23. Ambiguity and Vagueness: An Overview

Christopher Kennedy
Department of Linguistics
University of Chicago
1010 E. 59th St.
Chicago, IL 60637 USA
ck@uchicago.edu
+1.773.834.1988

May 5, 2009

23. Ambiguity and Vagueness: An Overview

Ambiguity and vagueness are two varieties of interpretive uncertainty which are often discussed together, but are distinct both in their essential features and in their significance for semantic theory and the philosophy of language. Ambiguity involves uncertainty about mappings between levels of representation with different structural characteristics, while vagueness involves uncertainly about the actual meanings of particular terms. This article examines ambiguity and vagueness in turn, providing a detailed picture of their empirical characteristics and the diagnostics for identifying them, and explaining their significance for theories of meaning. Although this article continues the tradition of discussing ambiguity and vagueness together, one of its goals is to emphasize the ways that these phenomena are distinct in their empirical properties, in the factors that give rise to them, and in the analytical tools that can be brought to bear on them.

1. Interpretive uncertainty

Most linguistic utterances display "interpretive uncertainty", in the sense that the mapping from an utterance to a meaning (Grice's 1957 "meaning$_{NN}$") appears (at least from the external perspective of a hearer or addressee) to be one-to-many rather than one-to-one. Whether the relation between utterances and meanings really is one-to-many is an open question, which has both semantic and philosophical

significance, as I will outline below. What is clear, though, is that particular strings of phonemes, letters or manual signs used to make utterances are, more often than not, capable of conveying distinct meanings.

As a first step, it is important to identify which kinds of interpretive uncertainty (viewed as empirical phenomena) are of theoretical interest. Consider for example (1a-b), which manifest several different kinds of uncertainty.

(1) a. Sterling's cousin is funny.

b. Julian's brother is heavy.

One kind concerns the kinds of individuals that the English noun phrases *Sterling's cousin* and *Julian's brother* can be used to pick out: the former is compatible with Sterling's cousin being male or female and the latter is compatible with Julian's brother being older or younger than him. However, this sort of uncertainty merely reflects the fact that *cousin* and *brother* are INDETERMINATE with respect to sex and age, respectively: both terms have conditions of application that specify certain kinds of familial relationships, and *brother*, unlike *cousin*, also imposes conditions on the sex of the individuals it applies to, but beyond these constraints these terms do not discriminate between objects as a matter of meaning. (Indeterminacy is also sometimes referred to as "generality"; see Zwicky & Sadock 1975, 1984, Gillon 1990, 2004.)

That this is so can be seen from the fact that distinctions of this sort do not affect judgments of truth or falsity. For example, assuming that the antecedents of the conditionals in (2a-b) specify the minimal difference between the actual world and the counterfactual worlds under consideration, the fact that (2a) is false and (2b) is true shows that a change in sex, unlike a change in familial relationships, does not affect the truth of the application of *cousin*.

(2) Lily is Sterling's cousin....

a. ...but if she were a boy, she wouldn't be Sterling's cousin anymore.

b. ...but if her mother weren't Sterling's father's sister, she wouldn't be Sterling's cousin anymore.

Likewise, the hypothetical change in age in (3a) doesn't affect the truth of the application of *brother*, though the change in sex in (3b) now does make a difference.

(3) Sterling is Julian's brother...

a. ...but if their ages were reversed, he wouldn't be Julian's brother anymore.

b. ...but if he were a girl, he wouldn't be Julian's brother anymore.

Indeterminacy reflects the fact that the meaning of a word or phrase typically does not involve an exhaustive specification of the features of whatever falls under its extension; some features are left open, resulting in the sort of flexibility of

application we see above. This is not to say that such features couldn't be specified: many languages contain cousin terms that do make distinctions based on sex (either through grammatical gender, as in Italian *cugino* 'cousin$_{masc}$' vs. *cugina* 'cousin$_{fem}$', or lexically, as in Norwegian *fetter* 'male cousin' vs. *kusine* 'female cousin'), and some contain male sibling terms that specify age (such as Mandarin *gēge* 'older brother' vs. *dìdi* 'younger brother'). Whether a particular distinction is indeterminate or not is thus somewhat arbitrary and language specific, and while it might be interesting to determine if there are cultural or historical explanations for the presence/absence of such distinctions in particular languages, the existence of indeterminacy in any single language is typically not a fact of particular significance for its semantic analysis.

