

Vagueness and linguistics

Robert van Rooij*

1 Introduction

An expression is vague, if its meaning is not precise. For vagueness at the sentence-level this means that a vague sentence does not give rise to precise truth conditions. This is a problem for the standard theory of meaning within linguistics, because this theory presupposes that each sentence has a precise meaning with respect to each context of use. The philosophical discussion on ‘vagueness’ concentrates on the notion of *tolerance*. A expression is vague, or has a tolerant meaning, if it is insensitive to small changes in the objects to which it can be meaningfully predicated. The discussion of ‘vagueness’ in linguistics mostly focusses on the interpretation of so-called ‘gradable adjectives’. Within that class a difference is made between relative adjectives like ‘tall’ and absolute adjectives like ‘flat’. An important difference between these two types of adjectives is that in contrast to relative adjectives, absolute adjectives allow for *natural precisifications*: if we fix a level of *granularity*, relative adjectives are still vague, but absolute adjectives are not. Still, also absolute adjectives give rise to vagueness. This suggests that vagueness also has something to do with what a natural, or appropriate, precisification is.

In this paper I first discuss the nature of vagueness, and contrast it with notions such as ambiguity and context-dependence. In section 3 I briefly discuss some reasons that could perhaps explain *why* vagueness is such a pervasive phenomenon in natural language. Section 4 reviews some more or less standard linguistic analyses of gradable adjectives. I will concentrate myself on approaches that take *comparison classes* into account. Because comparative constructions are ideally formed in terms of gradable adjectives, comparative ordering relations will be discussed as well. In section 5 it will be argued that one specific ordering relation is crucial for any analysis of vagueness that

*I would like to thank Johan van Benthem, David Etlin, Ewan Klein, Frank Veltman, and Tim Williamson for commenting on (sometimes much) earlier versions of this paper. Thanks also to Chris Kennedy, Manfred Krifka, and Roger Schwarzcild for discussion.

wants to capture the notion of ‘tolerance’: semi-orders. I will focus my attention here on contextuallist’ approaches that want to account for the Sorites paradox, because these approaches are most popular within linguistics. In the final main section, I concentrate on what some people have called ‘loose talk’. The main issue here is whether with loose use of language we say something that is strictly speaking false, but true enough in the particular conversational setting, or true, because the conversational setting loosens the requirements for a sentence to be true.

As a warning to the reader, I should emphasize that this paper is not just an unbiased overview of work in linguistics on vagueness. Especially the extensive use of comparison classes throughout the paper, the view that positive uses of adjectives are primary to their comparative use, the claim that semi-orders are crucial to model tolerance, and the ordering relation between structures used to model coarser-grained talk are at best not very common.

2 What is vagueness?

Vagueness is a pervasive feature of natural language. Members of almost any lexical category can be vague. Prototypical vague expressions are *adjectives* like ‘tall’, ‘fast’, ‘red’, and ‘adolescent’. The Sorites Paradox is the hallmark of vagueness and formulated in terms of a *noun*, ‘heap’. But also many *adverbs* (‘very’, ‘rather’, ‘probably’, ‘softly’, ‘well’) and *quantifiers* (‘many’, ‘a lot’, ‘a few’) give rise to vagueness. In fact, no linguistic expression whose meaning involves perception and categorization can be entirely free of vagueness. This is true for *proper names* and *definite descriptions* (‘Amsterdam’, ‘the border between Belgium and the Netherlands’), *verbs* like ‘start’, ‘finish’ and ‘understand’, but also for more abstract linguistic categories such as *tense* (past or future) and *aspect* (perfective or imperfective). If a vague term occurs in a complex expression, this complex expression is often vague as well. Because ‘very’ and ‘a heap’ are vague, the expressions ‘very sick’ and ‘not a heap’ are vague too. Some expressions turn vague expressions in complex expressions that are less vague. A measure phrase is a prototypical example (turning ‘tall’ into ‘3 feet tall’). Other expressions have the opposite effect: while ‘2 o’clock’ is not vague, when we combine it with a hedging expression like ‘approximately’, ‘about’, ‘almost’, ‘roughly’, etc. it becomes vague. Lakoff (1973) gives a list of more than 60 hedging expressions, and discusses in what sense they differ in meaning.

Whether one is tall depends on a *unique* gradient contingent fact, one’s height. Vagueness of being tall is then due to the fact that it is unclear whether one’s height counts as being tall or not. There is another reason for vagueness,

though. Whether one is a clever person depends, intuitively on more than one factor. If one scores well in some relevant respects but not in others one can be a borderline clever person. As discussed by Lakoff (1973), one can reduce this vagueness by using so-called *hedging* expressions like ‘In some respects’, ‘to some extent’, and ‘in a sense’.

With some vague content words (‘red’) we associate a *prototype*, i.e., a typical representative. For those expressions, membership is a matter of prototype resemblance. But having a prototypical representative is neither a necessary (‘tall’, ‘heavy’), nor a sufficient (‘bird’, ‘grandmother’) condition for being vague. Adverbs and quantifiers don’t have prototypes at all, and an adjective like ‘tall’ presumably has no prototype because there is in general no natural upper bound to how tall things can be (Kamp & Partee, 1995). Although a penguin is a much less prototypical bird than a robin, it is not less of a bird.

Vagueness in linguistics is a problem about the *meaning* of linguistic expressions. It seems natural to assume that to be a competent speaker of English one has to know what it means for ‘John came’ to be true or false. So, a minimal requirement for any theory of meaning seems to be that one knows the meaning of a declarative sentence if one knows under which circumstances it is, or would be true. According to truth conditional semantics – the most successful and productive linguistic theory of meaning we know so far – we should stick to this minimal requirement: identify the meaning of a declarative sentence with the *conditions*, or *circumstances* under which the sentence is true.

The phenomenon of vagueness poses a problem, or threat, to this initially appealingly simple picture of meaning.¹ Even if we know that John is 1.80 meters long, it is still not clear whether we should count ‘John is tall’ as being true or as being false. The phenomenon of *ambiguity* seems to pose a similar threat. The noun ‘bank’, for instance, can be interpreted in several ways: it is perhaps used most often to denote a financial institution, but it can also be used to denote the edge of a river (ambiguity). Assuming that the meaning expressed by a sentence is via Frege’s *principle of compositionality* determined by the meanings of its constituent expressions (and the way they are combined syntactically), the sentence ‘John went to the bank’ can be used to express (at least) two very different meanings. There is an important difference between ambiguity and vagueness: an expression is ambiguous when it has more than one semantically unrelated meaning, and thus tends to come with separate dictionary entries. This in contrast with expressions that are vague.² Based on this intuition, Lakoff (1970) suggested a test to distinguish ambiguity from

¹There are other problems as well, but they are not relevant for this paper.

²At the level of individual words, the contrast between ambiguity and vagueness can also be denoted by the distinction between *homonymy* and *polysemy*.

vagueness and underspecification based on V(erb) P(hrase) ellipsis. Suppose a sentence S has a finite number of interpretations. Lakoff proposed that we should continue the sentence with phrases like ‘and Mary did too’ or ‘and Mary isn’t either’. In case this new larger sentence has equally many interpretations as S itself, sentence S is ambiguous. Otherwise, the sentence is vague or underspecified. As an example, take the sentence ‘John walked to the bank’ containing the word ‘bank’. Suppose that this sentence has two interpretations: (i) John walked to the building in which the financial institution is housed, and (ii) John walked to the edge of the river. Intuitively, the whole sentence ‘John walked to the bank and Mary did too’ also has two interpretations: (i) both John and Mary walked to the building in which the financial institution is housed, and (ii) both John and Mary walked to the edge of the river. In particular, it doesn’t have an extra reading saying, for instance, that John walked to the edge of the river while Mary walked to the building in which the financial institution is housed. Thus, ‘John walked to the bank’ is ambiguous. The reason is that the two meanings of ‘bank’ are not semantically related at all, so ‘Mary did too’ has to pick up the original meaning. Now consider the sentence ‘John is not a bachelor’. Suppose that being a bachelor means that you have to be human, male, adult, and unmarried. Thus, our sentence can be true for four different reasons, and might thus be interpreted in four different ways. Now continue the sentence with ‘and Mary isn’t either’. Lakoff’s test predicts that the original sentence was not ambiguous, because the new sentence can be true if John is not a bachelor because he is not an adult, while Mary is not a bachelor because she is not male. Thus, the whole sentence might be true for more than four different reasons.

Although ‘John is not a bachelor’ is not ambiguous, it is also not vague. Rather, it is a *general* sentence, on a par with an example like ‘Katrin received a degree’. Just like the truth of the latter sentence does not specify the particular type of degree Katrin received, the truth of the former does not specify the reason why John is not a bachelor. In contrast to vague sentences, however, general sentences have determinate truth conditions, and do not pose a threat to truth conditional semantics.

Vagueness should also be contrasted with *context dependence*. Whether what is expressed by a sentence like ‘I am Robert’ is true or false obviously depends on who (of potentially infinitely many persons) utters it, a context dependent fact. Vagueness and context dependence are in principle independent properties, although they often co-occur. *Left* and *right* are context-dependent but not (very) vague, whereas nouns like *vegetable* and *bush* are vague but not (very) context dependent (Kamp & Partee, 1995). The fact that natural language is context dependent *complicates*, but does *not threaten* truth-conditional semantics.

A traditional way of thinking about vagueness is in terms of the *existence of borderline cases*. John is a borderline case of a tall man, if the sentence ‘John is a tall man’ is neither (clearly) true nor (clearly) false. The three-valued logic account of vagueness, as well as the core supervaluation theory-account³ are based on exactly this idea. These theories assume that predicates like ‘tall’ and ‘bald’ do not give rise to a two-fold, but rather to a three-fold partition of objects: the positive ones, the negatives ones, and the borderline cases. But proponents of fuzzy logic like Lakoff (1973), and authors like Wright (1975), Kamp (1981) have argued that the existence of borderline cases is inadequate to characterize vagueness: although by assuming a three-fold instead of a two-fold distinction one rightly rejects the existence of a clear border between the positive and the negative cases, one still assumes the existence of an equally unnatural border between, for instance, the positive and the borderline cases. What seems to characterize vagueness, instead, is the fact that the denotation of vague terms lack sharp boundaries. In the words of Sainsburry (1991), they are ‘boundaryless’: there is no sharp boundary that marks the things which fall under it from the things that do not, and no sharp boundary which marks the things which do definitely fall under it from those which do not definitely, and so on. Instead, all these boundaries are blurred. Fuzzy logicians model this by assuming an *infinite*, rather than a *finite* set of truth values. Arguably, however, this is not enough: even if we assume infinitely many truth values, there still has to be a sharp boundary between, for instance, those objects that do and those objects that don’t have the relevant property *P* to degree 1.⁴ According to Wright (1975) and Kamp (1981), instead, vagueness gives rise to *tolerance*: a vague predicate is insensitive to very small changes in the objects to which it can be meaningfully predicated. Obviously, on such a view, vagueness is intimately related with the Sorites paradox.

It is clear that the existence of borderline cases does not automatically lead to vagueness (what is vague about a clear three-fold distinction?), but it is not so clear that only predicates that give rise to tolerance are vague. Gaifman (1997) discusses the predicate ‘large number of fingers’. Because changes in fingers are in discrete units, there is hardly any scope for tolerance, and one finger can make the difference. Still, 5 and 6 are borderline cases of ‘large number of fingers’, and the predicate seems vague. One might propose that only those expressions give rise to tolerance that can be modified by a hedge,

³Roughly speaking, with the ‘core’ supervaluation theory account of vagueness, I mean the theory without something like Fine’s (1975) treatment of higher-order vagueness.

⁴Fuzzy logicians might respond by saying that with respect to a relative adjective like ‘tall’, no individual has this property to degree 1. But then, as noted by Williamson (p.c.), what to do with the sharp boundary between those objects that do and those objects that don’t have the relevant property *P* to degree > 0.5 ?

such as *about*, *sort of*, or *somewhat*. An obvious consequence of this suggestion is, however, that tolerance is the rule, rather than the exception.

Vagueness is standardly opposed to *precision*. Just as gradable adjectives like ‘tall’ and a quantity modifier like ‘a lot’ are prototypical vague expressions, mathematical adjectives like ‘rectangular’, and measure phrases like ‘1.80 meter’ are prototypically precise. But what does it mean for these latter expressions to be precise? On first thought it just means that they are precise, because they have an exact mathematical definition. However, if we want to use these terms to talk about observable objects, it is clear that these mathematical definitions would be useless: if they exist at all, we cannot possibly determine what are the rectangular objects in the precise geometrical sense, or objects that are exactly 1.80 meters long. For this reason, one allows for a margin of measurement error, or a threshold, in physics and other sciences. Notice, however, that once we allow for a margin of error, we could almost immediately construct a Sorites-series, meaning that adjectives like ‘rectangular’ and measurement phrases like ‘1.80 meter’ give rise to tolerance. According to Wright, this should mean that these phrases are vague. Pinkal (1995) stressed that our use of ‘precise’ measure phrases in natural language give rise to tolerance just like in the language of physics, but now to a much larger extent. But this means that it is not clear what is left of the absolute opposition between precision and vagueness.

A central problem posed by the Sorites paradox is whether vague predicates really give rise to inconsistency. Another conceptual problem is why vagueness is so pervasive in natural languages. A linguistic, or logical, problem is that although sorite predicates give rise to vagueness, it seems that sorite predicates are not distinct from non-vague predicates with respect to their inferential power: from ‘John is very tall’, for instance, we conclude to the truth of ‘John is tall’ and they give rise to *valid syllogisms*: ‘If some linguists are tall, and every linguist is smart, then some smart people are tall.’ Moreover, vague terms can be used to express uncontroversial true statements like ‘If John is an adult, he is a man’ and vague adjectives like ‘tall’ are typically used in *comparatives* that give rise to precise truth conditions. Moreover, we can reason with that: If John is taller than Mary, and Mary is taller than Sue, then John is taller than Sue.

3 Why vagueness?

It is standardly assumed that the existence of vagueness in natural language is *unavoidable* in ordinary discourse. Our powers of discrimination are limited and come with a margin of error, and it is just not always possible to draw

sharp borderlines. Sometimes we cannot be more precise even if we would want to. And indeed, the pervasive vagueness of natural language has often been regarded as a deficiency, especially by scientists, logicians, and philosophers. However, it seems that the use of vague expressions is not such a bad thing, and might even be beneficial compared to their precise counterparts.

First, on the standard view communicative success is defined as a 1-1 correspondence between what the speaker intends and how the listener interprets it. Communicating with vague expressions is then predicted to be bad. But the above definition of communicative success is both unreasonably strict, and more than required. Communication can still be successful in case speaker's intention and hearer's interpretation have sufficient overlap. How much overlap is sufficient depends on what is at stake, and could be utility-based, as suggested by R. Parikh (1994).

Standard Gricean pragmatic explanations of the use of language assumes that communication is a cooperative affair. In such a situation it never does any harm to be as precise as possible (disregarding processing costs). Thus, being vague can never be advantageous. This is in accordance with a standard game theoretical result saying that messages with precise meanings can be communicated successfully only in case the preferences of speaker and listener coincide. However, this result is based on the assumption that communication is *noiseless*. Recently, some game theorists (e.g. Myerson, 1991; de Jaegher, 2003) have shown that once the preferences of speaker and listener are not completely aligned, we can sometimes *communicate more* with vague, imprecise, or noisy information than with precise information. Vague, or *indirect*, use of language might be beneficial as well in case the speaker is *unsure* about the preferences of the hearer. In such circumstances a speaker might intent some of his messages to be diversely interpretable by cooperative versus non-cooperative listeners.

It is sometimes argued that it is useful to have vague predicates like 'tall' in our language, because it allows us to use language in a *flexible* way. Obviously, 'tall' means something different with respect to men than with respect to basketball players, which means that it has a very flexible meaning. This does not show, however, that *vagueness* is useful: vagueness is not the same as context-dependence, and the argument is consistent with 'tall' having a precise meaning in each context. But not any function from contexts to precise meanings will do. The meaning of 'tall' should be *learnable* and *computable* by us as boundedly rational agents. One requirement seems to be that 'tall' should at least behave *consistent* across these contexts: It shouldn't be possible that in one context x is counted as tall, but y is not, while in another context where we still look at things from the same perspective, but where we take more or less individuals into account, it is the other way around. But consistency is not

enough: for a predicate to be precise, it has to be learnable and computable what the extension of that predicate is in each context. Bosch (1983) argues this is too much to ask, because there are many potential contexts for which we wouldn't know how to distinguish the individuals to which we can apply the predicate from those to which we cannot.⁵

In the introduction I claimed that vagueness involves not only tolerance, but also the required level of fine-grainedness. It seems beneficial for both the speaker and the hearer to sometimes describe the world at a more coarse-grained level (see for instance Hobbs (1985) and Krifka (2006)). There is obviously something to this. Deciding which precise term to use may be harder than using a vague term. For the listener, information which is too specific may require more effort to analyze. Another reason for not always trying to be as precise as possible is that this would give rise to *instability*. As stressed by Pinkal (1995), in case one measures the height of a person in all too much detail, this measure might change from day to day, which is not very useful. Also, in case there exists no standard way to measure the length of a certain object (like a river), being too precise is only confusing if the method of measurement is not provided as well.

