then it is necessary to stipulate that they always have narrow scope with respect to other operators. This is an undesirable result, raising the following question: is the analysis of comparatives in terms of quantification over degrees correct, or should we look for an alternative analysis which does not introduce unrealized expectations regarding scope ambiguities.

One alternative to a quantificational analysis of comparatives (and other degree constructions) would be one in which the degree morpheme and gradable adjective combine directly to form a property-denoting expression which can be applied to the subject. Such an analysis would reject the claim that comparatives are quantificational, eliminating the expectation that they should show scope ambiguities. Exactly this type of approach will be developed in chapter 2, but it will require modifying a basic assumption about the interpretation of gradable adjectives. Specifically, the claim that gradable adjectives have a degree *argument* will be denied, in favor of an analysis in which gradable adjectives denote functions from objects to degrees.