A second type of uncertainty manifested by (1a) and (1b) is of much greater importance for semantic analysis, however, as it involves variability in the truth or satisfaction conditions that a particular bit of an utterance introduces into the meaning calculation. This kind of uncertainty is AMBIGUITY, which manifests itself as variation in truth conditions: one and the same utterance token can be judged true of one situation and false of another, or the other way around, depending on how it is interpreted. In (1a) and (1b), we see ambiguity in the different ways of understanding the contributions of *funny* and *heavy* to the truth conditions. (1a) can be construed either as a claim that Sterling's cousin has an ability to make people

laugh ("funny ha-ha") or that she tends to display odd or unusual behavior ("funny strange"). Similarly, (1b) can be used to convey the information that Julian's brother has a relatively high degree of weight, or that he is somehow serious, ponderous, or full of gravitas. That these pairs of interpretations involve distinct truth conditions is shown by the fact that we can use the same term (or, more accurately, the same bits of phonology) to say something that is true and something that is false of the same state of affairs, as in (4) (Zwicky & Sadock 1975).

(4) Sterling's cousin used to make people laugh with everything she did, though she was never in any way strange or unusual. She was funny without being funny. Lately, however, she has started behaving oddly, and has lost much of her sense of humor. Now she's funny but not very funny.

Both examples also manifest an ambiguity in the nature of the relation that holds between the genitive-marked nominal in the possessive construction and the denotation of the whole possessive DP (see article 46 *Possessives and Relational Nouns*). While the most salient relations are the familial ones expressed by the respective head nouns (*cousin of* and *brother of*), it is possible to understand these sentences as establishing different relations. For example, if Julian is one of several tutors working with a family of underachieving brothers, we could use (1b) as a way of saying something about the brother who has been assigned to Julian, without in

any way implying that Julian himself stands in the *brother of* relation to anyone. (He could be an only child.)

Even after we settle on a particular set of conditions of application for the ambiguous terms in (1a) and (1b) (e.g. that we mean "funny ha-ha" by *funny* or are using *heavy* to describe an object's weight), a third type of uncertainty remains about precisely what properties these terms ascribe to the objects to which they are applied, and possibly about whether these terms can even be applied in the first place. This is VAGUENESS, and is of still greater significance for semantic theory, as it raises fundamental questions about the nature of meaning, about deduction and reasoning, and about knowledge of language.

Consider, an utterance of (1b) in a context in which we know that *heavy* is being used to characterize Julian's brother's weight. If we take the person who utters this sentence to be speaking truthfully, we may conclude that Julian's brother's weight is above some threshold. However, any conclusions about how heavy he is will depend on a range of other contextual factors, such as his age, his height, information about the individuals under discussion, the goals of the discussion, the interests of the discourse participants, and so forth, and even then will be rough at best. For example, if we know that Julian's brother is a 4-year old, and that we're talking about the children in his preschool class, we can conclude from an utterance of (1b) that his weight is somewhere above some threshold, but it would

be extremely odd to follow up such an utterance by saying something like (5).

(5) Well, since that means he is at least 17.5 kg, we need to make sure that he is one of the carriers in the piggy-back race, rather than one of the riders.

(5) is odd because it presumes a specific cut-off point separating the heavy things from the non-heavy things (17.5 kg), but the kind of uncertainty involved in vagueness is precisely uncertainty about where the cut off is.

This can be further illustrated by a new context. Imagine that we are in a situation in which the relevant contextual factors are clear: Julian's brother is a 4-year old, we're talking about the children in his class, and we want to decide who should be the anchor on the tug-of-war team. In addition, we also know that Julian's brother weighs exactly 15.2 kg. Even with these details spelled out — in particular, even with our knowledge of Julian's brother's actual weight — we might still be uncertain as to whether (1b) is true: Julian's brother is a BORDERLINE CASE for truthful application of the predicate.

Borderline cases and uncertainty about the boundaries of a vague predicate's extension raise significant challenges for semantic theory. If we don't (and possibly can't) know exactly how much weight is required to put an object in the extension of *heavy* (in a particular context of use), even when we are aware of all of the potentially relevant facts, can we truly say that we know the meaning of the term?

Do we have only incomplete knowledge of its meaning? If our language contained only a few predicates like *heavy*, we might be able to set them aside as an interesting but ultimately insignificant puzzles. Vagueness is pervasive in natural language, however, showing up in all grammatical categories and across lexical fields, so understanding the principles underlying this type of uncertainty is of fundamental importance for semantic theory.

In the rest of this article, I will take a closer look at ambiguity and vagueness in turn, providing a more detailed picture of their empirical characteristics and the diagnostics for identifying them, and explaining their significance for theories of meaning. Although this article follows in a long tradition of discussing ambiguity and vagueness together, a goal of the article is to make it clear that these phenomena are distinct in their empirical properties, in the factors that give rise to them, and in the analytical tools that can be brought to bear on them. However, both present important challenges for semantics and philosophy of language, and in particular, for a compositional, truth conditional theory of meaning.

2. Ambiguity

2.2 Varieties of ambiguity

Ambiguity is associated with utterance chunks corresponding to all levels of linguistic analysis, from phonemes to discourses, and is characterized by the association of a single orthographic or phonological string with more than one meaning. Ambiguity can have significant consequences, for example if the wording of a legal document is such that it allows for interpretations that support distinct judgments. But it can also be employed for humorous effect, as in the following examples from the 1980s British comedy series *A Bit of Fry and Laurie* (created by Stephen Fry and Hugh Laurie).