A final reason why vague expressions are so prevalent in natural language might be that vague expressions are very useful to make *value judgments*.⁶ Even if you know that Quiza is 1.45 meters long, you might still learn something new when I tell you that she is *tall* for a Martian. This is mostly the case because whether one is tall depends on a gradient contingent fact, one's height, and it is not always very clear whether this particular height counts as being tall or not (for a Martian). A man for whom it is not clear whether he is tall or not, is standardly called a *borderline case* of tallness. John might be a borderline tall man although the speaker knows his height either because there is divergence in usage among competent speakers, or because the speaker still hesitates to call him tall or not. Lewis (1969), for instance, adopts the former characterization and suggests that languages themselves are free of vagueness but that the linguistic conventions of a population, or the linguistic habits of a person, select not a point but a fuzzy region in the space of precise languages.⁷ Proponents of the epistemic approach to vagueness (e.g. Williamson, 1994) and many adherents of the supervaluation theory (e.g. Fine, 1975; Kamp

⁵Bosch (1983) refers here to Waisman's (1968) notion of 'open texture', the feeling that for outlandish cases we wouldn't know yet how to apply the predicate. I believe this notion is closely related with the notion of 'unforeseen contingencies' in economics.

⁶Franks Veltman stresses this in his inaugural lecture Veltman (2002).

⁷Lewis (1970) takes this analysis of vagueness to be very similar to (what is now called) a supervaluation account. Burns (1991) argues (unconvincingly, we think) that the two are very different.

1975; Keefe, 2000) go for the second alternative: English, or its valuation function, is, or can be made, precise, but agents don't know, or don't bother to care, exactly what this valuation function is. On such a view, borderline cases of 'tall' are individuals that some speakers of a language consistent with the linguistic convention of English consider to be tall, while others don't (on Lewis's meta-linguistic account), or individuals of which an agent doesn't know, or doesn't care much about, whether they count as being tall or not, although the agent knows their precise height.

A semantic valuation function normally only determines how the facts are (e.g. whether John's height is 1.80 meters or 1.70 meters). One way to model the above intuition of borderline cases is to assume that a valuation function includes information about 'semantic' facts, and that a speaker takes several of such valuation functions to be possible. A valuation function includes information about semantic facts, if it not only says what John's height is, but also whether this particular height counts as being 'tall', i.e., whether somebody who is 1.80 meters should be considered to be tall or not. On this view, valuation functions, or worlds, should thus fulfill two roles, and those roles are exactly the roles a world can play according to Stalnaker's (1978) two-dimensional view on language. If we fix the meanings of the expressions, a sentence expresses a meaning, represented by a set of worlds, and if the actual world is a member of this set, what is said by the sentence is true, false otherwise. But if what is expressed by a (token of a) sentence depends on context, we can think of the world (together with the expression token) as determining how the expressions should be interpreted, and then it might be that in different worlds something different is said (i.e. different meanings are expressed) by the same sentential token. In this way we might explain why it might be unclear to the hearers what the speaker meant by saying 'I like you': the hearers might be unsure about the context that determines the reference of 'you' that the speaker had in mind. In the same way one might argue that also the meaning of nouns and adjectives depends on the utterance context. But if worlds fulfill the two roles suggested above, it seems natural to assume that they always fulfill the two roles *at the same time*. It follows that if we interpret a sentential token of a sentence φ in world w of which we consider it possible that it is the actual world, we use w both to determine *what is said* by φ , (denoted by $\llbracket\varphi\rrbracket_w$), and to determining whether what is said by φ in w is *true* in w , i.e., whether $w \in \llbracket\varphi\rrbracket_w$. The set of worlds in which what is said by sentence φ is *true* in w is also true in w – denoted by $\{w \in W : w \in \llbracket\varphi\rrbracket_w\}$ – is called the *diagonal proposition* by Stalnaker (1978). The diagonal proposition expressed by a sentence can be used to explain the so-called *evaluative mean-*

ing of vague predicates. As noted by Barker (2002),⁸ if one says that ‘John is tall’, one can make two kinds of statements: a *descriptive* one saying that John is above the relevant cut-off point (if it is clear in a context what the cut-off point for being tall is) and a *metalinguistic* one saying that the cut-off point for being tall is below the height of John (if it is clear in a context what John’s height is). The latter involves the evaluative meaning of ‘tall’ and can be accounted for straightforwardly in terms of diagonalization. There is nothing extraordinary about metalinguistic statements: identity statements can be used to express them as well. It is well-known that identity statements can be used (i) to state the identity of meaning of the two terms, but also (ii) to fix the meaning of one term in terms of the meaning of the other. The second reading is prominent in a sentence like ‘Deep throat is the person who was the source of Woodward and Bernstein’s Watergate information’, and should be understood metalinguistically.

4 Gradable adjectives

Although we have seen in section 2 that expressions of many lexical categories are vague, most research on vagueness concentrates on adjectives like ‘tall’ and ‘wide’. In linguistics these adjectives are known as *gradable* adjectives and should be distinguished from non-gradable adjectives like ‘pregnant’ and ‘even’. The latter adjectives do not give rise to (much) vagueness. There exist two major types of approaches to the analysis of gradable adjectives: degree-based approaches and delineation approaches. Degree-based approaches (e.g. Seuren, 1973; Cresswell, 1976; Bierwisch, 1984; von Stechow, 1984, Kennedy, 1999, 2007), analyze gradable adjectives as relations between individuals and degrees, where these degrees are thought of as scales associated with the dimension referred to by the adjective. Individuals can possess a property to a certain measurable degree. The truth conditions of sentences involving these adjectives are stated in terms of degrees. Delineation approaches (Lewis, 1970; Kamp, 1975; Klein 1980, 1991) analyze gradable adjectives like ‘tall’ as simple predicates, but assume that the extension of these terms are crucially context dependent. In this section we will discuss both types of approaches in somewhat more detail, and also see how they treat comparatives.

4.1 The degree based account

According to the degree-based approaches, relative adjectives are analyzed as relations between individuals and degrees, where these degrees are associated

⁸But see also Kyburg & Morreau (2000).

with the dimension referred to by the adjective. Individuals can possess a property to a certain measurable degree, and the truth conditions of comparative and positive sentences are stated in terms of degrees. According to the most straightforward degree-based approach, the absolute (1) is true iff the degree to which John is tall is (significantly) greater than a (contextually given) standard of height.⁹ while the comparative (2) is true iff *the* degree to which John is tall is greater than *the* degree to which Mary is tall.

- (1) John is tall.
- (2) John is taller than Mary.

But this straightforward degree-based approach has a problem with examples where the scope of the comparative contains an indefinite ('any'), or existential modal:

- (3) a. John is taller than anyone else.
- b. John is taller than allowed.

It is not easy to see how the above degree-based approach can account in a compositional way for the intuition that from (3-a), for instance, we infer that John is taller than *everybody* else, if we treat *any* as an indefinite. To account for this problem Seuren (1973) proposed that (2) 'John is taller than Mary' should be counted as true iff *there is* a degree d of tallness that John has but Mary does not: $\exists d[Tall(j, d) \wedge \neg Tall(m, d)]$.¹⁰ In this formalization, $T(j, d)$ means that John's degree of tallness includes *at least* d . This analysis easily accounts for the intuition concerning (3-a) and (3-b), by representing them by (4-a), and (4-b) respectively (treating 'any' as an existential quantifier):¹¹

- (4) a. $\exists d[T(j, d) \wedge \neg \exists x[x \neq j \wedge T(x, d)]]$.
- b. $\exists d[T(j, d) \wedge \neg \diamond T(j, d)]$.

Unfortunately, Seuren's analysis of comparatives meets a serious problem: we can't immediately account for the fact that from the truth of (2) we conclude

⁹For some adjectives, like 'tall' this standard is either defined as the (arithmetical or geometrical, possibly weighted) mean of the height of all individuals in the class, for others, like 'red' this standard can be thought of as the prototype representative of the class.

¹⁰The idea that the 'than'-clause of a comparative contains a negation goes back to Jespersen (1917).

¹¹Von Stechow (1984) proposed a somewhat different analysis to account for these sentences. It is fair to say that although Von Stechow's analysis has been improved on recently, it is the classic paper on comparatives, and still contains the most complete discussion of the analyses on the subject.

to (5) which involves its antonym:¹²

(2) John is *taller* than Mary.

(5) Mary is *shorter* than John.

There seems to be an easy way to account for this problem, however: just assume that ‘John is short’ is true if and only if he is not tall. Unfortunately, accounting for antonyms in this way gives rise to a serious problem. The problem – which von Stechow (1984) attributes to Hoepelman – is that it is predicted that the following two sentences are equivalent:

- (6) a. *John is tall and short.
b. John is neither tall nor short.

Indeed, the approach predicts that (6-a) is represented as $T(j) \wedge \neg T(j)$, while (6-b) is represented as $\neg T(j) \wedge \neg\neg T(j)$, which have the same meaning, and denote a contradiction. Although (6-a) seems inappropriate for exactly this reading, (6-b) seems to express a contingent proposition. The obvious conclusion from these examples is that we should not analyze ‘short’ as ‘not tall’. But if we can’t do this, it is not clear anymore how to account for the inference from (2) to (5) without making extra *ad hoc* assumptions.¹³

Von Stechow (1984) and Kennedy (1999) conclude that to account for this latter problem¹⁴ we have to assume that degrees are *directional*. Fortunately, it appears that in the alternative delineation approach to gradable adjectives, the above problem need not show up.

4.2 The delineation approach

Lewis (1970) and Kamp (1975) make use of supervaluation theory to account for the vagueness of adjectives like ‘tall’ and ‘bald’. These expressions are taken

¹²In an early, but still very relevant study on grading, Sapir (1944), makes a distinction between adjectives for which inferences like that between (2) and (5) go through, and adjectives for which they do not. In the first class are things like ‘good’-‘bad’ and ‘far’-‘near’, while in the second class are antonyms like ‘brilliant’-‘stupid’: From ‘John is it less/more brilliant than Mary’ we conclude that both John and Mary are brilliant, and we can’t continue this sentence with ‘but both are stupid’. This in contrast with the appropriate discourse ‘From the point of view of America, France is on the *near* side of Europe, though actually *far*’.

¹³Proponents of the degree-based approach can propose that in comparatives ‘short’ means ‘not tall’, but not in the positive use of the predicates: ‘John is short’ doesn’t mean the same as ‘John is not tall’. They can account for this by making use of the POS-operator they use to account for positive use of adjectives.

¹⁴And for a related problem that you can’t say ‘John is 1.80 meters *short*’.

to have specific cutoff-points with respect to each classical valuation function, or world, but it is undetermined, or unknown, what the actual cutoff-point for English is. Thus, a positive adjective is considered to be a simple predicate the extension of which depends on a world-dependent cutoff-point. For an adjective like ‘tall’, they propose that ‘John is taller than Mary’ is true just in case the set of worlds in which John is tall is a proper *superset* of the set of worlds in which Mary is tall.¹⁵ Because the proper superset relation gives rise to a strict partial order (irreflexive and transitive),¹⁶ it is correctly predicted that the comparative gives rise to a strict partial order as well. Still, the resulting analysis is not completely satisfactory.

First of all, as noted by Kamp (1975), in order to make this analysis work for comparisons like ‘taller than’ that give rise to orders that are not only irreflexive and transitive but satisfy other properties as well, Kamp and Lewis have to make a crucial *meaning postulate* to constrain the possible cutoff-points of a vague predicate P in the worlds: for all individuals x and y : either ‘ $P(x) \rightarrow P(y)$ ’ must be true in all worlds, or ‘ $P(y) \rightarrow P(x)$ ’.¹⁷ Although the meaning postulate required by Kamp and Lewis is natural, it is questionable whether by assuming it, we not already assume that we take the comparative to be more basic than the positive use of the adjective. But if so, this proposal is not in the ‘spirit’ of the delineation approach after all.

Second, the analysis gives rise to the same problem as Seuren’s (1973) analysis does: it cannot account for the equivalence of (6-a) and (6-b) without making ad hoc assumptions. As noted by von Stechow (1984), the reason is that the analyses of Kamp and Lewis are essentially the same as Seuren’s degree-based analysis. Lewis (1970) and Kamp (1975) assume that the worlds differ from each other in their cutoff-points of vague predicates. The comparative ‘John is taller than Mary’ is considered to be true if and only if there is a cutoff-point for ‘tall’ such that John is above it, while Mary is not. The easiest way to think of the cutoff-point of ‘tall’ in a world is as a particular *number*, a *degree*. But then we can assume that the predicate denotes a relation between individuals and degrees, and the delineation approach just claims that the comparative is true iff John has a degree of tallness that Mary does not have. This, of course, is exactly Seuren’s analysis of comparatives.

In the standard supervaluation theory used by Lewis and Kamp it is as-

¹⁵Well, this is Kamp’s (1975) proposal in case at least one of the two is considered to be a borderline tall individual.

¹⁶A relation R is irreflexive iff for all objects $x \in I : \neg R(x, x)$. It is transitive iff for all objects $x, y, z \in I : R(x, y) \wedge R(y, z) \rightarrow R(x, z)$.

¹⁷Notice that by adopting this constraint we can reformulate the above analysis of comparisons in terms of existential quantification: John is taller than Mary is true in a supervaluation frame just in case there is a complete valuation function, or world, in which John is tall, but Mary is not.

sumed that vague words like ‘bald’ simply have an extension in a world: each world has a unique cutoff-point from whereon individuals are not counted as bald anymore. But it might be, of course, that John can be counted as tall when we compare him with other men, but not tall when we compare him with (other) basketball players. Thus, whether someone of 1.80 meters is tall or not depends (partly) on who that person is compared with. In case the adjective occurs in attributive position it is natural to assume that the comparison class is at least partly determined by the *noun* to which the adjective is attached. For John to be *a tall man* it must be that John is being tall for a man, but that doesn’t mean that John is a tall basketball player. This suggests that whether we count an object to be *tall for an X*, we compare this object with the other objects that are *X* and decide whether the former object is tall. Because adjectives need not be attached to nouns, however, this can’t be enough. That is why we will assume with Wheeler (1972) and Klein (1980) that every relative adjective should be interpreted with respect to a *comparison class*.¹⁸ A comparison class is just a set of objects/individuals and is contextually given. In particular if the adjective stands alone, we might assume that the contextually given comparison class helps to determine what counts as being *tall*.¹⁹ The truth of the positive sentence (1)

(1) John is tall.

depends on the contextually given comparison class: (1) is true with respect to comparison class *c* (in model *M*) iff John is counted as tall in this class (in model *M*). The proposition expressed by a comparative like (2) is context independent.²⁰

(2) John is taller than Mary.

and the sentence is true iff there is a comparison class according to which John counts as tall, while Mary does not.²¹ This analysis is obviously close to the

¹⁸Graff (2000) has argued that a comparison class should be an *intensional* instead of an extensional object. If we allow for individuals of different worlds to be part of the same comparison class, I don’t think this crucially undermines Klein’s proposal.

¹⁹Although it does not necessarily uniquely determine what counts as being *tall*. See the end of this section for some discussion.

²⁰But see section 6.

²¹Equatives can be analyzed in terms of comparison classes as well. Klein (1980) proposes that (i-a) should be interpreted as (i-b).

- (i) a. John is *as tall as* Mary.
 b. In every context where Mary is tall, John is tall as well.

Klein (1980) notes that on this analysis, the negation of (i-a), i.e. (ii-a), is correctly predicted

Kamp/Lewis approach to comparatives and to Seuren’s degree-based account. Indeed, one can easily show that Klein (1980) can account for (3-a), and (3-b) in almost exactly the same way as Seuren (1973) could. But Klein’s (1980) analysis seems better suited to account for the fact that from (2) we conclude that Mary is shorter than John. We have concluded above that to account for this problem we are not allowed to analyze ‘short’ as ‘not tall’. Fortunately, Klein (1980) doesn’t need to analyze ‘short’ in this way in order to account for the equivalence of (2) and (5).

(5) Mary is *shorter* than John.

The reason is that if there is any comparison class in which John but not Mary counts as tall, this is also the case in the comparison class containing just John and Mary.²² But this means, intuitively, that *in this context* Mary is short, while John is not. From this we can conclude that we can account for the intuition that (2) and (5) have the same truth conditions without assuming that we should analyze ‘short’ as meaning *not tall*.

Notice that both the degree-based and the delineation account assume that the meaning of comparatives are context independent, but that the meaning of sentences containing positive adjectives are context dependent. On the comparison-class approach context-dependence is (partly) accounted for by the selection of the relevant comparison class, while on the degree approach the truth conditions of a sentence like (1) depend on the contextually selected standard of height. In this sense, the two approaches are the same. But there are also four differences that (we feel) are in favor of the comparison-class approach.

First, degree-based approaches need something like comparison classes as well. Such approaches claim that a positive sentence like ‘John is tall’ is true iff the degree to which John is tall is (significantly) greater than a given standard of height. But, obviously, this standard for basketball players differs from the standard for men.

to be equivalent with (ii-b):

- (ii) a. John is *not* as tall as Mary.
- b. Mary is taller than John.

Standard pragmatics can explain why in the context of question ‘How tall is John?’, (i-a) would come to mean that John and Mary are equally tall.

²²Later we will see that this is guaranteed by van Benthem’s (1983) Downward Difference constraint (DD). However, some authors have argued that objects indistinguishable with respect to a small comparison class, might be distinguishable with respect to a larger comparison class. According to proponents of this view (see section 5.2), constraint (DD) is not valid.