(6) FRY: You have a daughter, I believe.

LAURIE: Yeah, Henrietta.

FRY: Did he? I'm sorry to hear that. That must've hurt.

(5) illustrates a case of PHONOLOGICAL AMBIGUITY, playing on the British comedians' pronunciations of the name *Henrietta* and the sentence *Henry ate her*. (6) makes use of the LEXICAL AMBIGUITY between the name *Nancy* and the British slang term *nancy*, which means weak or effeminate when used as an adjective.

(7) FRY: Something I've always been meaning to ask you: How did you manage

to keep Nancy for so long?

LAURIE: I've never been nancy, John.

Sometimes the humor is unintended, as in the classified advertisement in (8) (cited in Pinker 1994, p. 102).

(8) FOR SALE: Mixing bowl set designed to please a cook with round bottom for efficient beating.

This example illustrates a case of STRUCTURAL AMBIGUITY: whether the cook or the mixing bowl set has a round bottom (and whether the round bottom supports efficient beating of eggs, flour, etc. or efficient beating of the cook) depends on the structural relationships among the constituents of the sentence, in particular whether *with a round bottom* is parsed as a syntactic modifier of the nominal headed by *mixing bowl set* or the one headed by *cook*.

SCOPE AMBIGUITY is illustrated by (9), which can have either the interpretation in (9a) or the one in (9b), depending on whether the quantifier *every chef* is understood as taking scope above or below negation.

(9) Every chef wasn't a madman.

a. No chef was a madman.

b. Not every chef was a madman.

This example is actually part of a larger chunk of discourse in which it becomes clear that the intended interpretation is (9b):

(10) Every chef wasn't a madman. Most weren't, in fact. But many were and are, and the very best chefs, I knew, as I wrote my book at what my chef, Chef Pardus, would call production speed, were a little twisted in the dark spaces of their brain. (From Michael Rulman, *Soul of a Chef*, p. 133)

But the sentence could also be used to make the stronger claim paraphrased in (9a), indicating a real truth conditional distinction. Scope ambiguities involving quantifiers and other logical expressions (negation, other quantifiers, modals, intensional verbs, and so forth) have played a significant role in linguistic theory, since different methods of accounting for them involve different assumptions about the syntax-semantics interface (see article 6 *Compositionality*, article 69 *Scope and Binding*, and article 92 *Syntax and Semantics*), a point I will come back to in more detail below.

2.2 Testing for ambiguity

Zwicky & Sadock (1975) provide comprehensive discussion of several different tests for ambiguity, some of which distinguish particular types of ambiguity from each other; here I will focus on the two tests that are most commonly employed in semantic argumentation. The first and most straightforward test is what Zwicky and

Sadock call the TEST OF CONTRADICTION, which involves determining whether the same string of words or phonemes (modulo the addition/subtraction of negation) can be used to simultaneously affirm and deny a particular state of affairs. We saw this test illustrated with *funny* in (4) above; the fact that (11) is read as a true contradiction shows that a merely indeterminate term like *cousin* does not allow for this.

(11) Lily's mother is Sterling's father's sister. Since Lily is a girl, she is Sterling's cousin but not his cousin.

A second important test involves IDENTITY OF SENSE ANAPHORA, such as ellipsis, anaphoric *too*, pronominalization and so forth. As pointed out by Lakoff (1970), such relations impose a parallelism constraint that requires an anaphoric term and its antecedent to have the same meaning (see article 78 *Ellipsis*, for a discussion of parallelism in ellipsis), which has the effect of reducing the interpretation space of an ambiguous expression. For example, consider (12), which can have either of the truth-conditionally distinct interpretations paraphrased in (12a) and (12b), depending on whether the subject *the fish* is associated with the agent or theme argument of the verb *eat*.

(12) The fish is ready to eat.

a. The fish is ready to eat a meal.

b. The fish is ready to be eaten.

When (12) is the antecedent in an identity of sense anaphora construction, as in (13a-b), the resulting structure remains two-ways ambiguous; it does not become four-ways ambiguous: if the fish is ready to do some eating, then the chicken is too; if the fish is ready to be eaten, then the chicken is too.

(13) a. The fish is ready to eat, and the chicken is ready to eat too.

b. The fish is ready to eat, but the chicken isn't.

That is, these sentences do not have understandings in which the fish is an agent and the chicken is a theme, or vice-versa.