Second, the comparison-class account just assumes that the meaning of the comparative ‘taller than’ is a function of the meaning of ‘tall’. This is much in line with Frege’s principle of compositionality,²³ and also accounts for the fact that in a wide variety of languages the positive is formally unmarked in relation to the comparative (Klein, 1980). The degree-based approach, however, treats the comparative as basic, and states the truth conditions of a sentence involving a positive adjective like (1) by making use of the comparative relation ‘taller than’: John should be (significantly) taller than the norm, or normal size of individuals, or taller than indifferent, or taller than we expected him to be. But why then, does at least the vast majority of languages express the comparative in more terms than the positive?²⁴

The third obvious difference is that in contrast to the comparison-class approach, the degree approach makes essential use of *degrees* from the very beginning. Because we can only make sense of degrees once we adopt a system of full-blooded quantitative measurement, it is predicted that already for very simple comparatives like (2) it is crucial to make use of quantitative measurement. But this seems wrong: making comparisons precedes measurement and counting, both in science and in individual development. A child may observe and inform others, for instance, that an individual is *taller* than he was before, or that John is *taller* than Mary, before it is able to count or to handle a measure tape. Thus, it seems preferred to favor a system that is able to account for simple comparatives that doesn’t make use of precise measures, i.e. degrees. Of course, one also has to make sense of measurement phrases, but it would be preferred to find a system that accounts for comparisons that doesn’t explicitly mention measures without degrees. Moreover, the introduction of degrees should give rise to a system that is a *natural extension* of the theory of comparatives that doesn’t make use of degrees.

Suppose that one already has degrees with respect to a gradable adjective ‘*P*’ at one’s disposal. According to a degree-based analysis, a positive sentence of the form ‘*x* is *P*’ is now considered to be true iff *x* has property *P* to a degree

²³But see Von Stechow (1984) for an argument saying that also the degree-based approach is in line with Frege’s principle.

²⁴Sapir (1944, 125) argues that we must regard grading from different points of view. “It is very important to realize that psychologically all comparatives are primary in relation to their corresponding absolutes (‘positives’) [...] Linguistic usage tends to start from the graded concept, e.g. *good* (= better than indifferent), *bad* (= worse than indifferent) [...] for the obvious reason that in experience it is the strikingly high-graded or low-graded concept that has significance, while the generalized concept which includes all the members of a graded series is arrived at by a gradual process of striking the balance between these graded terms. The purely logical, the psychological, and the linguistic orders of primacy, therefore, do not necessarily correspond.” Perhaps this is so, but this doesn’t mean that making use of a theory that reflects such a correspondence wouldn’t be preferred.

higher than a contextually given cutoff-point. But one may wonder why anyone then would use the positive sentence at all? Why not immediately state the precise degree to which x has property P ? One valid reason can be, of course, that the speaker doesn't know the precise degree. Another reason might be that by using the positive sentence the speaker can make a value judgement (see section 3). However, it seems clear that there are more obvious reasons why the vague positive sentence is used. Consider expressions like 'probable' and 'useful'. They behave much like gradable adjectives. I take it that these types of expressions have been used (also to give value judgments) long before they have been given a measure-theoretic interpretation in the 20th century. It is unnatural to claim that a speaker who uses such an expression believes that its positive use comes with a precise cardinal cutoff-point, but that he or she just doesn't know what this exact cutoff-point is. In fact, most economists in the 19th and early 20th century used an expression like 'useful' even though they strongly *disbelieved* that 'usefulness' could be given a precise measure-theoretic interpretation.

A final problem for the degree-based account is that it is unclear how it accounts for the vagueness of sentences that involve, for instance, *nouns*, and not relative adjectives. Of course, prototype-based approaches of the meanings of nouns make use of something like degrees: the degree to which an object is a (typical) bird is inversely related to the *distance* of this object to the prototypical bird. But if one uses degrees here, it is unclear why we don't have simple comparatives like ' x is more heap than y ' that involve nouns.^{25,26}

Comparison classes can not only be used to help to determine the meaning of adjectives like 'tall', they can be used to help determining the meaning of degree modifier like 'very' and 'fairly' as well (cf. Wheeler, 1972; Klein, 1980). Suppose John is considered to be tall with respect to comparison class c . Then we might say that John is *very* tall just in case John is considered to be tall

²⁵Of course, we do have ' x is more *of* a heap than y ', but that seems to have a somewhat different meaning (or not?).

²⁶There exists a striking syntactic difference between common nouns like 'bank' on the one hand and verbs and adjectives like 'tall' on the other: whereas common nouns combine with a determiner to form a noun phrase, verbs and adjectives must be nominalized first, before they can play that role. It has been argued that there corresponds a semantic difference with this syntactic difference: it is in general not determinate how to count things like 'the tall ones' or 'the red ones' (a red grapefruit, for instance, won't have the same color as a red tomato), nor is it determinate how a thing which is tall or is red must be individuated and reidentified. In contrast to what falls under a common noun like 'cat' or 'bank', the general terms 'tall' and 'red' do not *by themselves* determine units which could underlie the possibility of counting: arbitrary many parts of a red object are red objects again. Of course, we can count red things, but only once we have determined beforehand what counts as an individual thing. Notice, though, that some philosophers have argued that even for nouns like 'cats' it is unclear exactly what should be counted (see Lewis (1993) for discussion).

among the tall ones in c , i.e., iff $John \in Tall(Tall(c))$. Similarly, we might say that John is *fairly* tall if he is tall, but not very tall. Assuming that predicates denote (in each model) functions from sets of individuals (i.e. comparison classes) to sets of individuals (i.e., comparison classes) is crucial for such an analysis of degree modifiers.

Klein (1980) argued that because the positive adjective ‘tall’ is formally unmarked in relation to the comparative ‘taller than’, the meaning of the latter should be a function of the meaning of the former. Before we will look at how the meaning of the comparative ‘taller than’ can be defined in terms of the meaning of the adjective ‘tall’, we will take a look at positive and comparative sentences involving *absolute* adjectives first, because it turns out that their behavior can be characterized easier.

4.3 Absolute terms and comparison classes

Consider adjectives like ‘full’, ‘flat’, or ‘straight’. Just like the meaning of ‘tall’, also the meaning of these adjectives is vague. These adjectives are also perfectly acceptable in comparatives: there is nothing wrong with saying that one surface is *flatter* than another, or that one bottle is *fuller* than another. In this respect they differ from adjectives like ‘pregnant’ and ‘even’, and are on a par with other gradable adjectives like ‘tall’. However, as observed by Unger (1975) and also discussed by Rothstein & Winter (2004) and Kennedy & McNally (2005), while with relative adjectives one can easily say something like

(7) John is tall, but not the tallest/ but somebody else is taller.

this can’t be done (so naturally) with (maximal) absolute adjectives.²⁷

- (8) a. *My glass is *FULL*, but it could be fuller.
 b. *This line is *STRAIGHT*, but you can make it straighter.

What this contrast shows is that sentences with absolute adjectives generate entailments that sentences with relative adjectives lack: it is inconsistent to say that something is flatter than something that is flat.²⁸ Thus, from (9-a) we conclude (9-b) that the pavement is not flat:

- (9) a. The desk is flatter than the pavement.

²⁷According to Unger, stress forces a precise interpretation of the absolute adjective.

²⁸Bolinger (1972) and others working on degree words observed another contrast between absolute and relative adjectives: while relative adjectives combine well with degree adverbs like ‘very’ and ‘rather’, absolute adjectives combine well with other adverbs like ‘completely’, ‘almost’, ‘hardly’, and ‘nearly’.

- b. The pavement is not flat.

Comparatives with relative adjectives, however, don't give rise to entailments at all:

- (10) a. John is taller than Mary. $\not\Rightarrow$ John/Mary is tall/not tall.

The natural way to account for absolute adjectives in a degree-based approach is to assume with Kennedy & McNally (2005) and others that we start with a 'fuller than' or 'flatter than'-relation between degrees, but assume that in contrast to the 'taller than'-relation formed with the relative adjective 'tall', the 'fuller than' and 'flatter than'-relations have *maximal elements*. The positive use of the adjective has then this maximal degree as its cut-off point.

But now suppose that we would like to assume with Klein (1980) that not only comparatives with relative adjectives like 'taller than' should be defined in terms of its corresponding positive adjective, but that this also holds for 'flatter than'. How should we proceed? Intuitively – and in contrast to the meaning of *relative* adjectives –, the meaning of the positive use of 'flat' is context-*independent*. We will suggest that this means that in the positive use of 'flat', the adjective should always be interpreted with respect to the *maximal* comparison class: the whole domain. If 'flat' is used in a comparative, however, we will assume that its meaning is relativized to a *smaller* comparison class. However, 'flat' selects with respect to each comparison class intuitively its 'maximal' element.

Saying that the meaning of an adjective depends on a comparison class, and defining the comparative relation by saying that ' x is P -er than y ' iff x has property P in comparison class $\{x, y\}$, but y has not is not enough to give a proper analysis of the comparative relation. The reason is that once we relativize the meaning of adjectives to comparison classes, this leaves room for the most diverse behavior. It might be, for instance, that although x is considered flat in set $\{x, y\}$ and y is not, x is not considered to be flat in set $\{x, y, z\}$, while y is. In order to assure that the meaning of the adjective in different comparison classes behaves properly, we have to constrain this behavior. Fortunately for us, the formal problem we face here is exactly the same as the problem discussed already by Arrow (1959) of how to derive a preference ordering relation from the assumption that the notion of *choice* is primitive. Assuming a choice function – a function that selects elements from each finite set of options –, Arrow showed how this can generate an ordering relation if we put some natural constraints on how this choice function should behave on different contexts. Let us define a *context structure*, M , to be a triple $\langle I, C, V \rangle$, where I is a non-empty set of individuals, the set of contexts, C , consists of all finite subsets of I , and the valuation V assigns to each context $c \in C$ and

each property P those individuals in c which are to count as ‘being P in c ’. Say that $P(c)$ denotes the set of individuals in c that have property P with respect to c : $P(c) = \{x \in c : x \in V(P, c)\}$ (In the context of social choice theory, P is thought of as a choice function, selecting the *best* elements of c .) We follow Arrow by proposing the following principle of choice (C), and the cross-contextual constraints (A1) and (A2) to limit the possible variation:

- (C) $\forall c \in C : P(c) \neq \emptyset$.
(A1) If $c \subseteq c'$, then $c \cap P(c') \subseteq P(c)$.
(A2) If $c \subseteq c'$ and $c \cap P(c') \neq \emptyset$, then $P(c) \subseteq P(c')$.

If we say that ‘ x is P -er than y ’, $x >_P y$, $\text{iff}_{def} x \in P(\{x, y\}) \wedge y \notin P(\{x, y\})$, Arrow (1959) shows that the comparative as defined above gives rise to a *weak order*. A structure $\langle I, R \rangle$, with R a binary relation on I , is a weak order just in case R is irreflexive (IR), transitive (TR), and almost connected (AC).

Definition 1.

A weak order is a structure $\langle I, R \rangle$, with R a binary relation on I that satisfies the following conditions:

- (IR) $\forall x : \neg R(x, x)$.
(TR) $\forall x, y, z : (R(x, y) \wedge R(y, z)) \rightarrow R(x, z)$.
(AC) $\forall x, y, z : R(x, y) \rightarrow (R(x, z) \vee R(z, y))$.

We can now define the relations ‘being as P as’, ‘ \sim_P ’, and ‘being at least as P as’, ‘ \geq_P ’, as follows: $x \sim_P y$ iff_{def} neither $x >_P y$ nor $y >_P x$, and $x \geq_P y$ iff_{def} $x >_P y$ or $x \sim_P y$. The relation ‘ \sim_P ’ is predicted to be an equivalence relation, while ‘ \geq_P ’ is predicted to be reflexive, transitive, and strongly connected (meaning that $\forall x, y : x \geq_P y \vee y \geq_P x$).²⁹

Although weak orders seem appropriate for the analysis of comparatives,³⁰ by just looking at the comparison class involving the desk and the pavement, we can’t account for the contrast between the acceptability of (11-a) versus the (at least according to Unger (1975)) unacceptability of (11-b):

- (11) a. The desk is not flat, but it is flatter than the pavement.
b. *The pavement is FLAT, but the desk is flatter.

The asymmetry suggests that the comparison class involved in the first conjunct should play a role in the discourse as well: the comparison class c involved

²⁹It is standard in the literature to also denote a structure $\langle I, \geq_P \rangle$ with ‘ \geq_P ’ reflexive, transitive, and strongly connected by a weak order.

³⁰The reader probably doubts whether weak orders are also appropriate for the analysis of vagueness. Indeed, I will argue in section 5 that for vagueness, we need semi-orders, rather than weak orders.

in the analysis of the conjunctive sentences (11-a) and (11-b), and thus their first clauses, should be a *superset* of the comparison involved in the analysis of the second conjunct, i.e. $c \supseteq \{d, p\}$. Thus, a comparative like ‘The desk is flatter than the pavement’ is true in context c iff (i) $\{d, p\} \subseteq c$ and (ii) $d \in Flat(\{d, p\})$, but $p \notin Flat(\{d, p\})$. With this extra constraint we predict by (A1) that from the truth of this sentence in c , we conclude that the pavement is not flat in c , because $\{d, p\} \cap Flat(c)$ is by (A1) required to be a subset of $\{d\}$. From this it follows that (11-b) denotes a contradiction. Sentence (11-a), however, does not denote a contradiction, because the desk can be flat compared to the pavement, without it being the case that the desk is flat when compared to the other members of $c \supseteq \{d, p\}$.

Kennedy & McNally (2005) and Kennedy (2007) observe that although all sentences involving absolute adjectives give rise to inferences, there exists an important distinction between the two opposing absolute adjectives ‘dry’ and ‘wet’: whereas ‘dry’ behaves the same as ‘flat’ with respect to their entailment behavior in comparatives, ‘wet’ behaves quite differently, and gives rise to a *positive* inference:

- (12) a. The floor is drier than the table. \Rightarrow The table is not dry.
 b. The floor is wetter than the table. \Rightarrow The floor is wet.

How can we account for this difference between ‘dry’ and ‘wet’? Assuming that absolute adjectives satisfy axiom (A1) will prove useful here. But another distinction between relative and absolute adjectives is crucial as well. Cruse (1986) observes that in contrast to the case of relative adjectives, with antonym pairs involving absolute adjectives, the negation of one form (normally) entails the positive statement of the other (i.e., they are contradictory to one another).³¹

- (13) a. The door is not open. \Rightarrow The door is closed.
 b. The table is not wet. \Rightarrow The table is dry.
 c. The baby is not awake. \Rightarrow The baby is asleep.

Based on this observation, we will assume that ‘wet’ is defined as ‘not being dry’. It follows by axiom (A1) that if there is a subset c' of c in which the floor is not considered to be dry, the floor won’t be considered dry in comparison class c either. So, we can conclude that the floor is wet in c .

But not only the inference pattern in comparatives with ‘wet’ is different from that of its antonym ‘dry’, the same is true for acceptable discourses. For adjectives like ‘flat’ and ‘dry’, we have seen above that (11-a) is appropriate, but (11-b) is not, and this was explained by our constraint on accessible

³¹This is not always the case, it doesn’t hold for the pair *full-empty*.

comparison classes, and our assumption that absolute predicates obey constraint (A1). But the same constraint also explains the contrastive behavior of appropriate discourses for the adjective ‘wet’.

- (14) a. *?The floor is not wet, but wetter than the table.
 b. The table is wet, but the floor is wetter.

The appropriateness of (14-b) is easily accounted for: the floor is not considered dry w.r.t. comparison class c . In the subset {table, floor} of c , however, objects not dry in c might be considered dry, and thus the floor might be dry with respect to {table, floor}, and thus not wet. This doesn’t have to be the case for the table, and so (14-b) is predicted to be appropriate. As for the inappropriateness of (14-a), note that if the floor is not wet with respect to comparison class c , it can’t be counted as wet in any subset c' of c either, which is enough to rule out (14-a).

Although this analysis of absolute adjectives is appealing, the analysis given above cannot explain yet another contrast between absolute and relative adjectives. As observed by Sedivy et al. (1999) and discussed by Kennedy (2007), there is nothing wrong with using a relative adjective as part of a definite description in a sentence like ‘Please give me the *long* nail’ to distinguish two nails neither of which is particularly long. The use of absolute adjectives in such definite descriptions is much less appropriate.³² In terms of our framework, what this suggests is that the *use* of the *positive* absolute adjective (in contrast to the use of the adjective in a comparative) demands that the comparison class with respect to which the adjective is interpreted is simply the whole domain, and thus that only those individuals can be called ‘flat’, that are the flattest of all the individuals in the whole (context independent) domain. We will see in section 6 how this analysis might still account for the vagueness of absolute adjectives.

4.4 Comparison classes and relative adjectives

Let us now consider relative adjectives like ‘tall’ again. Just like for ‘flat’, we want to derive a weak ordering relation for ‘tall’ in terms of how this adjective behaves across comparison classes. However, it is easily seen that although ‘tall’ seems to obey axiom (A2), axiom (A1) is much too strong for relative adjectives: (A1) demands that if both x and y are considered to be tall in the context of $\{x, y, z\}$, both should be considered to be tall in the context $\{x, y\}$ as

³²It should be noted, though, that Kennedy (2007) bases this general claim on the difference between the behavior of ‘long’ versus ‘full’. But the absolute adjective ‘full’ behaves crucially different from other claimed absolute adjectives, in that its antonym ‘empty’ is not contradictory with ‘full’. This might well be a crucial difference.

well. But that is exactly what we don't want: in the latter context, we want it to be possible that only x , or only y , is considered to be tall. We should conclude that if we want to characterize the behavior of relative adjectives, we should give up on (A1). Unfortunately, by just constraints (C) and (A2) we cannot guarantee that the comparative behaves as desired. In particular, we cannot guarantee that it behaves as almost connected ((C) assures that the relation 'at least as P as' is complete, while (A2) assures that the comparative behaves as a strict partial order, i.e., irreflexive and transitive.).