This fact can be exploited to demonstrate ambiguity by constructing test examples in such a way that one of the expressions in the identity of sense relation is compatible only with one interpretation of the ambiguous term or structure. For example, since potatoes are not plausible agents of an eating event, the second conjunct of (14a) disambiguates the first conjunct towards the interpretation in (12b). In contrast, (14b) is somewhat odd, because children are typically taken to be agents of eating events, rather than themes, but the context of a meal promotes the theme- rather than agent-based interpretation of the first conjunct, creating a conflict. This conflict is even stronger in (14c), which strongly implies either agentive potatoes or partially-cooked children, giving the whole sentence a certain dark humour.

(14) a. The fish is ready to eat, but the potatoes are not.

b. ?The fish is ready to eat, but the children are not.

c. ??The potatoes are ready to eat, but the children are not.

Part of the humorous effect of (14c) is based on the fact that a sensible interpretation of the sentence necessitates a violation of the default parallelism of sense imposed by ellipsis, a parallism that would not be required if there weren't two senses to begin with. The use of such violations for rhetorical effect is referred to as SYLLEPSIS, or sometimes by the more general term ZEUGMA.

2.2 Ambiguity and semantic theory

Ambiguity has played a central role in the development of semantic theory by providing crucial data for both building and evaluating theories of lexical representation and semantic composition. Cases of ambiguity are often "analytical choice points" which can lead to very different conclusions depending on how the initial ambiguity is evaluated. For example, whether scope ambiguities are taken to be structural (reflecting different Logical Forms), lexical (reflecting optional senses, possibly derived via type-shifting), or compositional (reflecting indeterminacy in the application of composition rules) has consequences for the overall architecture of a theory of the syntax-semantics interface, as noted above.

Consider also the ambiguity of adjectival modification structures such as (15)

(first discussed in detail by Bolinger 1967; see also Siegel 1976; McConnell-Ginet 1982; Cinque 1993; Larson 1998 and article 57 *Adjectives*), which is ambiguous between the "intersective" reading paraphrased in (15a) and the "nonintersective" one paraphrased in (15b).

(15) Olga is a beautiful dancer.

a. Olga is a dancer who is beautiful.

b. Olga is a dancer who dances beautifully.

Siegel (1976) takes this to be a case of lexical ambiguity (in the adjective *beautiful*), and builds a theory of adjective meaning on top of this assumption. In contrast, Larson (1998) argues that the adjectives themselves are unambiguous, and shows how the different interpretations can be accommodated by hypothesizing that nouns (like verbs) introduce a Davidsonian event variable, and that the adjective can take either the noun's individual variable or its event variable as an argument. (The first option derives the interpretation in (15a); the second derives the one in (15b).)

In addition to playing an important methodological role in semantic theory, ambiguity has also been presented as a challenge for foundational assumptions of semantic theory. For example, Parsons (1973) develops an argument which aims to show that the existence of ambiguity in natural language provides a challenge for the hypothesis that sentence meaning involves truth conditions (see Saka 2007 for a

more recent version of this argument). The challenge goes like this. Assume that a sentence S contains an ambiguous term whose different senses give rise to distinct truth conditions p and q (as in the case of (15)), such that (16a-b) hold.

(16) a. S is true if and only if p

b. S is true if and only if q

c. p if and only if q

But (16a-b) mutually entail (16c), which is obviously incorrect, so one of our assumptions must be wrong; according to Parsons, the problematic assumption is the one that the meaning of S can be stated in terms of Tarskian truth definitions like (16a-b). In other words, the presence of ambiguity in natural language shows that sentence meaning is not truth conditional. Note that this argument extends to any theory in which truth conditions are a part of sentence meaning, as in mainstream semantic theory for example, where truth conditions are joined by presuppositions, implicatures, expressive meaning, context change potential and possibly other kinds of information. Even if we adopt this richer view of sentence meaning, it is still the case that truth conditions constitute both an analytical and methodological foundation, playing crucial roles in the way that we go about building hypotheses about semantic competence and constructing the data that we use to test them. Parson's challenge based on ambiguity is therefore an important one.

This sketch of the challenge is not quite complete, however, because it omits the crucial fact that an argument based on (16a-c) has weight only relative to specific assumptions about the grammatical principles that regulate the mapping between sound (or manual gestures or orthography, depending on the modality of communication) and meaning. (16) hides the fact that S is a syntactic object (something Parsons accepts), and as such needs to be mapped both to a meaning of the appropriate type via a finite set of recursively defined composition rules, and to a (modality-dependent) pronunciation. The fact that a particular pronunciation may be consistent with more than one meaning is a problem for a truth conditional view of meaning only if the mapping principles necessarily relate that pronunciation to a single syntactic object, which must then (somehow) be mapped onto distinct sets of truth conditions, giving us the situation in (16). If, on the other hand, the mapping principles allow for the possibility of relating the pronunciation to distinct syntactic objects, we end up with S in (16a) and S' in (16b), and the problem disappears.

Parsons acknowledges this in her discussion of lexical ambiguity when she says that "it may be that 'bank' (financial institution) and 'bank' (wall of a river channel) can be distinguished on the basis of a good syntax", and Saka (2007) does the same, but both are skeptical that the full range of ambiguity phenomena can be handled in this way. In order to make the skeptical case, however, one would need to address actual proposals about these relations within linguistic theory, and show

that none provide a coherent basis for handling the challenge of uncertainty. While it is not possible to address all plausible linguistic analyses of these phenomena, a quick look at a few reasonably well-established approaches to them suggests that current linguistic theory can take us fairly far in meeting the challenge presented by ambiguity for truth conditional theories of meaning.

Let's begin with lexical and structural ambiguity. The latter is straightforward: the fact that syntactic representations (or well-formedness derivations/proofs, if we are working in a theory that eschews levels of representation) have hierarchical structure but phonological representations have only linear structure ensures that two structurally distinct representations may have the same pronunciation. This is the case in (17a-b), which are both pronounced /a b c/, assuming that syntactic precedence relations determine linear order of pronunciation.

(17) a. 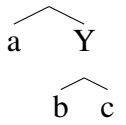

b. 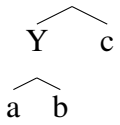

If we further assume a (conservative) set of composition rules whereby the denotation of any constituent α is a function of the denotations of its immediate subconstituents (its daughters), we end up with structural ambiguity: (17a) and (17b) can have different interpretations but the same pronunciations.

In fact, lexical ambiguity (homonymy) is handled in essentially the same way — in terms of representational properties that are obscured in the mapping to phonol-

ogy — a point made by Gillon (1990). This fact is often obscured by a convenient notational shortcut, however: the use of orthographic units (written words) to represent terminal nodes in a syntactic representation. Because of this convention, a representation like (18) gives the impression that the first occurrence of *bank* and the second occurrence are the same objects.

(18)

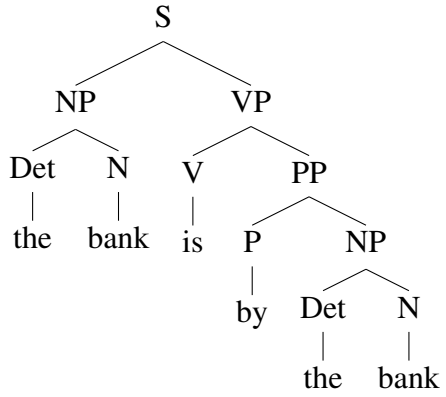

But this impression is incorrect. In fact, in most theories of syntax all nodes in a syntactic representation have the same basic formal properties: they are structured bundles of features. In particular, a syntactic object σ is at least a triple of the form $\langle P, S, D \rangle$ (depending on the theory, they may have more features), where P is a set of phonological features (the pronunciation of σ), S is a set of morphosyntactic features (category, case, number, etc.), and D is a set of semantic features: whatever is used to characterize denotations in the context of a broader theory of compositional interpretation. The exact properties of these features vary across framework and assumptions, but the overall architecture is the same. Crucially, nothing forbids a language from containing objects — either simple or complex — that have

the same phonological features but distinct semantic (or syntactic) features. Such objects sound the same, but are formally distinct, and their use in syntactic structures that are otherwise identical entails that those structures are distinct syntactic objects, which may in turn be mapped onto distinct meanings. Some researchers, such as Gillon (1990), attempt to capture this fact through the use of indices (*bank*$_1$ vs. *bank*$_2$), though the indices themselves have no theoretical significance. They are merely notational devices that let us distinguish an object of the form $\langle P, S, D\rangle$ from one of the form $\langle P, S, D'\rangle$ without having to write much more than what is specified by P.

Polysemy is a harder problem. (See article 21 *Sense Relations*, for additional discussion.) The most bare-bones way of handling it is to bite the bullet and assume that polysemy has the same representational status as homonymy: polysemous terms involve syntactic objects that have identical phonological features but distinct semantic features, and so are formally distinct. This view is completely consistent with a broader theory of lexical organization that explains why e.g. (19a-b) are related in a way that (20a-b) are not.

(19) a. $\langle$/run/, V, *manner of locomotion*$\rangle$

b. $\langle$/run/, V, *compete for elected office*$\rangle$

(20) a. $\langle$/bank/, N, *financial institution*$\rangle$

b. $\langle$/bank/, N, *wall of a river channel*$\rangle$

However, there is a wide body of work in lexical semantics that has attempted to more directly capture the differences between polysemy and homonymy, typically in terms of semantic underspecification (see e.g. Pinkal 1999; this strategy can also be applied to representations larger than words, as explained in article 24 *Semantic Underspecification*). On this view, a polysemous term may have disjunctive or incomplete semantic features, but in the course of putting together a syntactic representation a particular option must be selected, possibly via interactions with the semantic features of other terms. For example, Pustejovsky (1995) shows how the various senses of the verb *enjoy* in (21) are typically determined by semantic properties of its nominal argument, and develops a generative grammar of lexical semantic feature composition to account for these patterns.