To assure that the comparative behaves as desired, we add to (C) and (A2) the Upward Difference-constraint (UD), proposed by Van Benthem (1982). To state this constraint, we define the notion of a difference pair: $\langle x, y \rangle \in D_P(c)$ iff_{def} $x \in P(c)$ and $y \in (c - P(c))$. Now we can define the constraint:

$$(UD) \ c \subseteq c' \text{ and } D_P(c') = \emptyset, \text{ then } D_P(c) = \emptyset.$$

In fact, van Benthem (1983) states the following constraints: No Reversal (NR), Upward Difference (UD), and Downward Difference (DD) (where c^2 abbreviates $c \times c$):

$$(NR) \ \neg \exists c, c' \in C, x, y \in I : \langle x, y \rangle \in D_P(c) \wedge \langle y, x \rangle \in D_P(c').$$

$$(UD) \ c \subseteq c' \text{ and } D_P(c') = \emptyset, \text{ then } D_P(c) = \emptyset.$$

$$(DD) \ c \subseteq c' \text{ and } D_P(c) = \emptyset, \text{ then } D_P(c') \cap c^2 = \emptyset.$$

Van Benthem (1982) shows that if constraints (NR), (UD) and (UD) are satisfied, the relations ' \geq_P ', ' \sim_P ' and ' $>_P$ ' as defined before still have the same properties as before: ' \geq_P ' is reflexive, transitive, and strongly connected; ' \sim_P ' is still predicted to be an equivalence relation, while the comparative ' $>_P$ ' is still predicted to be (i) irreflexive, (ii) transitive, and (iii) almost connected, just as in the case of absolute adjectives. For adjectives like 'tall', this seems just what we want.

It is important to realize that the above conditions constrain, but do *not* (at all) *uniquely determine* the behavior of the relevant predicate P across comparative classes. First, it might be that different context structures give rise to quite different ordering relations ' $>_P$ '. But even if two context structures gives rise to the same ' $>_P$ '-ordering, predicate P might still behave quite differently in those two models on larger comparison classes. This is due to the fact that to prove that ' $>_P$ ' behaves as desired, we look at comparison classes of at most 3 elements. Once we have larger subsets of I , there are at least two context structures that give rise to the same ordering that satisfy the above constraints. Thus, whether a particular individual of a given comparison class counts as a P -individual or not might still depend on the context structure

M .³³ This is as it should be, if we want to account for the observation of Kamp (1975) and Graff (2000) that even if the comparison class is identified, the meaning of the adjective might still depend on context. Graff (2000) observes, for instance, that with a sentence like ‘Fido is old for a dog’, where the comparison class is made explicit, one can not only attribute elderliness to Fido, but also extreme longevity.³⁴

‘Tall’ is known as a *one-dimensional* adjective. Kamp (1975) calls an adjective one-dimensional if we can associate with the adjective a unique measurable aspect that membership depends upon. Examples of one-dimensional adjectives are ‘tall’, ‘heavy’ and ‘hot’. For the adjective ‘tall’, for instance, this aspect is ‘height’, for ‘heavy’ it is ‘weight’, and for ‘hot’ it is ‘temperature’. These adjectives can easily be modified by adverbs of degree like ‘very’ and ‘quite’: ‘*very tall*’, ‘*quite heavy*’; easily undergo comparative formation (‘*taller*’, ‘*heavier*’, ‘*warmer*’), and superlative formation (‘*tallest*’, ‘*heaviest*’, ‘*hottest*’). One-dimensional adjectives have only one contrary predicate: its antonym. According to Kamp (1975), most adjectives are *more-dimensional*. The extensions of color adjectives like ‘blue’, for instance, are determined by three dimensions: its *brightness*, *hue*, and *saturation*. This is still relatively unproblematic. For other more-dimensional adjectives, things are less clear. There is no unique way to determine, for instance, whether John is *cleverer* than Mary: it doesn’t depend only on the ability of solving problems, and even if it did, it is not clear how the problem solving abilities in different contexts should be weighed against each other. Adjectives like ‘large’ and its antonym ‘small’ behaves just like ‘clever’: it is unclear whether it is its height, its volume, or its surface, or a combination of these which decides whether an object is large.

Important for us is that for more-dimensional adjectives P , the assumption that the comparative ‘ P -er than’ is almost-connected is problematic. The reason is that such adjectives induce orderings that give rise to *incomparability*

³³In terms of the analysis of value judgments sketched in section 3, one might think of each context structure M as a precise valuation function. Although all these structures give rise to the same ordering relation, they give the predicate that gives rise to this ordering relation different extensions with respect to the same comparison classes.

³⁴As argued by many opponents of analyses like ours, making the meaning of the adjective context dependent doesn’t eliminate its vagueness. The phrase ‘old for a dog’ is just as vague as the adjective ‘old’ is. Thus, making the meaning of P , the positive form, dependent on both a comparison class and a structure M is not enough. We will come back to this in section 5.

(cf. Klein, 1980),³⁵ and thus do not have a unique antonym.³⁶ Let us assume, for the sake of argument, that there are only two properties/dimensions associated with being clever: an ability to manipulate numbers, and an ability to manipulate people. Let us say that John is cleverer than Mary iff John is better both in manipulating numbers and in manipulating people. But now consider Sue. Sue is worse than John in manipulating numbers but better in manipulating people. Thus, neither John is cleverer than Sue, nor Sue is cleverer than John. For the *cleverer than*-relation still being almost connected, it has to be the case that Sue is cleverer than Mary. But it is well possible that although Sue is better in manipulating people than Mary (and John), Mary is better than Sue in manipulating numbers. Thus, if one doesn't fix a particular dimension, one cannot claim that 'cleverer than' denotes a relation that is almost connected. To solve this problem one could simply claim that the comparative can only be used if one fixes a particular dimension. According to two other proposals, we can also compare in a multi-dimensional way, but if we do so, almost connectedness is not valid anymore. We can make sense of that by either giving up constraint (UD) and try to replace it with another one that still guarantees that ' $>_P$ ' behaves transitive. Another way to give up almost connectedness would be to give up the assumption that C consists of *all* finite subsets of I , but to demand for all $c, c' \in C$ that if c and c' have a non-empty intersection, then their union is an element of C as well.³⁷

4.5 Degrees and measures

Remember that Arrow (1959) and Van Benthem (1982) showed that given their constraints on context structures, the comparative as defined by Klein (1980) and others is (i) irreflexive, (ii) transitive, and (iii) almost connected, where R is almost connected iff $\forall x, y, z : xRy \rightarrow (xRz \vee zRy)$. Relations with these properties are well-known in semantics: Lewis' (1973) relation of comparative similarity is one of them. It is also well-known that such relations R can be turned into *linear* orderings R^* of equivalence classes of individuals

³⁵The standard way to make the distinction between *indistinguishability* and *incomparability* is by starting with a structure like $\langle I, >_P, \sim_P \rangle$ where I is a set of objects, ' $>_P$ ' a primitive preference relation, and ' \sim_P ' a primitive indistinguishability relation. Given such a structure, and the natural definition of ' $<_P$ ' in terms of ' $>_P$ ', it is possible that $>_P \cup \sim_P \cup <_P \neq I \times I$. Thus, it is possible that for two elements x and y of I , it is neither the case that $x >_P y$, nor $y >_P x$, nor $x \sim_P y$. In that case, we call x and y *incomparable*. Now it is easy to rule out incomparability: just demand that $>_P \cup \sim_P \cup <_P = I \times I$.

³⁶For more-dimensional predicates, comparative formation is arguably more difficult as well. Kamp (1975), for instance, claims that 'This is bluer than that' is most of the time not a meaningful statement.

³⁷See van Rooij (2009) for making precise these two approaches.

that are connected ($\forall v, z : vR^*z \vee zR^*v$). First, say that x is R -equivalent to y , $x \sim_R y$, iff_{def} neither xRy nor yRx . Take $[x]_{\sim_R}$ to be the equivalence class $\{y \in I | y \sim_R x\}$. Then we say that $[x]_{\sim_R} R^* [y]_{\sim_R}$ iff_{def} xRy . Linear orderings like R^* form the basic input for any form of quantitative measurement, i.e., of *degrees!*

The idea of measurement theory (Krantz et al, 1971) is that we represent properties of and relations between elements of certain abstract ordering structures in terms of properties of and relations between real numbers that we already understand much better. A *quantitative* measure based on a linear ordering like $\langle I, >^* \rangle$ is a representation of the qualitative ordering relation (like *taller than*, represented by ' $>^*$ ') in terms of the quantitative ordering relation *greater than*, ' $>$ ', between real numbers. This measurement is defined in terms of a (homomorphic) function f that assigns each element of I to a real number such that $\forall x, y \in I : x >^* y$ if and only if $f(y) > f(x)$. In general, there are many alternative mappings to f that would numerically represent the qualitative relation $>^*$ equally well. However, this mapping f is unique up to a certain group of transformations. For instance, the mappings f and g to the (ordered) set of real numbers represent the same (*ordinal*) ordering structure $\langle I, >^* \rangle$ in case they can be related by a *monotone transformation*: for any $x, y \in I : f(x) > f(y)$ iff $g(x) > g(y)$. To faithfully represent *more informative* ordering structures, the different mappings should be unique up to a *more limited* group of transformations. For instance, suppose that the ordering structure can make sense not only of sentences like ' x is taller than y ', but also of things like ' x is taller than y by more than v is taller than w '. For a mapping f to represent this latter type of information, it should be the case that $f(x) - f(y) > f(v) - f(w)$. Two mappings f and g faithfully represent the same such an ordering structure, iff f and g are the same up to a *positive linear, or affine, transformation*, i.e., for any $x \in I : g(x) = \alpha(f(x)) + \beta$, where α, β are real numbers, and $\alpha > 0$. Indeed, it are linear transformations that preserve such differences.³⁸ The number α represents the fact that the *unit* of measurement is arbitrary. The number β represents the fact that the ordering structure doesn't have a fixed zero-point. Quantitative measures that are unique up to such linear transformations are called *intensional magnitudes*. The best known intensional magnitude is 'warmth'. We can measure it in terms of degrees Celsius and degrees Fahrenheit, and it is well-known that we can transform the one to the other by means of a linear transformation: x degrees Celsius is $32 + \frac{9}{5}x$ degrees Fahrenheit. Notice that one degree Celsius warmer is not the same as one degree Fahrenheit warmer (but $\frac{9}{5}$ degrees Fahrenheit

³⁸Such mappings not only preserve differences, but also ratios between differences: $\frac{f(x)-f(y)}{f(v)-f(w)} = \frac{g(x)-g(y)}{g(v)-g(w)}$.

warmer), and that $0\text{ }^{\circ}\text{C}$ is not at all the same as $0\text{ }^{\circ}\text{F}$, but rather as $32\text{ }^{\circ}\text{F}$.

To account for measure phrases like ‘John is 1.80 meter tall’, we need a theory of degrees that involves *addition*. Standard measurement theory gives us that for so-called *extensional* or *additive magnitudes*. Ordering structures that give rise to extensional measurement must come with an operation of *concatenation*, a specified procedure for joining two objects, denoted by ‘ \circ ’. Thus, the ordering structures must be of the form $\langle I, >^*, \circ \rangle$, where operation ‘ \circ ’ satisfies certain conditions, like associativity. For mapping f to faithfully represent such an ordering structure it must be the case that for all $\forall x, y \in I : f(x \circ y) = f(x) + f(y)$ (of course, it must also hold that $x >^* y$ iff $f(x) > f(y)$). Just as before, this mapping is unique up to a certain group of transformations. The group of transformations are now all of the form $g(x) = \alpha f(x)$, with $\alpha > 0$. The existence of α means that the unit of measurement is still arbitrary, but the disappearance of the constant β used in intensional magnitudes reflects the idea that the zero is *not arbitrary* anymore.

A striking fact about natural language is that certain adjectives combine well with measure phrases while others do not. For instance, it is ok to say that John is 1.80 meters *tall*, but it is not good to say that Mary is 5 degrees *happy*. A natural explanation of this is to claim that while ‘tall’ denotes an additive magnitude, ‘happy’ denotes only an ordinal one. A more interesting example is a *negative adjective* like ‘short’. The antonyms of positive adjectives generally don’t allow for measure phrases in their ‘bare’ use (to say ‘John is 1.80 meters short’ is inappropriate), but they combine well with measure phrases in comparatives: ‘Mary is 2 cm shorter than John’, and then they mean the same as ‘John is 2 cm taller than Mary’. Sassoon (2008) has recently given an appealing explanation of this fact by claiming that negative adjectives like ‘short’ are not extensional, but *intensional* magnitudes. This makes a lot of sense: if l is a mapping from individuals to their height, one can define a mapping s from individuals to their ‘shortness’ in terms of it by means of a linear (though not positive) transformation: $\forall x \in I : s(x) = \alpha l(x) + \beta$. The existence of the β reflects the idea that ‘shortness’ doesn’t have a zero-point. If one sets α to -1 – as one intuitively should –, this means that $s(x) = \beta - l(x)$, which explains many, if not all, of the striking linguistic data involving ‘short’ discussed by Kennedy (1999). But why does ‘shortness’ not have a zero-point? A natural answer would be that zero-point β is simply the height of the longest object or individual, and that we don’t know this height.³⁹

³⁹Schwarzchild (2006) noted that measure terms occur with some, but not all positive ‘measure’ adjectives. They occur in English with adjectives like ‘old’, ‘tall’, ‘high’, and ‘thick’, but not with ‘warm’, ‘heavy’, and ‘big’. We can say, for instance, ‘John is 5 years old’, and ‘The ice was 5 cm thick’, but we don’t say in English ‘The water was 75 degrees warm’, ‘The suitcase is 20 kilos heavy’, or ‘The apartment is 1000 square feet big’. We

5 Tolerance and the Sorites paradox

In section 4.2 we have argued that the meaning of P , the positive form, is dependent on a context dependent comparison class and on the context structure M . It is standardly argued, however, that this double context dependence is not enough. It still fails to explain the fact that the positive form of the predicate allows for tolerance and gives rise to the Sorites paradox. This problem will be addressed in this section.

5.1 Semi-orders

Consider a long series of people ordered in terms of their height. Of each of them you are asked whether they are tall or not. We assume that the variance between two subsequent persons is always *indistinguishable*. Now, if you decide that the first individual presented to you, the tallest, is tall, it seems only reasonable to judge the second individual to be tall as well, since you cannot distinguish their heights. But, then, by the same token, the third person must be tall as well, and so on indefinitely. In particular, this makes also the last person tall, which is a counterintuitive conclusion, given that it is in contradiction with our intuition that this last, and shortest individual, is short, and thus not tall.

This so-called Sorites reasoning is elementary, based only on our intuition that the first individual is tall, the last short, and the following inductive premise, which seems unobjectable:

[P] If you call one individual tall, and this individual is not visibly taller than another individual, you have to call the other one tall too.

Our above Sorites reasoning involved the predicate ‘tall’, but that was obviously not essential. Take any predicate P that gives rise to a complete

might try to account for this by saying that while adjectives like ‘old’ work on additive magnitudes, an adjective like ‘warm’ does not. Unfortunately, we can say things like ‘It is twice as warm in Amsterdam as it is in Berlin’. Moreover, there are cross-linguistic differences. In German some of the counterparts of the inappropriate examples above are acceptable: ‘Das Konzert war nur 40 Minuten lang’, ‘Der Koffer ist 20 Kilo schwer’, and ‘Die Wohnung ist 90 Quadratmeter gross’ are all appropriate. To account for this, Schwarzchild (2006) proposes a syntactic solution, but one might wonder whether a pragmatic solution is not more suitable. It seems that it is for some measure phrases much more natural to use the relative adjective than for others. It doesn’t make a lot of sense for mph, or degrees, or kilos, because saying that John went 50 mph doesn’t leave much room for the relevant (relative) adjective. The same for ‘It is 50 degrees’, and ‘The suitcase is 20 kg’. It is different with centimeters, because there are many adjectives measured in terms of centimeters, so adding tall, or wide, makes perfect sense.

ordering ‘being at least as P than’. Let us assume that ‘ \sim_P ’ is the indifference relation between individuals with respect to predicate P . Now we can state the inductive premise somewhat more formally as follows:

$$[\mathbf{P}] \text{ For any } x, y \in I : (P(x) \wedge x \sim_P y) \rightarrow P(y).$$

If we assume that it is possible that $\exists x_1, \dots, x_n : x_1 \sim_P x_2 \wedge \dots \wedge x_{n-1} \sim_P x_n$, but $P(x_1)$ and $\neg P(x_n)$, the paradox will arise. Recall that if $P(x_1)$ and $\neg P(x_n)$, it is required that $x_1 >_P x_n$. In section 4.2 we have defined the relation ‘ $>_P$ ’ in terms of the behavior of predicate P : $x >_P y$ iff $x \in P(\{x, y\})$ and $y \notin P(\{x, y\})$ and we will assume that $x \sim_P y$ holds iff_{def} neither $x >_P y$ nor $y >_P x$. We can assure that ‘ $>_P$ ’ and ‘ \sim_P ’ behave natural by putting constraints on the behavior of P across comparison classes. The constraints discussed in sections 4.3 and 4.4, however, did not allow for the possibility that $\exists x_1, \dots, x_n : x_1 \sim_P x_2 \wedge \dots \wedge x_{n-1} \sim_P x_n$, but $P(x_1)$ and $\neg P(x_n)$. Fortunately, there is a well-known ordering with this property. According to this ordering the statement ‘ $x >_P y$ ’ means that x is *significantly* or *noticeably* greater than y , and the relation ‘ $>_P$ ’ is irreflexive and transitive, but need not be almost connected. The indistinguishability relation ‘ \sim_P ’ is reflexive and symmetric, but need not be transitive. Thus, ‘ \sim_P ’ does not give rise to an equivalence relation. The structure that results is what Luce (1956) calls a *semi-order*. A structure $\langle I, R \rangle$, with R a binary relation on I , is a semi-order just in case R is irreflexive (IR), satisfies the interval-order (IO) condition, and is semitransitive (STr).⁴⁰

Definition 2.