(21) a. enjoy a martini (*enjoy drinking it*)

b. enjoy a cigarette (*enjoy smoking it*)

c. enjoy a recording of Britten's opera *Billy Budd* (*enjoy listening to it*)

d. enjoy a debate (*enjoy participating in it* or *enjoy debating*)

Whether Pustejovsky's specific analysis is ultimately the best way of accounting for these facts remains to be determined (it could be the case that they involve coercion of the sort discussed below; see Pylkkänen & McElree 2006 for recent work on

this); what is important for the current discussion is that it is the type of analysis that needs to be addressed in order to make the case that polysemy is a problem for truth conditional semantics, since it provides exactly the sort of representational basis for distinguishing sentences that involve distinct senses of polysemous terms, as well as empirical arguments for the representational status of that distinction.

Saka (2007) extends Parsons' argument by pointing to various sorts of "coercion" phenomena, such as deferred reference, metonymy, metaphor, type/token alternations, and so forth, which are not so clearly amenable to the type of representational analysis posited above for ambiguity and polysemy (see article 25 *Coercion* and article 26 *Metaphor and Metonymy*). For example, there is no reason to assume that the syntactic object pronounced *ham sandwich* in (22a) is a different lexical item (or is composed of different lexical items) from the one in (22b).

(22) a. This is the ham sandwich. (waitress holding up a ham sandwich)

b. I am the ham sandwich. (raising my hand and beckoning to the waitress)

In fact, as both Nunberg (1995) and Ward (2004) argue, there is good reason to believe that the basic meaning of *ham sandwich* is preserved, since it can be straightforwardly targeted by discourse anaphora, as in (23).

(23) The ham sandwich seems to be enjoying *it*. (*it* = the ham sandwich)

This is particularly striking, considering that access to part of a meaning (which we might take deferred reference to involve, since *the ham sandwich* in (23) is actually being used to pick out an individual distinct from the sandwich) is generally bad. (24), for example, does not readily permit an interpretation in which *it* refers only to the ham, not to the whole sandwich (Postal 1969; Ward, Sproat & McKoon 1991).

(24) The ham sandwich didn't go down well because *it* was two years old.

However, even these facts can be accommodated by assuming with Ward and Nunberg that the relevant readings are derived by mapping the denotation of *ham sandwich* (or the verb, in Ward's case) to a new one of the right sort; formally, we can implement this by positing a type-shifting rule that maps properties to properties. The broader motivation for type-shifting rules as part of the compositional interpretation system is very well established (e.g., to account for the fact that the conjunction *and* can combine with categories of any type without having to posit a large set of lexical entries; see Partee & Rooth 1983 and article 96 *Type Shifting*). What is relevant for the current discussion is that they map meanings into new meanings, providing just the representational distinction we need to ensure that deferred reference can be handled by a truth conditional semantics. For example, if (25a) is the basic representation of *ham sandwich*, (25b) is the type-shifted representation, where f is a context-sensitive function mapping properties into proper-

ties.

(25) a. $\langle$ham sandwich, S, *ham sandwich*$\rangle$

b. $\langle$ham sandwich, S, f(*ham sandwich*)$\rangle$

This analysis explains why discourse anaphora is possible — the core meaning is still part of the representation — and also has no problem with examples like (26), where the shifted version of the nominal provides a "regular" restriction for the quantifier *every* in the usual way.

(26) a. Every ham sandwich is enjoying his meal.

b. for all x such that x is a f(ham sandwich), x is enjoying x's meal

f a contextually salient function from ham sandwiches to individuals (the *eater-of* function)

Referential ambiguity does not pose a particular challenge, since it can be handled straightforwardly either by assuming a semantics with variables and assignment functions and a syntax in which syntactic features distinguish one variable from another (here indices do have theoretical significance; Heim & Kratzer (1998) provides a good overview of such a system), or by assuming a semantics without variables and letting sentences that contain anaphoric terms denote incomplete propositions (as in Jacobson 1999, 2000). On this latter view, sentences with anaphoric terms aren't actually assigned truth conditions, and two uses of e.g. *He*

is from Chicago that are "about" different individuals will in fact have identical meanings: they denote a function from individuals to truth values that is true of an individual if it is from Chicago. So this sort of analysis bypasses the problem of interpretive uncertainty completely by denying that fixing the reference of pronouns is part of semantics.

In each of the cases discussed above, the key to responding to Parsons' challenge was to demonstrate that standard (or at least reasonable) assumptions about lexico-syntactic representation and composition support the view that observed variability in truth-conditions has a representational basis: in each case, the mappings from representations to meanings are one-to-one, but the mappings from representations (and meanings) to pronunciations are sometimes many-to-one. But what if standard assumptions fail to support such a view? In such a case, Parsons' challenge would reemerge, unless a representational distinction can be demonstrated.