A *semi-order* is a structure $\langle I, R \rangle$, with R a binary relation on I that satisfies the following conditions:

$$\begin{aligned} (IR) \quad & \forall x : \neg R(x, x). \\ (IO) \quad & \forall x, y, v, w : (R(x, y) \wedge R(v, w)) \rightarrow (R(x, w) \vee R(v, y)). \\ (STr) \quad & \forall x, y, z, v : (R(x, y) \wedge R(y, z)) \rightarrow (R(x, v) \vee R(v, z)). \end{aligned}$$

Just as weak orders, also semi-orders can be given a measure theoretical interpretation. If P is the predicate ‘tall’, one can think of R as the semi-order relation (significantly) ‘taller than’, and say that ‘ $x >_P y$ ’ is true iff the height of x is higher than the height of y plus some fixed (small) real number ϵ . In the same way ‘ $x \sim_P y$ ’ is true if the difference in height between x and y is less than ϵ . In case $\epsilon = 0$, the semi-order is a weak order.⁴¹

⁴⁰Any relation that is irreflexive and satisfies the interval-order condition is called an *interval order*. All interval orders are also transitive, meaning that they are stronger than strict partial orders.

⁴¹cf. Scott & Suppes (1958).

5.2 Contextual solutions to the Sorites paradox

Because it is such an age-old problem, there exist many proposed solutions to the Sorites paradox. It is fair to say that within linguistics, the so-called *contextualist*' solution is most popular. The contextualist' solution is really a whole family of solutions, but in all of its versions it is crucially based on the assumption that the meaning (or extension) of vague expressions depends on context. As such, it is in line with standard linguistic accounts of comparatives (whether they be degree-based, or comparison class-based). The contextualist solution to the Sorites paradox was first proposed by Kamp (1981). Other proponents of this type of solution include Bosch (1983), Pinkal (1984), Veltman (1987), van Deemter (1994), Raffman (1994, 1996), Gaifman (1997), Soames (1999), Graff (2000), and Shapiro (2006). In the following I will first discuss the central ideas of most of these contextualist' solutions, and the problems they give rise to, in a rather condensed way. Afterwards, I will discuss my favorite contextual solution in somewhat more detail.

The standard reaction to the Sorites paradox taken by proponents of fuzzy logic and/or supervaluation theory is to say that the argument is valid, but that the inductive premise [**P**] (or one of its instantiations) is *false*. But why, then, does it at least *seem* to us that the inductive premise is true? According to the standard accounts of vagueness making use of fuzzy logic and supervaluation theory, this is so because the inductive premise is *almost* true (in fuzzy logic), true in almost all complete valuations (in supervaluation theory), or that almost all its instantiations are true.

Linguists (e.g. Kamp, 1975; Klein (1980), Kamp & Partee, 1995; Pinkal, 1995) typically don't like the fuzzy logic approach to vagueness, because they can't account for what Fine (1975) called 'penumbral' connections. The standard objection of supervaluationists to the analysis of vagueness in terms of (standard) fuzzy logic is that 'John is *P* or he is not *P*' and 'John is *P* and he is not *P*' are not predicted to be tautological and contradictory, respectively. But this objection is less than convincing: if John is a borderline case of a tall man, what is wrong with saying, for instance, that he is to a certain extent tall and to a certain extent not tall? But other objections that also involve the assumed truth functionality are more serious: Suppose 'Bert is tall' is true to degree 0.5 and 'Fred is tall' is true to degree 0.4. This can only be the case, intuitively, if Bert is taller than Fred. Now consider (a) 'Fred is tall and Bert is tall' and (b) 'Fred is tall and Bert is not tall'. Adopting the standard fuzzy logic-treatment of negation, these sentences are predicted to have the same value (0.4, if also the standard 'min'-analysis of conjunction is adopted). But it seems that this is false: (a) can be true (to some extent), but if Fred is

shorter than Bert, sentence (b) can, intuitively, be true to no positive degree.⁴²

The treatment of vagueness and the Sorites paradox in supervaluation theory is not unproblematic either. It doesn't seem to have a good answer to the question what it means that the negation of $[\mathbf{P}]$, $\exists d, d'[d \sim_P d' \wedge P(d) \wedge \neg P(d')]$, is predicted to be supertrue. This statement seems to deny that there are borderline cases, and why aren't we able to say which one (or more) of its instances is not true.⁴³ This problem is related to Dummett's (1975) complaint: the use of complete refinements in supervaluation theory assumes that we *can* always make sharp cutoff-points: vagueness exists only because in daily life we are *too lazy* to make them. But this assumption seems to be wrong: vagueness exists, according to Dummett (1975), because we *cannot* make such sharp cutoff-points even if we wanted to. In terms of what we discussed above, this means that the relation ' $>_P$ ' should be thought of a semi-order, rather than as a weak order. But if this view is adopted, it is unnatural to claim that $[\mathbf{P}]$ is false.

Indeed, in contrast to the fuzzy logic and supervaluation theory treatments, the contextualist's solution of the Sorites paradox is based on the idea that $[\mathbf{P}]$ *must* be true, because this is a principle in accordance with the way we use language. We can think of $[\mathbf{P}]$ either as one universal sentence, or as a collection of individual conditional premises. Normally, this doesn't make any difference: the one universal sentence is classically equivalent with the collection of the individual instances. For many proponents of the contextualist's solution, however, this distinction is crucial. I will call $[\mathbf{P}]$ read as one universal sentence its *collective* reading. If $[\mathbf{P}]$ is really seen as a set of individual conditional premises, I will call it the *distributive* reading.

In his classic article, Dummett (1975) claimed that inductive premise $[\mathbf{P}]$ (in both of its readings) should be considered to be true, and not just indefinite or 'near to truth', because the corresponding inferences are essential for our reasoning with natural language concepts. Dummett's conclusion was that

⁴²As another example: consider the conditionals (c) 'If Fred is tall, Bert is tall' and (d) 'If Fred is tall, Bert is not tall'. If the values of these sentences depend only on the values of their parts, they should have the same value. But this prediction is wrong: while (c) is plausibly true, (d) is certainly wrong. In defense of fuzzy logic, one might claim that conjunction and negation should not be analyzed as standardly assumed by Lakoff (1973), for instance. In fact, there exist many alternative ways to analyze the connectives in fuzzy logic. Arguably, however, these alternatives are less natural than the standard one (cf. Dubois & Prade, 1980).

⁴³Williamson (1994) proposed that we should replace $[\mathbf{P}]$ by the weaker statement $[\mathbf{P}_w]$: For any $x, y \in I : (\Box P(x) \wedge x \sim_P y) \rightarrow P(y)$, where ' $\Box\varphi$ ' means that φ is known. But, as Graff (2000) points out, this proposal does not explain why we are so inclined to believe $[\mathbf{P}]$. Shapiro (2006) proposed to weaken $[\mathbf{P}]$ in yet another way. Making a difference between a classical ' \neg ' and a three-valued ' $\dot{\neg}$ ' negation, his principle of tolerance $[\mathbf{P}_{ws}]$ says that for any $x, y \in I : (P(x) \wedge x \sim_P y) \rightarrow \neg\dot{\neg}P(y)$. I believe we want something stronger than this.

thus natural language is inconsistent. Kamp (1981) follows Dummett’s claim with respect to the distributive reading of $[\mathbf{P}]$, but not his conclusion. To get around inconsistency, he (i) makes use of a sophisticated mechanism of *context change* and (ii) adopts a non-truth conditional analysis of conditional sentences, and proposes a weak, but non-standard notion of entailment. The idea of context change is that once it is explicitly accepted within the discourse that x has property P , for any vague predicate, the initial contextually given valuation function V changes into (possibly) new valuation function V' such that indistinguishable, or at least sufficiently similar individuals to x must be counted as having property P as well according to new valuation function (and context) V' . In other words, what Kamp proposes is that each of the inductive premises is true in case its antecedent is verified, because of context change.

Although they can be seen as improvements on Kamp’s initial proposal, the approaches by Pinkal (1984), Raffman (1994, 1996) and Soames (1999) are still very similar, in that they essentially assume that the interpretation of a vague predicate systematically changes as we proceed along a Sorites series. Stanley (2003) has argued that such contextual solutions to the Sorites paradox break down in its ‘elliptical’ version. In the sentence ‘John likes you, and Bill does too’, the occurrence of ‘you’ in the Verb Phrase ellipsis (henceforth, VP ellipsis) must be interpreted as referring to the same person picked out by the overt ‘you’ in the first clause. This fact not only holds for the pronoun ‘you’, but for all context dependent expressions. According to Stanley (2003) this means that if we assume that the content of a vague predicate changes as we proceed along a sorites series, the contextualist’s solution cannot handle the paradox stated in its ‘elliptical’ version:

- (15) If that₁ is tall, that₂ one is too, and if that₂ is, that₃ is too, . . . , that _{n} is too.

If the meaning of the word ‘tall’ is context dependent, it cannot shift its denotation in any of the different conditionals in (15), which means – according to Stanley (2003) – that the standard formulation of the contextualist’s solution can’t solve all versions of the Sorites paradox.

In his argumentation, Stanley (2004) crucially assumes that ‘context-dependence’ means ‘indexical’. But there exists a well-known argument due to Klein (1980) that the context dependence of the meaning of a word like ‘tall’ should *not* be indexical. Counterintuitive predictions result, Klein argued, if we would make that assumption. Look at the following example (16-a) due to Ludlow (1989):

- (16) a. That elephant is large and that flea is too.
 b. That elephant is large for an elephant and that flea is large for an elephant.

In case the adjective ‘large’ is treated as an indexical, the intuitively natural sentence (16-a) would be interpreted as something like (16-b), which is absurd. Klein (1980) and Ludlow (1989) conclude that the meaning of the adjective should not (in general) remain constant under VP ellipsis. Together with the assumption that VP ellipsis remains constant for indexical expressions, one has to conclude that adjectives like ‘tall’ and ‘large’ are not indexical.⁴⁴ But if the meaning of the adjective is not purely indexical, what is the property that remains constant under VP-ellipsis? One proposal would be that the comparison class with respect to which the adjective is interpreted can be dependent on the subject term involved. Thus, in (16-a) the property could be $\lambda x.L(x, f(x))$, where x is an individual variable, while f is a context dependent function from individuals to comparison classes.⁴⁵ Confronted with a Sorites series of objects x_1, \dots, x_n as in example (15), the constant meaning of the predicate is $\lambda x.T(x, f(x))$, and the context-dependent (partial) function f could be recursively defined as follows: $f(x_1) = \{x_1, x_n\}$, and $f(x_{i+1}) = f(x_i) \cup \{x_{i+1}\}$.

While most contextualists follow Kamp in validating the distributive reading of **[P]** by means of context change, they normally seek to improve on (ii) by making the resulting logic more classical. The latter is normally done by changing the notions of context and indistinguishability that are involved. One proposal along these lines is sketched by Veltman and Muskens (described in

⁴⁴This means that comparison classes are not scope-less (cf. Ludlow, 1989). Soames (1999) treats gradable adjectives as indexical, nevertheless. But also he might have good reasons for doing so, because even for indexicals constancy under VP ellipsis is disputable. Ellis (2004) gave the following example suggesting that this is not the case:

Thirty friends are standing in the middle of a very large field. One of them has the following idea: ‘Why don’t we each go and stand in any place we choose, and see where everyone goes. Jill, you go first.’ Jill walks a good distance away from the group and shouts, ‘I’m going to stand here!’ It’s Tom’s turn next, and being the tag-along Tom is, he goes straight for Jill and stands right next to her. Jill exclaims humorously, ‘And I guess Tom is too!’ Sally then goes and stands on the other side of Jill, who now says ‘And apparently, so is Sally!’ Then Bill goes and stands behind Jill (‘and so is Bill’), and then Ann stands in front of Jill (‘and Ann’). Each of the other twenty-nine people walks towards Jill and stands as close to her as s/he can without touching anyone else. In each case, Jill amusingly shouts ‘And so is s-and-so!’

In this case, the interpretation of ‘here’, which appears in VP ellipsis, varies.

⁴⁵This suggestion is close to one proposed by Kennedy (1999). He notes that something like this is needed to account for the intuition that for ‘Everybody in my family is tall’ to be true, we compare different members of my family to different sets: we compare men with other men, women with other women, and children with other children. Peter Bosch (p.c.) gave also the following example, which makes the same point: ‘Everything in America is big: The cars, the buildings, and even the Turkeys.’

Veltman, 1987). They make use of a construction first mentioned by Russell (1940), and elaborated by Luce (1956), Goodman (1966) and Dummett (1975), that makes the notion of ‘indistinguishability’ context dependent. According to this idea, a context can be modeled by a comparison class, and the elements of such a comparison class can be used by an observer as resources to distinguish between, for instance, the heights of individuals. If x is indistinguishable from y , which in turn is indistinguishable from z , it might still be that x is distinguishable from z (i.e., the notion of ‘indistinguishability’ is not transitive). What is proposed is that if z is available in context, it can be used to *indirectly* distinguish x from y : because in contrast to y , x is (directly) distinguishable from z . Thus, it might be that in one context, x and y are not distinguishable, while they are so in a larger context. (The notion of ‘indirectly distinguishability’ is transitive, and gives rise to an equivalence relation.) Now Veltman & Muskens propose that $[\mathbf{P}]$ is true under its distributive reading, but only with respect to their new notion of *indirect* indistinguishability.

To use this construction to solve the Sorites paradox, they propose that in each point in the discourse, a unique comparison class is relevant for the interpretation of a sentence, and thus predicate, at that point. For each sentence (formula) this unique class is simply the set of individuals mentioned by this sentence. Now consider any instance of $[\mathbf{P}]$ of the form ‘ $P(x) \rightarrow P(y)$ ’ with x and y not directly distinguishable. Because only x and y are mentioned by this sentence, the relevant comparison class is just $\{x, y\}$. But with respect to this class, x is not *indirectly* distinguishable from y either, and thus the premise is predicted to be true. This holds for all inductive premises, if treated separately.⁴⁶ But this doesn’t mean that we can derive a contradiction. In order to do that we would have to conjoin all inductive premises into one large inductive premise. But by the way the comparison classes relevant for each sentence are determined, the relevant comparison class of this large conjoined premise will contain many individuals. In fact, it will contain more than enough individuals such that at least for some set $\{x, y\}$ whose elements are indistinguishable in context $\{x, y\}$, its elements will be distinguishable in this larger context. Thus, what is wrong in the Sorites reasoning, on this account, are not the individual inductive premises, but the way these premises have been put together to produce a contradiction.

Van Deemter (1995) highlights that what Veltman & Muskens (1987) effectively propose is that inductive premise $[\mathbf{P}]$ (on both its collective and its distributive reading) is actually *ambiguous* between two readings. According to one reading, we should read the premise as $[\mathbf{P}_1]$:

⁴⁶In this respect, their analysis is very similar to the proposals of Raffman (1994, 1996) and Graff (2000).

[**P**₁] For any $x, y \in I, c \in C : (P(x, c) \wedge x \sim_P^{\{x, y\}} y \wedge y \in c) \rightarrow P(y, c)$.

It is clearly the case that if we formalize the inductive premise as [**P**₁] and read it collectively, the Sorites paradox will arise. Thus, [**P**₁] should be rejected on its collective reading. However, we can also read [**P**] in a different way such that it doesn't give rise to a contradiction. In fact, Van Deemter (1995) claims that [**P**] seems so natural because we don't read it as [**P**₁], but rather as [**P**₂]:

[**P**₂] For any $x, y \in I, c \in C : (P(x, c) \wedge x \sim_P^c y) \rightarrow P(y, c)$.

Notice that the only difference between [**P**₁] and [**P**₂] is that whereas [**P**₁] assumes that the similarity, or indistinguishability relation only considers x and y , while for [**P**₂] it might involve (many) more objects. We assume that $x \sim_P^c y$ is true for $x, y \in c$ if and only if the following statement is true: $x \in P(c)$ iff $y \in P(c)$. Notice that if c is a proper superset of $\{x, y\}$, i.e. $c \supset \{x, y\}$, it might be that $x \sim_P^{\{x, y\}} y$, but $x \not\sim_P^c y$. Comparison class c might have enough individuals available such that we can *indirectly* distinguish x from y . Now it is easy to see that this 'small' switch from $x \sim_P^{\{x, y\}} y$ to $x \sim_P^c y$ makes a crucial distinction: in contrast to [**P**₁], [**P**₂] is *analytically true*. Assume that the two statements in the antecedent of [**P**₂] are true: $x \in P(c)$ and $x \in P(c)$ iff $y \in P(c)$. From this we can conclude immediately that $y \in P(c)$, which is what the consequent says. Thus, [**P**₂] is true, because its antecedent is false for at least one pair of *directly* indistinguishable objects. Now, why does this new reading of the inductive premise not give rise to the Sorites paradox? The reason is obvious: it doesn't give any information on top of the first premise that $P(x_1)$ and $\neg P(x_n)$, so we cannot derive a contradiction. According to the above reasoning, the Sorites paradox arises due to an *equivocation* of comparison classes.

Appealing as this analysis might be, it is perhaps not exactly what we are looking for. The main worry is that the notion of 'indistinguishability' that Veltman & Muskens and Van Deemter propose is disputable. One might argue that Dummett's (1975) and Kamp's (1981) argument in favor of [**P**] is crucially based on the assumption that 'distinguishability' should really be *directly* distinguishability. It is doubtful whether reading [**P**₂] of [**P**] really implements Dummett's (1975) conviction that the latter is a regulative principle of our use of language.⁴⁷ We don't seem to hold [**P**] to be true if we are confronted with a Sorites series, just because its antecedent is false for at least one pair of objects. For it to be a *substantial* regulative principle of the way we use

⁴⁷In a very real sense, the solution Veltman & Muskens propose is based on the same intuition as the solution proposed by supervaluationists: the Sorites paradox arises, because we equivocate 'similarity' with 'sameness'. But Dummett's point was to take the former notion more seriously.

our language, it seems that we *should* read $[\mathbf{P}]$ as $[\mathbf{P}_1]$. But both Veltman & Muskens (1987) and Van Deemter (1995) predict that $[\mathbf{P}_1]$ on its collective reading is *false*.⁴⁸ But if we want to obey this principle, how, then, can we be saved from contradiction?

Dummett (1975) is mostly known for suggesting a rather pessimistic, or nihilistic, conclusion: real vagueness gives rise to contradiction. But in line with Wittgenstein’s radically pragmatic solution to the problem posed by vagueness in his *Philosophische Untersuchungen*⁴⁹ he also suggests that *in practice* contradiction is avoided. In normal discourse, we talk about relatively few objects, all of which are easily discernible from the others. In those circumstances, $[\mathbf{P}]$ will not give rise to inconsistency, but serves its purpose quite well. Only in exceptional situations i.e., when we are confronted with long sequences of pairwise indistinguishable objects — do things go wrong. But in such situations, we should not be using vague predicates like ‘tall’ but precisely measurable predicates involving, in this case, millimeters. We will discuss how a weak version of this reaction can be formalized naturally in terms of comparison classes. The idea is that it only makes sense to use a predicate P in a context — i.e. with respect to a comparison class —, if it helps to clearly demarcate the set of individuals that have property P from those who don’t. Following Gaifman (1997), we will implement this idea by assuming that any subset of I can only be an element of the set of *pragmatically appropriate* comparison classes C just in case the gap between the last individual(s) that have property P and the first that do(es) not must be between individuals x and y such that x is clearly P -er than y .⁵⁰ This is not the case if the graph of the relation ‘ \sim_P ’ is closed in $c \times c$.⁵¹ Indeed, it is exactly in those cases that the Sorites paradox arises.

How does such a proposal deal with the Sorites paradox? Well, it claims that *in all contexts in which P can be used appropriately*, $[\mathbf{P}_1]$ is true (in both its distributive *and* its collective reading). If we assume in addition that the first element x_1 of a Sorites series is the absolute most P -individual, and the last element x_n the absolute least P -individual, it also claims that *in all contexts c in which it is appropriate to use predicate P in combination with x_1 and x_n* , ‘ $P(x_1, c)$ ’ is *true* and ‘ $P(x_n, c)$ ’ is *false*. Thus, in all *appropriate*

⁴⁸For a related recent critique of contextualist’s solutions to the Sorites paradox, see Keefe (2007).

⁴⁹See in particular section 85-87: ‘A rule stands like a signpost ... The signpost in order in normal circumstances it fulfils its purpose. See also Waismann’s (1968) notion of ‘open texture’.

⁵⁰Also Graff (2000) has argued that not all sets of individuals can figure as appropriate comparison classes. Her reasoning, however, is quite different from ours. She argues, for instance, that comparison classes need to form a *kind*.

⁵¹Notice that also in discrete cases the relation ‘ \sim_P ’ can be closed in $c \times c$. In just depends on how ‘ \sim_P ’ is defined.

contexts, the premises of the Sorites argument are considered to be true. Still, no contradiction can be derived, because using predicate P when *explicitly* confronted with set of objects that form a Sorites series is *inappropriate*. Thus, in contrast to the original contextualist approaches of Kamp (1981), Pinkal (1984), and others, the Sorites paradox is not avoided by assuming that the meaning (or extension) of the predicate changes as the discourse proceeds. Rather, the Sorites paradox is avoided by claiming that the use of predicate P is inappropriate when confronted with a Sorites series of objects.⁵² As a result, Stanley’s problem for earlier contextualist’ solutions does not arise.^{53,54}

In our approach, we analyze adjectives like ‘tall’ with respect to a comparison class. But neither in this nor the previous sections have we discussed how a sentence like $P(x)$ should be interpreted with respect to model M and comparison class c in case ‘ x ’ does not denote an individual in c . One proposal would be to declare such a sentence neither true nor false. Another proposal would be to interpret the sentence not w.r.t. c but rather w.r.t. $c \cup \{x\}$. According to a yet different proposal, $P(x)$ should be true if there is a $y \in P_M(c)$ such that $M \models x >_P y$, and false otherwise. Now consider comparison class $c = \{x, y, z\}$ with $x \sim_P y$ and $y >_P z$. According to our account, we predict that x and y are considered to be P -individuals in c , while z is not. But now consider object $v \notin c$. Suppose that $x >_P v$, but $y \sim_P v$ and $v \sim_P z$. Should we count

⁵²Although Graff (2000) seems to adopt a version of the standard contextualist’ approach to vagueness, her analysis can be thought of as being close to our pragmatic analysis as well. She makes two claims: (i) there exists cutoff-points, but (ii) if x is significantly P -er than z , $x >^1_P z$, and y is similar to x , it must be the case that y is significantly P -er than z as well, $y >^1_P z$. Although the similarity relation Graff (2000) seems to assume behaves like the similarity relation used in semi-orders, the relation ‘ $>^1_P$ ’ can obviously not be the corresponding relation of a semi-order. Instead, the relation ‘ $>^1_P$ ’ should have the properties of a *weak order*. In fact, if we start with a semi-order $\langle I, > \rangle$, we can define a weak order $\langle I^*, >^* \rangle$ that behaves just like the one Graff (2000) seems to assume. Given that ‘ \sim ’ is defined in terms of ‘ $>$ ’ as usual, we can define a new relation ‘ \approx ’ in terms of it: $x \approx y$ iff_{def} $\exists z \in I^n : x \sim z_1 \sim \dots \sim z_n \sim y$. Obviously, ‘ \approx ’ is an equivalence relation. In terms of ‘ \approx ’ we define equivalence classes like $[x]_{\approx}$ as usual, and take I^* to be the set $\{[x]_{\approx} : x \in I\}$. Now we define the order relation $>^*$ between the elements of I^* as follows: $X >^* Y$ iff_{def} $X \neq \emptyset \wedge Y \neq \emptyset \wedge \forall x \in X : \forall y \in Y : x > y$. One can show that ‘ $>^*$ ’ indeed is a weak order. Now, why is (our reformulation of) Graff’s analysis close to our pragmatic analysis? The reason is that in order to assume that the first element of a Sorites series has property P but the last one has not, she has to assume (if our reformulation of her ideas is faithful) that I is not closed under ‘ \sim_P ’, and thus that the series allows for a cutoff-point.

⁵³On this proposal, Stanley’s elliptical conjunction is claimed to be inappropriate. Or better, if we assume that $[P]$ holds and that the property that remains constant under VP-ellipsis is $\lambda x.T(x, f(x))$, where f was defined as $f(x_1) = \{x_1, x_n\}$ and $f(x_{i+1}) = f(x_i) \cup \{x_{i+1}\}$, it is predicted that the sentence at the one but last step in the Sorites series is inappropriate.

⁵⁴It turns out to be possible to characterize semi-orders in terms of the way relative adjectives behave with respect to *appropriate* comparison classes (see Van Rooij, 2009).

‘ $P(v)$ ’ true with respect to M and c , or not? According to the second proposal mentioned above, we should not. But in that case, $M, c \models P(y) \wedge \neg P(v)$, although $y \sim_P v$. Notice that this is consistent with our approach, if we assume that *pragmatic* appropriateness conditions are *independent* of *semantic* truth conditions:⁵⁵ although ‘ $P(y) \wedge \neg P(v)$ ’ is true, it is not appropriate to assert it.

We have seen above that each context structure $M = \langle I, C, V \rangle$ gives rise to a relation of ‘indistinguishability’. Now I would like to argue that this latter notion should be considered more generally: $x \sim_P y$ iff there is no *relevant* difference between x and y concerning property P .⁵⁶ But, obviously, whether the P -distinction between x and y is relevant or not depends on context: on the *goals* of the participants of the discourse. Graff (2000) even suggests that the existence of vagueness in natural language is partly a consequence of the vagueness of our purposes. To argue in favor of this view, she points out that it depends on Smith’s purpose whether we are inclined to accept the following *sorites sentence*: ‘For any n , if n grains of coffee are enough for Smith’s purpose, then so are $n - 1$.’ Normally we are, but if very *fine distinctions* may tip the balance, things differ.⁵⁷ This suggests that in case fine distinctions matter, fewer individuals should be considered to stand in the tolerance relation ‘ \sim_P ’, while the relation ‘ $>_P$ ’ should contain more pairs of individuals. We will see in section 6 that what Graff (2000) has in mind is closely related to many other phenomena.

5.3 Boundaryless concepts and higher order vagueness

Vagueness is traditionally defined by the possession of borderline cases. More recently, it has been argued that this is not sufficient: a predicate might have borderline cases without being vague. What characterizes predicates as being vague, according to this alternative view, is the fact that these predicates have no clear borderline cases. A predicate is vague, if it lacks sharp boundaries. One way to model such *boundaryless* concepts is by making use of semi-orders. Just assume that the relation ‘ \sim_P ’ is closed under $I \times I$. In such a case, the topmost P individuals would semantically (definitely) fall under the extension of P , while the bottommost P -individuals would semantically (definitely) not

⁵⁵This is generally assumed in two-dimensional theories of presuppositions and conversational implicatures.

⁵⁶Luce’s (1956) intention was not to just model the notion of ‘indistinguishability’, but rather that of a more general notion of ‘indifference’. Also Pinkal (1995) and others have argued that tolerance should be based on ‘indifference’ rather than ‘indistinguishability’.

⁵⁷See Pinkal (1984, 1995) for a similar argument.

fall under the extension of P .⁵⁸ But whenever we want to look for the dividing line between the P - and the $\neg P$ individuals, we can't find it: For any two consecutive individuals x and y , if we focus on only those two individuals, we won't find a significantly enough difference between them to draw here the dividing line. But this is not only true for the dividing line between P and $\neg P$ individuals, but also for dividing lines between, for instance, the *definitely* P individuals and the not definitely P individuals, etc.

The view that vague concepts are boundaryless is closely related with the view that vague expressions give rise to *higher-order vagueness*.⁵⁹ According to this view, there is not only no clear (1st-order) borderline between the (clearly) P versus the (clearly) non P -individuals, there is also no clear (2nd-order) borderline between the clearly P versus the 1st-order borderline cases of P individuals, etc. This phenomenon is known as *higher-order vagueness*. Williamson (1994) and others have argued that one can account for higher-order vagueness within a modal logic, if one takes the accessibility relation in terms of which one defines the meaning of 'definitely' or 'clearly' to be *reflexive* and *symmetric*, but *not transitive*. Making use of semi-orders, it is quite easy to do this. Let us assume that W is the set of worlds, and that ' $>_P$ ' is an ordering between the elements of W such that $\langle W, >_P \rangle$ is a semi-order. One can think of $w >_P v$ to be true iff the cutoff-point for being a P -individual in w is (significantly) higher than that cutoff-point in v . Now define an accessibility relation R between worlds as follows: v is accessible from w , i.e. wRv , iff neither $w >_P v$ nor $v >_P w$. From section 5.1 we know that relation R must be reflexive, and symmetric, but need not be transitive. Intuitively, we can say that wRv iff the cutoff-points according to which one has property P in w and v differ in at most a fixed margin ϵ . Let us now add a *definitely* operator ' Δ ' to the language. We say that ' $\Delta\varphi$ ' is true iff ' φ ' is true on all R -accessible worlds. Now we say that x is a 1st-order borderline case of P -ness iff ' $\neg \Delta P(x) \wedge \neg \Delta \neg P(x)$ ' is true. Similarly, we say that x is a 2nd-order borderline case of P -ness if ' $\neg \Delta \Delta P(x) \wedge \neg \Delta \neg \Delta P(x)$ ' is true, etcetera.

Fine (1975) and Keefe (2000) propose a rather different account of higher-order vagueness. They suggest that to account for higher-order vagueness we have to realize that what counts as an admissible (total) valuation is itself vague. I just want to suggest a very simple – and no doubt simplistic – way

⁵⁸I have argued above that in such a case it doesn't make sense to *use* predicate P in context I . This doesn't rule out, of course, that P can be used appropriately with respect to comparison classes smaller than I .

⁵⁹Whether the assuming that predicate P give rise to (unlimited) higher order vagueness fully captures the intuition that the concept denoted by P is boundaryless is controversial, however.

to implement their basic idea.

Fine (1975) and Keefe (2000) make use of both partial and total valuation functions. Recall that both partial and total valuation functions are interpretation functions of a *language* with respect to a *domain* of quantification. Suppose that a language \mathcal{L} contains individual constants of all and only all individuals of domain I . In contrast to the partial valuation function, the total valuation functions do not allow for gaps – i.e., individuals in I that neither have property P nor property $\neg P$. Each individual that can be referred to by an individual constant, or quantified over, has property P or has not. Thus, according to each total valuation, there is a clear borderline between the individuals that have and the individuals that do not have property P . But now suppose that we (slightly) adjust the model by *extending its domain*, and that we adjust the language accordingly by introducing some new individual constants that can refer to those individuals. Now it will be the case that a valuation function that was total with respect to \mathcal{L} and I is only partial with respect to the new language \mathcal{L}' and new domain I' . Of course, these new partial functions can in turn be extended by new total functions, or valuations. But now the same game can be played again, and perhaps so indefinitely. Perhaps the most natural way to extend a domain, is to think of (some of) its elements as (something like) measures, or *degrees*. According to one model, we measure the height of individuals up to a decimeter precise, while at another up to a centimeter, or a millimeter precise. You will object that in this way we are not really talking about vagueness anymore, but about levels of *granularity*. But, then, perhaps vagueness is crucially related with granularity.

6 Vagueness and granularity

In a series of papers (e.g. Hobbs, 1985) Hobbs argues that both in reasoning and in natural language use it is crucial that we conceptualize and describe the world at different levels of granularity. A road, for instance, can be viewed as a line, a surface, or a volume. The level of granularity that we make use of depends on what is relevant. When we are planning a trip, we view the road as a line. When we are driving on it, we view it as a surface, and when we hit a pothole, it becomes a volume to us. In our use of natural language we even employ this fact by being able to describe the same phenomenon at different levels of granularity within the same discourse. Thus, we sometimes explicitly shift perspective, i.e., shift the level of granularity to describe the same situation. This is perhaps most obviously the case when we talk about time and space. Consider examples (17) and (18):

(17) It is two o'clock. In fact, it is two minutes after two.

- (18) The point of this pencil is actually an irregular surface with several peaks. (Asher and Vieu, 1995).

In (17) we shift to describing a time-point in a more specific way, while in (18) we shift from a description of a pencil point as a point to its being a surface.

6.1 Absolute terms revisited

At the end of section 4.3 we have suggested that with the positive use of absolute adjectives like ‘full’, ‘flat’, and ‘straight’, the only relevant comparison class can be the set of *all* objects, i.e. *I*. In this way, we can also account for the fact that in contrast to *relative* adjectives like ‘tall’, absolute adjectives don’t give rise to the Sorites paradox with distinguishable objects, because the second premise is quite naturally judged to be false (cf. Kennedy, 2007).

- (19)
- a. P1. A theater in which every seat is occupied is full.
 - b. P2. Any theater with one fewer occupied seat than a full theater is full.
 - c. C. Therefore, any theater in which half of (none of, etc.) the seats are occupied is full.

It also accounts for Kennedy’s (2007) observation that just like measure phrases, but in contrast to relative adjectives, absolute adjectives allow for natural precisifications.

- (20)
- a. We need a 10 meter long rod for the antenna, but this one is 1 millimeter short of 10 meters, so unfortunately it won’t work.
 - b. The rod for the antenna needs to be *straight*, but this one has a 1 millimeter bend in the middle, so unfortunately it won’t work.
 - c. ??We need a long rod for the antenna, but since *long* means ‘greater than 10 meters’ and this one is 1 millimeter short of 10 meters, unfortunately it won’t work.

An unfortunate consequence of this analysis, or so it seems, is that we have to conclude that absolute adjectives can hardly ever be used. This, however, seems to contradict actual practice. Moreover, the analysis cannot explain why absolute adjectives like ‘flat’ and ‘full’ allow for valuable interpretation and give rise to the Sorites paradox: just how bumpless should a table be to be called ‘flat’?, and how much liquid should a bottle contain before it can be called ‘full’? Thus, the puzzle then is to explain our daily use of absolute terms, and why they give rise to vagueness. We will take over Lewis’s (1979) suggestion here:

Peter Unger has argued that hardly anything is flat. Take something you claim is flat; he will find something else and you get to agree that it is even flatter. You think the pavement is flat – but how can you deny that your desk is flatter? But “flat” is an *absolute term*: it is inconsistent to say that something is flatter than something that is flat. Having agreed that your desk is flatter than the pavement, you must concede that the pavement is not flat after all. Perhaps you now claim that your desk is flat; but doubtless Unger think of something that you will agree is even flatter than your desk. And so it goes.

[...] The right response to Unger, I suggest, is that he is changing the score on you. When he says that the desk is flatter than the pavement, what he says is acceptable only under raised standards of precision.

Lewis suggests that although Holland has hills, a sentence like ‘Holland is flat’ can still be used truly and appropriately in daily conversations, because whether the sentence can be used appropriately or not might depend on a contextually determined *standard of precision*.