One of the strongest cases of this sort comes from the work of Charles Travis, who discusses a particular form of truth conditional variability associated with color terms (see e.g., Travis 1985, 1994, 1997). The following passage illustrates the phenomenon:

> A story. Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist

> friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood. (Travis 1997, p. 89)

This scenario appears to show that distinct utterances of the words in (27), said in order to describe the same scenario (the relation between the leaves and a particular color), can be associated with distinct truth values.

(27) The leaves are green.

Following the line of reasoning initially advanced by Parsons, Travis concludes from this example that sentence meaning is not truth conditional; that instead, the semantic value of a sentence at most imposes some necessary conditions under which it may be true (as well as conditions under which it may be used), but those conditions need not be sufficient, and the content of the sentence does not define a function from contexts to truth.

Travis' skeptical conclusion is challenged by Kennedy & McNally (to appear), however, who ask us to consider a modified version of the story of Pia and her leaves. Now she has a pile of painted leaves of varying shades of green (pile A) as well as a pile of naturally green leaves, also of varying shades (pile B). Pia's artist friend walks in and asks if she can have some green leaves for a project. Pia invites her to sort through the piles and take whichever leaves she wants. In sorting

through the piles, the artist might utter any of the sentences in (28) in reference either to leaves from pile A or to leaves from pile B, as appropriate based on the way that they manifest green: the particular combination of hue, saturation, and brightness, extent of color, and so forth.

(28) a. These leaves are green.

b. These leaves are greener than those.

c. These leaves aren't as green as those.

d. These leaves are less green than those.

e. These leaves are not green enough.

f. These leaves are too green.

g. These leaves are completely green.

h. These leaves are perfectly green.

i. These leaves are pretty/really green.

j. These leaves are not so green.

What is important to observe is that for the artist, who is interested in the colors of the leaves in her composition, any of these sentences would in principle be felicitous. Furthermore, (28a) is true of all of the leaves — both the painted ones and the natural ones — provided they are "green enough". The only issue is how green they are, or maybe how much of each of them is green; why they are green (i.e. because

they are naturally or artificially so) is irrelevant.

The situation is different for the botanist. She is perfectly justified in continuing to reject (the words in) (28a) as a false description of the painted leaves, while accepting it as true of the natural leaves. However, if these are her judgments about (28a), then none of the examples in (28b-i) are acceptable as descriptions of any of the leaves. That is, she cannot point to pile B (the naturally green leaves) and utter (28a) with the intended meaning (that the leaves are naturally green), and then strengthen or reiterate her point by pointing to pile A and uttering (28e) or (28j). Similarly, there is no way for her to use (28b) to justify her selection of the naturally green leaves over the painted ones, or (28c-d) to justify rejection of the latter, strictly on the basis of their biological properties. In short, once she starts using sentences that involve some notion of degree or comparison, the painted/natural distinction is out of the picture; all that is relevant is the relative degree of color.

What these facts show is that there is a semantic difference between occurrences of *green* that are used to distinguish between objects on the basis of *why* they are green (e.g., chlorophyll vs. paint) and instances that are used to distinguish between objects on the basis of *how* they are green (depth of hue, proximity to a prototype, extent of color, etc.). Each of (28b-j) involves the combination of the color adjective with a different element from the set of English degree morphemes, all of which require the adjective they combine with to be gradable. The fact that (28b-j)

are acceptable when (28a) is true of both sorts of leaves shows that on this use, it is gradable; the fact that (28b-j) are unacceptable when (28a) is true only of the naturally green leaves (in a context in which both piles contain objects with the same range of objective color features) shows that, on this use, it is nongradable.

The gradable/nongradable distinction is a matter of meaning, typically cashed out as a distinction of semantic type (see Kennedy 1999 for discussion, and article 57 *Adjectives*). It follows, then, that the two utterances of *green* in Travis' story about Pia and her painted leaves involve utterances of distinct terms with distinct meanings, and therefore the sentences in which they are uttered are distinct sentences which may have distinct conditions for truth. This example therefore poses no more of a challenge for truth conditional semantics than other cases of lexical ambiguity. It does, however, highlight the importance of a detailed and comprehensive linguistic analysis, since it shows that some cases of lexical ambiguity are revealed only through a close examination of the distribution and interpretation of the terms of interest in a variety of syntactic and morphological contexts.

3. Vagueness

3.3 The challenge of vagueness

It is generally accepted that the locus of vagueness in sentences like (29) is the predicate headed by the gradable adjective *expensive*: this sentence is vague because, intuitively, what it means to count as expensive is unclear.

(29) The coffee in Rome is expensive.