6.2 Standards of precision

Until now we have assumed that the comparison-class account always works with a *fixed* model, or context structure $M = \langle I, C, V \rangle$. We have seen that on the basis of such a context structure we can define an ordering relation, like ‘*P-er than*’ or ‘ $>_P$ ’, between individuals, and in terms of the relation ‘*being as P as*’ or ‘ \sim_P ’. In a different context structure M' , however, the relations ‘ $>_P$ ’ and ‘ \sim_P ’ might come out very differently. In particular, it might be that according to a different context structure, the resulting comparative relation ‘ $>_P$ ’ is more fine-grained.⁶⁰

⁶⁰In section 4 we have assumed that although statements involving predicates like ‘tall’ occurring positively are vague, comparatives are *not*. The vagueness of ‘tall’ was accounted for by interpreting sentences where the predicate is used positively with respect to a contextually given comparison class. Comparatives were not vague, because for their interpretation we existentially quantified over comparison classes. It is quite common among linguists to assume that comparatives are not vague. Philosophers like Kamp (1975), Williamson (1994) and Keefe (2000), however, have claimed otherwise. Kamp (1975) noted already that comparatives associated with more-dimensional predicates – for example ‘cleverer than’ – are typically vague. They have borderline cases: pairs of people about whom there is no fact of the matter about who is cleverer, or whether they are equally clever. This is particularly common when comparing people who are clever in different ways. Keefe (2000), argues that there can also be borderline cases of one-dimensional comparatives. She argues that

Let us now look at a *set* of context structures \mathcal{M} . Let us say that in all context structures M, M' of \mathcal{M} , the set of individuals, I , and the set of contexts, C , is the same: $I_M = I_{M'}$ and $C_M = C_{M'}$. This means that M and M' of \mathcal{M} only differ with respect to their valuation functions, $V_M \neq V_{M'}$. Now we can define a *refinement* relation between models M and M' as follows: we say that model M' is a refinement of model M with respect to predicate P only if $\exists x, y \in I, M \models x \sim_P y$, but $M' \not\models x \sim_P y$. So, M' is more fine-grained than M with respect to predicate P iff there is at least one pair of individuals equally P in M that is not equally P in M' . There is a natural constraint on the ordering between models: if Mary is taller than Sue, but shorter than John in fine-grained model M' , it cannot be the case that John and Sue are counted as equally tall in the more coarse-grained model M , but still taller than Mary. Formally: M' is a refinement of M w.r.t. P only if $\forall x, y, z \in I$: if $M' \models x \geq_P y \wedge y \geq_P z$ and $M \models x \sim_P z$, then $M \models x \sim_P y \wedge y \sim_P z$. This follows if we define refinements w.r.t. predicate P as follows: M' is a refinement of M with respect to P , $M \leq_P M'$, iff $V_M(>_P) \subseteq V_{M'}(>_P)$. Notice that if $V_M(>_P) \subseteq V_{M'}(>_P)$, it follows that $V_M(\sim_P) \supseteq V_{M'}(\sim_P)$, which is what we desired. Moreover, it follows that in more coarse-grained models, the relation ‘being at least as P ’ will be larger: $M \leq_P M'$ only if $V_M(\geq_P) \supseteq V_{M'}(\geq_P)$.

Consider now the following two sentences.

- (21) a. John is 2 meters tall.
 b. John is *exactly* 2 meters tall.

Both sentences are true if and only if the (maximal) height of John is the height denoted by the phrase ‘2 meters tall’. Intuitively, however, (21-a) would count in more contexts as being true, or appropriate, than (21-b), if John is a little bit more than 2.00 meters tall, e.g. if he is 2.02 meters tall. The question is how we should account for this latter difference. Part of the answer to this question involves the issue whether even (21-a) can be *true* in case John is 2.02 meters tall, or whether in such a case, (21-a) is (strictly speaking) false, but still appropriate, because *close enough to the truth*.⁶¹

although there is a single dimension of height, people cannot always be exactly placed on it and assigned an exact height. For what exactly should count as the top of one’s head? Consequently there may also be borderline cases of *taller than*. Even more interesting from our point of view is the fact that *taller than* can also be vague due to indeterminacy over exactly what should count as a point (or an equivalence class) in the tallness ordering: individuals whose height is 2 micro-millimeter apart normally count as equally tall, although this would not be the case in contexts in which every micro-millimeter is important.

⁶¹In a recent paper, Sauerland & Penka (2007) studied the difference between modifiers like ‘exactly’ and ‘definitely’. Both could be used to modify measure phrases, but they do so in different ways: ‘exactly’ says something about how precise the measure phrase should be interpreted, while ‘definitely’ measures the speaker’s epistemic certainty. This is

Lasersohn (1999) argues that in case John is 2.02 meters tall, (21-a) is false. In fact, (21-a) and (21-b) are claimed to be semantically equivalent. Still, sentence (21-a) can be used felicitously, because minor kinds of falsehood are pragmatically permissible. The distinction between 2.00 meters and 2.02 meters may not be relevant to the purposes of our discourse, so we can ignore it, counting the sentence as close enough to the truth for its context, even though not really true. But if (21-a) is true iff John is 2.00 meters tall, how should we then account for the above mentioned contrast between (21-a) and (21-b)? Lasersohn proposes that the function of ‘exactly’ is to put higher demands on the situations under which a sentence is pragmatically felicitous, i.e., under which circumstances (21-a) – the sentence without the word ‘exactly’ – counts as close enough to the truth.

To formally account for this idea, Lasersohn associates with each expression not only its denotation, but also a context dependent set of items understood to differ from the denotation only in ways which are pragmatically ignorable in that context. This set is called the *pragmatic halo* of the item. The pragmatic halo of a complex expression is determined compositionally from the pragmatic halos of its parts⁶² Assuming that we should analyze pragmatic vagueness in terms of pragmatic halos, (if α is an element of the pragmatic halo of expression \mathcal{A} of type $\langle b, a \rangle$, and β an element of the pragmatic halo of expression \mathcal{B} of type b , then $\alpha(\beta)$ is an element of the pragmatic halo of expression \mathcal{AB} of type a). Obviously, the actual denotation, or meaning, of an expression is also an element of the expression’s pragmatic halo, and for some types of expressions (e.g. logical words, presumably) this is the only element of the set. The result of this is that a sentence might be ‘true’ in its halo, although the sentence is actually (in its denotation) false. Now Lasersohn says that φ is close enough to the truth in a context iff one of the elements of φ ’s pragmatic halo in this context is true. To account for the distinction between (21-a) and (21-b), Lasersohn (1999) claims that the function of ‘exactly’ is to *contract* the halo. The result of this contraction is that if the proposition expressed by ‘John is 2.02 meters tall’ is an element of the pragmatic halo of (21-a), this proposition need no longer be an element of the pragmatic halo of (21-b). As a result, although (21-a) will be false, but true enough in the context, (21-b) is neither true, nor true enough in that context.

Lasersohn (1999) assumes that (21-a) is true iff John is exactly 2.00 meters tall. He assumes that the phrase ‘2 meters’ has a context-independent semantic interpretation, such that semantically this is never compatible with

corroborated by the fact that the adjective ‘tall’ can be modified by ‘definitely’ but not by ‘exactly’.

⁶²It is questionable whether compositionality should be assumed here, and it contrasts standard Gricean treatments of other pragmatic phenomena.

2.02 meters. But what is this context-independent semantic interpretation of ‘2 meters’? Presumably, or at least the only one that seems consistent with his general methodology, Lasersohn’s answer would be that this interpretation can never be compatible with 1 millimeter, or even 1 micro-millimeter, higher than 2.00 meters. Lasersohn assumes a context independent fixed measure of height, which is such that it is practically, or even physically, impossible to refine it. But this means that measure phrases like ‘2 meters’ denote (physically) non-extensive points, which has the result that a sentence like (21-a) is (almost) necessarily false (on an exactly-reading of the measure phrase). On this line of reasoning, the actual denotation of an expression in the end doesn’t really count, and the only thing that matters is the pragmatic halo. This conclusion suggests that the assumptions Lasersohn makes are questionable. But it gives rise to an empirical problem as well: Suppose that up to the minimal point we can measure, John and Mary are equally tall. There is nothing incomprehensible about my saying: ‘but had we had a better measurement instrument, we would have found out after all that John is taller than Mary’. Because Lasersohn predicts it be incomprehensible, he seems to make an empirically incorrect prediction.

Lasersohn crucially assumes that measure phrases are interpreted with respect to a *fixed* underlying structure of heights. The obvious alternative – an alternative assumed by Lewis (1979) and I guess by most semanticists –, is to assume that the underlying height-structure is not fixed in advance, but determined by context. One context can require a more fine-grained measurement than another, and what counts as a particular height in one context, counts as an interval of heights in a context that requires more fine-grained measurements. On this alternative view, context determines the underlying measure structure, or level of fine-grainedness, with respect to which sentences should be interpreted, and there is no need for pragmatic halos. Pragmatics is not independent of semantics, but rather helps to determine whether a sentence is true or false. But this means that (21-a) can be true in one context where John is 2.02 meters, but false at another. Now, what is the function of the word ‘exactly’ in (21-b)? Like Lasersohn (1999), we would like to claim that its function is most of all pragmatic. But where for Lasersohn its function is to contract the pragmatic halo, we claim that it is an indication of the fine-grainedness of the underlying structure. If (21-a) is true with respect to an underlying structure where 2.02 meters is not distinguished from 2.00 meters tall, the use of ‘exactly’ in (21-b) may indicate that the speaker has a richer measurement structure in mind.

We have assumed that in a more coarse grained model, measure phrases denote *intervals* rather than points of a more fine-grained model. But how do the intervals denoted by different measurement phrases relate to each other?

Until now we have assumed the simplest view – perhaps suggested by the phrase ‘granularity’. Also in the coarse-grained model, a measure phrase denotes an equivalence class, but such an equivalence class now contains more elements of I than in a fine-grained model. Thus, just like the fine-grained model, also the coarse-grained model is a weak order. This picture might be somewhat simplistic, though. In section 2 we saw already that in physics one assumes that measure phrases allow for a margin of error. But then one can assume that these phrases denote intervals, and this almost immediately means that the denotations of two different measure phrases might overlap, and that there are areas where it is unclear which of two measure phrases can be truly attributed to the points in this area, if not both. Assuming that all measure phrases allow for the same margin of error, the resulting structure will be a semi-order. If different measure phrases might allow for different margins of error, the resulting structure will be an interval order. A structure $\langle I, R \rangle$, with R a binary relation on I , is an interval order just in case R is irreflexive (IR) and satisfies the interval-order (IO) condition.

Definition 3.

An interval order is a structure $\langle I, R \rangle$, with R a binary relation on I that satisfies the following conditions:

$$(IR) \forall x : \neg R(x, x).$$

$$(IO) \forall x, y, v, w : (R(x, y) \wedge R(v, w)) \rightarrow (R(x, w) \vee R(v, y)).$$

In section 6.2 we have defined the refinement relation between models with respect to relation R as follows: M' is an R -refinement of M iff $V_M(R) \subset V_{M'}(R)$. We have seen that this works in case ‘ R ’ denotes a weak order in both models M and M' . However, it is the appropriate definition as well in case it denotes, for instance, an interval order in both models, or if it denotes an interval order in model M and a semi-order, or a weak order in model M' . It is well-known that we can think of a linear order as a refinement of a weak order: every linear order is a weak order, but not every weak order is a linear order. But in the same way, interval orders are refinements of strict partial orders, semi-orders are refinements of interval orders, and weak orders are refinements of semi-orders.

Interval orders are relevant for the analysis of vagueness. In most discussions on vagueness where an indifference, or margin of error principle is involved, this indifference, or margin of error, is standardly taken to be the same for all expressions.⁶³ Jerry Hobbs and Manfred Krifka have recently argued that we should give up this assumption. Hobbs (2000) observes that if there are 920 people at the meeting, the following sentence is intuitively false,

⁶³But see Williamson (1994) for some discussion.

if $X = 980$, but possibly true, in case $X = 1000$:

(22) There are *about* X people at the meeting.

This is surprising, because ‘1000’ is higher than ‘980’, and the latter instantiation of X seems thus closer to the truth. To account for these intuitions, also Hobbs suggests that different objects in the same model need not necessarily have the same grain-size. This is exactly what we see in interval orders. Krifka (2007) has argued that something very similar happens even without the use of an approximator like ‘about’.

It is well-known that (21-a) allows for much more variation than (23):

(21-a) John is 2 meters tall.

(23) John is 2.02 meters tall.

Whereas (21-a) can be true (or appropriate) in case John is exactly 2.03 meters tall, (23) cannot. Krifka (2007) argues that because ‘2.02 meters’ is a more complex expression than ‘2 meters’, it has a more complex meaning (derived by Horn’s division of pragmatic labor). One way to account for this observation is to assume that the complexity of the meaning involves the fine-grainedness of the underlying structure of measurement. Although the denotations of ‘2 meters’ and ‘2.02 meters’ each denotes a point in the measure system used for the interpretation of the sentence (21-a) and (23), respectively, what denotes a point in the measurement system underlying (21-a) denotes an interval (i.e., a set of more fine-grained points) in the measuring system underlying (23). That’s why, according to this analysis, the expression ‘2 meters’ is more vague than the expression ‘2.02 meters’. On this analysis, the underlying ordering relation is always a weak order. But there is another way to account for the same intuition. According to this alternative analysis, both ‘2 meters’ and ‘2.02 meters’ have a denotation *in the same* (coarser-‘grained’) *model* M' such that (21-a) is true in case John is exactly 2.03 meters tall, while (23) is not. This can be accounted for by assuming that numbers denote *intervals*, but that some intervals might be proper parts of others, i.e., that we look at *interval orders*. In our case, what is denoted by ‘2.02 meters’ in coarse-‘grained’ model M' is an interval that is a proper part of the interval denoted by the phrase ‘2 meters’ in model M' .

6.3 Granularity and relevance

Our way to relate different models in terms of a coarsening relation made use of a standard technique. We assume an ordering relation and define an equivalence relation in terms of it. In a coarser-grained model we just let an

equivalence class of individuals of the fine-grained model be represented by a single individual. We did this (in sections 4.5 and 6.2) by means of the ordering *taller than*. Another classical case (e.g. Kamp, 1979) is to define intervals as equivalence classes of simultaneous events defined in terms of the temporal inclusion relation. To define a coarsening relation between models, however, we don't need to start with such an ordering.

Hobbs (1985) argues that to represent or conceptualize the world at a coarser-grained level, we can just restrict ourselves by looking only at the *relevant* predicates of our original language. Consider a model $M = \langle I, V \rangle$ for the first-order language \mathcal{L} , and take \mathcal{L}' to be a sublanguage of \mathcal{L} containing only its 'relevant' predicates. In terms of the monadic predicates of \mathcal{L}' we can now define an equivalence relation ' $\sim_{\mathcal{L}'}$ ' with respect to language \mathcal{L}' : $a \sim_{\mathcal{L}'} b$ iff $a, b \in I_M$ and for all monadic predicates P of \mathcal{L}' : $M \models P(a) \Leftrightarrow M \models P(b)$. In terms of this equivalence relation, Hobbs (1985) proposed to construct a coarse-grained model M' as follows: (i) the domain $I_{M'}$ is just the set of equivalence classes $I_{M'} = \{\{y \in I : y \sim_{\mathcal{L}'} x\} : x \in I_M\}$, and (ii) the valuation function is such that for all monadic predicates $P \in \mathcal{L}'$, $M' \models P([a])$ iff $M \models P(a)$, where $[a]$ denotes the equivalence class containing $a \in I_M$.⁶⁴ Except for examples like (17) and (18), Hobbs (1985) suggests that in this way we can also account for some examples discussed by Nunberg (1985).

Nunberg (1985) has argued that an account of definite reference and the use of the phrase 'the same' is greatly simplified if we assume that people construct models of what is going on at different levels of granularity. An intuitive account of definite reference says that 'my P' refers to the *unique* (relevant) individual with property P owned by me. An example like 'My leg hurts' seems an obvious counterexample, however, because people have (usually) two legs. Nunberg (1985) suggests that the distinction between the two legs is not pragmatically relevant and can be represented by a single object in a more coarse-grained model, and thus satisfying the uniqueness requirement after all. Perhaps slightly less controversial is his proposed use of levels of granularity to give a semantics for 'the same'. It is obvious that we can use this phrase to denote not only *token*-identity, but *type*-identity as well. Thus, we can say

(24) I own a Ford Falcon. The same car is owned by Enzo.

meaning that Enzo and I owned the same *type* of car, not the same *token*. Nunberg (1985) and Hobbs (1985) suggest to account for *type*-identity as identity at a more coarse-grained level of description. In this way, they propose, we can account for the fact that we cannot say 'A Ford Falcon was heading

⁶⁴Van Lambalgen (2001) noted that this construction needs to be generalized, because it does not generalize to predicates of higher arity.

south on U.S. 101, went out of control, and crashed into the same car' to mean that it hit another Ford Falcon. The reason is that *type*-level identity is just indistinguishability, but only restricted to distinguishable predicates that are *relevant*.⁶⁵

Hobbs' suggestion is very appealing. Unfortunately, Lasersohn (2000) showed that the truth definition at the coarse-grained level proposed by Hobbs (1985) does not capture the intuitive motivation. It is clear that discourse (24) should be interpreted with respect to a coarse-grained model. According to Hobbs' construction, $M' \leq M$ just in case if for every monadic predicate $P \in \mathcal{L}'$, if $P([a])$ is true in coarse-grained model M' , it has to be the case that $P(b)$ is true in fine-grained model M , for *every* $b \in [a]$. However, it is clear that in (24) the predicates 'Owned by me' and 'Owned by Enzo' are relevant, and thus part of \mathcal{L}' . Because in M' it is the same car that has both of these properties, Hobbs' construction predicts that every token of this car should have both properties in M as well. This, however, is obviously not the case. Lasersohn (2000) concludes that we should thus not account for the type-token distinction in terms of levels of granularity. Lasersohn might well be right with this conclusion, but this doesn't mean that Hobbs' (1985) project could not be saved by modifying Hobbs' ordering relation between models of different grain. Instead of making use of *universal* quantification as proposed by Hobbs (1985), we *could* make use of *existential* quantification. In that case we should not *define* the domain of the coarse-grained model M' in terms of equivalence classes of M as suggested by Hobbs, but rather take it to be a primitive domain. Moreover, we assume a surjective function f from the domain of M , I_M , to the domain of coarser-grained model M' , $I_{M'}$,^{66,67} that preserves (but not necessarily anti-preserved) each relevant predicate P : if $x \in P_M$, then $f(x) \in P_{M'}$, but if $x \notin P_M$, it need not be that $f(x) \notin P_{M'}$.⁶⁸ The other direction, however, follows by contraposition: if $x \notin P_{M'}$, then there is no $y \in f^{-1}(x)$ such that $y \in P_M$. In this case it is not problematic that the predicates 'Owned by me' and 'Owned by Enzo' are relevant, and thus part of \mathcal{L}' . It is just important that their negatives like 'Not owned by me' and 'Not owned by Enzo' are not part of \mathcal{L}' . To capture the idea of simplification, or coarsening, it is natural to assume that f is *not injective*: it might be that $f(x) = f(y)$, although $x \neq y$. Of course, we want refinements to preserve all the predicates and

⁶⁵According to Tim Williamson (p.c.), a sentence like 'Enzo owns two of the same car' is ok. This would be problematic for the Nunberg/Hobbs proposal.

⁶⁶A function f from D to D' is surjective iff the range of f is D' .

⁶⁷If one desires, one can think of this function as a counterpart function used in quantified modal logic.

⁶⁸Technically, f is just a homomorphism from M to M' . In general, homomorphisms don't preserve negative sentences.

relations of the restricted language \mathcal{L}' , but this preservation is now stated as follows: $M' \leq_{\mathcal{L}'} M$ just in case if $x \in P_{M'}$, then $\exists y \in f^{-1}(x) \in P_M$, for each $P \in \mathcal{L}'$.⁶⁹ To a certain extent, the refinement relation used here is similar to the refinement relation used in supervaluation theory. But these refinement relations are obviously not the same: in contrast to what we assumed now, in supervaluation theory one assumes a one-to-one relation between the objects that exist in the different models.⁷⁰

The question how to relate the truth conditions of the same sentence with respect to models of different levels of granularity is a deep one, and I am not aware of any *general* satisfying answer. The question is closely connected with the issue in tense logic how to relate truth conditions of sentences with respect to *intervals* or events, to that of models that take *instants* to be basic. For ‘John slept’ to be true in interval i , is it enough that he slept at *some* instant in i , or should he have been sleeping in *most*, or even *all* instants in i ? Although the first proposal seems very weak, Kamp (1979) convincingly argued that the latter proposal is much too strong for a sentence like ‘John wrote an article’. Kamp concludes that we should not reduce judgments about intervals to judgments about instants. More in general, according to a tradition going back to Russell (1914) and Wiener (1914), and taken up by Kamp (1979), van Benthem (1984), and Thomason (1984), we should rather try to reduce our talk about instants to talk about intervals and events: absolute time does not exist independently of our experiences and is at best a construction out of them.⁷¹ Notice how much this view differs not only from Lasersohn’s (1999) position, but also from all those analyses that propose to account for the meaning of ‘tall’ in terms of exact cutoff-points.

7 Conclusion

In this paper I focussed my discussion on relative adjectives. The meaning of a relative adjective is context dependent. I accounted for this context dependence by assuming that such adjectives should be interpreted with respect to

⁶⁹In general, the truth conditions of sentences in course-grained model M' are defined in terms of their truth conditions in fine-grained model M as follows:

$$M', g \models P(x) \quad \text{iff} \quad \exists d \in f^{-1}(\llbracket x \rrbracket^{M',g}) : M, g[x/d] \models P(x)$$

$$M', g \models \neg\varphi \quad \text{iff} \quad M', g \not\models \varphi$$

$$M', g \models \varphi \wedge \psi \quad \text{iff} \quad M', g \models \varphi \text{ and } M', g \models \psi$$

$$M', g \models \forall x\varphi \quad \text{iff} \quad \text{for all } d \in I_{M'} : M', g[x/d] \models \varphi.$$

Notice that $M', g \models \neg P(x)$ iff $\forall d \in f^{-1}(\llbracket x \rrbracket^{M',g}) : M, g[x/d] \not\models P(x)$.

⁷⁰Instead, the system described here is much closer to the ‘inverse system’ used by Van der Does and Van Lambalgen (2000) to account for the logic of vision.

⁷¹In fact, this tradition goes back all the way to Leibniz. Technically, the actual world, the *ding an sich*, is seen as the inverse limit.

so-called comparison classes. In terms of the interpretation of adjectives with respect to comparison classes, I characterized the difference between relative and absolute adjectives in terms a difference of constraints these two types of adjectives have to fulfill. More importantly, I have argued that the notion of ‘tolerance’ can be accounted in terms of Luce’s semi-orders. Semi-orders, in turn, can be characterized in terms of constraints on comparison classes as well. I have argued that the level of tolerance a predicate allows for depends on the context of use.⁷² If the stakes are higher, a predicate allows for less tolerance. The level of tolerance, in turn, determines how precise one can be. In this way, the discussion in section 6 is linked closely to the discussion in section 5.

In this paper I focussed on the context dependence and vagueness of *adjectives*, and tried not to make use of numbers, like degrees, or probabilities. I suggested that many other vague expressions can be analyzed by make use of the same tools, but I don’t think it can do all of the work. The most natural way to account for vague quantifiers like ‘many’ and ‘few’ (cf. Fernando & Kamp, 1996), and the use of adjectives in *exclamatives* like ‘How tall you are!’ (Graff, 2000), for instance, seems to be by making use of probabilities. Whether we are forced to take probabilities into account to account for such cases remains to be seen.

References

- [85] Arrow, K. (1959), ‘Rational choice functions and orderings’, *Economica*, **26**: 121-127.
- [2] Asher, N. and L. Vieu (1995), ‘Towards a geometry of common sense: A semantics and complete axiomatization of Mereotopology’, *Proceedings of IJCAI-95*, Montreal, pp. 846-852.
- [3] Barker, C. (2002), ‘The dynamics of vagueness’, *Linguistics and Philosophy*, **25**: 1-36.
- [4] Benthem, J. van (1982), ‘Later than late: on the logical origin of the temporal order’, *Pacific Philosophical Quarterly* **63**: 193-203.
- [85] Bierwisch, M. (1984), ‘The semantics of gradation’, in Bierwisch, M. and E. Lang (eds.), *Dimensional Adjectives*, pp. 71-262, Springer Verlag.
- [85] Bolinger, D. (1972), *Degree words*, The Hague/Paris: Mouton.

⁷²In a sense this is just what Parikh (1994) argued for as well.

- [85] Bosch, P. (1983), ‘Vagueness’ is context-dependence. A solution to the Sorites paradox’, in T. Ballmer and M. Pinkal (eds.), *Approaching Vagueness*, North-Holland Linguistic Series 50, North Holland, Amsterdam, pp. 189-210.
- [85] Burns, L. C. (1991), *Vagueness: An investigation into Natural Language and the Sorites Paradox*, Kluwer, Dordrecht.
- [85] Cresswell, M. (1976), ‘The semantics of degree’, in Partee, B. et al. (eds.), *Montague Grammar*, pp. 261-292, Academic Press.
- [85] Cruse, D.A. (1986), *Lexical Semantics*, Cambridge University Press, Cambridge.
- [85] Deemter, K. van (1995) ‘The Sorites fallacy and the context-dependence of vague predicates’, in M. Kanazawa, et al., *Quantifiers, Deduction, and Context*, CSLI Publications, Stanford, Ca., pp. 59-86.
- [85] Does, J. van der and M. van Lambalgen (2000), ‘A logic of vision’, *Linguistics and Philosophy*, **23**: 1-92.
- [85] Dubois, D. & H. Prade (1980), *Fuzzy Sets and Systems. Theory and Applications*, Academic Press, New York.
- [85] Dummett, M. (1975), ‘Wang’s paradox’, *Synthese*, **30**: 301-324.
- [85] Ellis, J. (2004), ‘Context, indexicals, and the Sorites’, *Analysis*, bf 64: 362-364.
- [85] Fernando, T. and H. Kamp (1996), *Proceedings of Salt 6*.
- [85] Fine, K. (1975), ‘Vagueness, truth and logic’, *Synthese*, **30**: 265-300.
- [85] Gaifman, H. (1997), ‘Vagueness, tolerance and contextual logic’, manuscript, Columbia University.
- [85] Goodman, (1966), *The structure of appearance*, Bobbs-Merill Publ.
- [85] Graff, D. (2000), ‘Shifting sands: An interest relative theory of vagueness’, *Philosophical Topics*, **28**: 45-81.
- [85] Hobbs, J. (1985), ‘Granularity’, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-85)*.
- [85] Hobbs, J. (2000), ‘Half orders of magnitude’, manuscript.

- [85] Jaegher, K. de (2003), ‘A game-theoretic rationale for vagueness’, *Linguistics and Philosophy*, **26**: 637-659.
- [85] Jespersen, O. (1917), ‘Negation in English and other languages’, in *Selected writings of Otto Jespersen*, London, Tokyo.
- [85] Kamp, H. (1975), ‘Two theories of adjectives’, in E. Keenan (ed.), *Formal Semantics of Natural Language*, Cambridge University Press, Cambridge, pp. 123-155.
- [85] Kamp, H. (1979), ‘Events, instants and temporal reference’, in R. Bäerle et al. (eds.), *Semantics from Different Points of View*, Springer-Verlag, Berlin, pp. 376-417.
- [85] Kamp, H. (1981), ‘The paradox of the heap’, in U. Mönnich (ed.), *Aspects of Philosophical Logic*, D. Reidel, Dordrecht, pp. 225-277.
- [85] Kamp, H. and B. Partee (1995), ‘Prototype theory and compositionality’, *Cognition*, **57**: 129-191.
- [85] Keefe, R. and P. Smith (1996), *Vagueness. A Reader*, MIT Press, Cambridge, Massachusetts.
- [85] Keefe, R. (2000), *Theories of Vagueness*, Cambridge University Press, Cambridge.
- [85] Keefe, R. (2007), ‘Vagueness without context change’, *Mind*, **116**: 462.
- [85] Kennedy, C. (1999), *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. Garland Press, New York.
- [85] Kennedy, C. (2007), ‘Vagueness and grammar: the semantics of relative and absolute gradable adjectives’, *Linguistics and Philosophy*.
- [85] Kennedy, C. and L. McNally (2005), ‘Scale structure and the semantic typology of gradable predicates’, *Language*, **81**: 1-37.
- [85] Klein, E. (1980), ‘The semantics of positive and comparative adjectives’, *Linguistics and Philosophy*, **4**: 1-45.
- [85] Klein, E. (1991), ‘Comparatives’, in von Stechow, A. and D. Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, Walter de Gruyter.
- [85] Krantz, D., R. Duncan Luce, P. Suppes, and T. Tversky (1971), *Foundations of Measurement, Vol. 1*, Academic Press, New York and London.

- [85] Krifka, M. (2007), ‘Approximate interpretation of number words: A case of strategic communication’, in G. Bouma et al. (eds.), *Cognitive Foundations of Interpretation*, Proceedings of the KNAW colloquium, Amsterdam, pp. 111-126.
- [85] Kyburg, A. and M. Morreau (2000), ‘Fitting worlds: vague language in context’, *Linguistics and Philosophy*, **23**: 577-597.
- [85] Lakoff, G. (1970), ‘A note on vagueness and ambiguity’, *Linguistic Inquiry*, **1**: 357-359.
- [85] Lakoff, G. (1973), ‘Hedges: a study in meaning criteria and the logic of fuzzy concepts’, *Journal of Philosophical Logic*, **2**: 458-508.
- [85] Lambalgen, M. van (2001), ‘Conditional quantification, or poor man’s probability’, *Journal of Logic and Computation*, **11**: 295-335.
- [85] Lasersohn, P. (1999), ‘Pragmatic Halos’, *Language*, **75**: 522-551.
- [85] Lasersohn, P. (2000), ‘Same, Models and representations’, In B. Jackson & T. Matthew (eds.), *Proceedings from SALT X*, pp. 83-97, CLC Publications.
- [85] Lewis, D. (1969), *Convention*, Harvard University Press, Cambridge, Mass.
- [85] Lewis, D. (1970), ‘General semantics’, *Synthese*, **22**: 18-67.
- [85] Lewis, D. (1973), *Counterfactuals*, Blackwell, Oxford.
- [85] Lewis, D. (1979), ‘Scorekeeping in a language game’, *Journal of Philosophical Logic*, **8**: 339-359.
- [85] Lewis, D. (1993), ‘Many, but almost one’, in K. Campbell et al (eds.), *Ontology, Causality, and Mind: Essays on the Philosophy of D.M. Armstrong*, Cambridge, Cambridge University Press.
- [85] Luce, R. D. (1956), ‘Semiordeers and a theory of utility discrimination’, *Econometrica*, **24**: 178-191.
- [85] Ludlow, P. (1989), ‘Implicit comparison classes’, *Linguistics and Philosophy*, **12**: 519-533.
- [85] Myerson, R. B. (1991), *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge MA.

- [85] Nunberg, G. (1985), ‘Individuation in Context’, *Proceedings of WCCFL*, Santa Cruz, California, pp. 203-217.
- [85] Parikh, R. (1994), ‘Vagueness and utility: The semantics of common nouns’, *Linguistics and Philosophy*, **17**: 521-535.
- [85] Pinkal, M. (1984), ‘Consistency and context change: The Sorites paradox’, in F. Landman & F. Veltman (eds.), *Varieties of Formal Semantics, Proceedings of the fourth Amsterdam Colloquium*, Groningen-Amsterdam Studies in Semantics, Foris Publications, Dordrecht.
- [85] Pinkal, M. (1995), *Logic and Lexicon*, Studies in Linguistics and Philosophy, Kluwer, Dordrecht.
- [85] Raffman, D. (1994), ‘Vagueness without paradox’, *The Philosophical Review*, **103**, 41-74.
- [85] Raffman, D. (1996), ‘Vagueness and Context-Relativity’, *Philosophical Studies*, **81**: 175-192.
- [85] Rooij, R. van (2009), ‘Revealed preference and satisficing behavior’, in G. Bonano et. al. (eds.), *Proceedings of LOFT 2008*, Amsterdam University Press.
- [85] Rothstein, C. and Y. Winter (2004), ‘Total adjectives vs. partial adjectives’, *Natural Language Semantics*, **12**: 259-288.
- [85] Russell, B. (1914), *Our Knowledge of the External World*, Chicago, Norton.
- [85] Russell, B. (1940), *An Inquiry into Meaning and Truth*, George Allen and Unwin, London.
- [85] Sapir, E. (1944), ‘Grading: A study in semantics’, *Philosophy of Science*, **11**: 93-116.
- [85] Sassoon, G. (2008), *Vagueness, Gradability and Typicality*, PhD thesis, Tel Aviv University.
- [85] Sauerland, U. and D. Penka (2007), ‘Scalar vs. epistemic vagueness’, in M. Gibson & T. Friedman (eds.), *Proceedings SALT 17*, Ithaca, N.Y, CLC Publications.
- [85] Scott, D. and P. Suppes (1958), ‘Foundational aspects of theories of measurement’, *Journal of Symbolic Logic*, **23**: 113-128.

- [85] Schwarzschild, R. (2005), ‘Measure phrases as modifiers of adjectives’, *Recherches Linguistiques de Vincennes*, **35**: 207-228.
- [85] Sedivy, J. M. Tannenhaus, C. Chambers, and G. Carlson (1999), ‘Achieving incremental semantic interpretation through contextual representation’, *Cognition*, **71**: 109-147.
- [85] Seuren, P. (1973), ‘The comparative’, in Kiefer, F. and N. Ruwet (eds.), *Generative Grammar in Europe*, pp. 528-564, Reidel.
- [85] Shapiro, S. (2006), *Vagueness in Context*, Clarendon Press, Oxford.
- [71] Stalnaker, R. (1978), ‘Assertion’, In: P. Cole (ed.), *Syntax and Semantics, vol. 9: Pragmatics*, pp. 315-332.
- [85] Stanley, J. (2003), ‘Context, interest relativity, and the Sorites’, *Analysis*,
- [85] Soames, S. (1999), *Understanding Truth*, Oxford University Press, New York.
- [85] Stechow, A. von (1984), ‘Comparing semantic theories of comparison’, *Journal of Semantics*, **3**: 1-77.
- [85] Suppes, P. and J.L. Zinnes (eds.), *Basic Measurement Theory*, Handbook of Mathematical Psychology, Willey, New York.
- [85] Thomason, R. (1984), ‘On constructing instants from events’, *Journal of Philosophical Logic*, **13**: 85-96.
- [85] Unger, P. (1975), *Ignorance*, Clarendon Press, Oxford.
- [85] Veltman, F. (1987), *Syllabus Logische Analyse 2: Vaagheid*, Universiteit van Amsterdam.
- [85] Veltman, F. (2002), *Het verschil tussen vaag en precies*, Inaugural lecture, Amsterdam University Press, Amsterdam.
- [85] Waismann, F. (1968), ‘Verifiability’, in A. Flew (ed.), *Logic and Language*, Basic Blackwell, Oxford.
- [85] Wheeler, S. C. (1972), ‘Attributives and their modifiers’, *Nous*, **6**: 310-334.
- [85] Wiener, N. (1914), ‘A contribution to the theory of relative position’, *Proc. Camb. Phil. Soc.*, **17**: 441-449.

- [85] Williamson, T. (1994), *Vagueness*, Routledge, London.
- [85] Wilson, D. and D. Sperber (2002), 'Truthfulness and relevance', *Mind*: **443**: 583-632.
- [85] Wright, C. (1975), 'On the coherence of vague predicates', *Synthese*, **30**: 325-365.