Sentences like (29) have three distinguishing characteristics, which have been the focus of much work on vagueness in semantics and the philosophy of language. The first is contextual variability in truth conditions: (29) could be judged true if asserted as part of a conversation about the cost of living in Rome vs. Naples (*In Rome, even the coffee is expensive!*), for example, but false in a discussion of the cost of living in Chicago vs. Rome (*The rents are high in Rome, but at least the coffee is not expensive!*). This kind of variability is of course not restricted to vague predicates — for example, the relational noun *citizen* introduces variability because it has an implicit argument (*citizen of x*) but it is not vague — though all vague predicates appear to display it.

The second feature of vagueness is the existence of borderline cases. For any context, in addition to the sets of objects that a predicate like *is expensive* is clearly

true of and clearly false of, there is typically a third set of objects for which it is difficult or impossible to make these judgments. Just as it is easy to imagine contexts in which (29) is clearly true and contexts in which it is clearly false, it is also easy to imagine a context in which such a decision cannot be so easily made. Consider, for example, a visit to a coffee shop to buy a pound of coffee. The Mud Blend at $1.50/pound is clearly not expensive, and the Organic Kona at $20/pound is clearly expensive, but what about the Swell Start Blend at $9.25/pound? A natural response is "I'm not sure"; this is the essence of being a borderline case.

Finally, vague predicates give rise to the Sorites Paradox, illustrated in (30).

(30) *The Sorites Paradox*

P1. A $5 cup of coffee is expensive (for a cup of coffee).

P2. Any cup of coffee that costs 1 cent less than an expensive one is expensive (for a cup of coffee).

C. Therefore, any free cup of coffee is expensive.

The structure of the argument appears to be valid, and the premises appear to be true, but the conclusion is without a doubt false. Evidently, the problem lies somewhere in the inductive second premise; what is hard is figuring out exactly what goes wrong. And even if we do, we also need to explain both why it is so hard to detect the flaw and why we are so willing to accept it as valid in the first place.

These points are made forcefully by Fara (2000), who succinctly characterizes the challenges faced by any explanatorily adequate account of vagueness in the form of the following three questions:

(31) a. *The Semantic Question*

If the inductive premise of a Sorites argument is false, then is its classical negation — the SHARP BOUNDARIES CLAIM that there is an adjacent pair in a sorites sequence such that one has the property named by the vague predicate and the other doesn't — true?

(i) If yes, how is this compatible with borderline cases?

(ii) If no, what revision of classical logic and semantics must be made to accommodate this fact?

b. *The Epistemological Question*

If the inductive premise is false, why are we unable to say which of its instances fail, even in the presence of (what we think is) complete knowledge of the facts relevant to judgments about the predicate?

c. *The Psychological Question*

If the inductive premise is false, why are we so inclined to accept it in the first place? What makes vague predicates tolerant in the relevant way? Why do they seem "boundaryless"?

These questions provide a set of evaluation criteria for theories of vagueness: one theory can be preferred over another to the extent that it provides satisfactory answers to these questions. Of particular importance is answering the Epistemological and Psychological questions: it is fairly straightforward to construct a theory that answers the Semantic Question, but many such theories fail to say anything about the other two, and so fail as explanatory theories of vagueness.

In particular, this is the case with most linguistic analyses of the class of vague predicates most commonly discussed by semanticists: gradable adjectives like *expensive*. A fruitful and rich line of research, primarily on comparatives, superlatives and other complex expressions of quantity and degree, analyzes the meaning of gradable adjectives as relations between objects and degrees (see e.g., Seuren 1973; Cresswell 1977; von Stechow 1984; Heim 2000; Bierwisch 1989; Schwarzschild & Wilkinson 2002; Kennedy 1999, 2001; Kennedy & McNally 2005; Rotstein & Winter 2004; and see article 58 *Comparatives and Superlatives*). The adjective *expensive*, on this view, denotes the relation in (32), which is true of an object x and a degree (of cost) d just in case the cost of x is at least as great as d.

(32) $\quad [\![expensive]\!] = \lambda d \lambda x.\textsc{cost}(x) \succeq d$

When it comes to analyzing the positive (unmarked) form of a gradable predicate, which is what we see in examples like (29) and on which a Sorites argument

606 is based (i.e., the vague form), the usual strategy is to hypothesize that the degree
607 argument is saturated by a contextually determined STANDARD OF COMPARISON,
608 which represents the "cut off point" between the positive and negative extensions
609 of the predicate, possibly relativized to a COMPARISON CLASS of objects deemed
610 somehow similar to the target of predication. (For discussion of standards of com-
611 parison and comparison classes, see Wheeler 1972; Rips & Turnbull 1980; Klein
612 1980; Ludlow 1989; Bierwisch 1989; Kamp & Partee 1995; Fara 2000; Kennedy
613 2007.) The standard of comparison is usually treated as a free variable over degrees
614 whose value is determined by a special assignment function (see Barker 2002 for an
615 explicit statement of this idea), though it is sometimes linked to a particular value,
616 such as the average degree to which the objects in the comparison class manifest
617 the relevant gradable property (as in e.g., Bartsch & Vennemann 1972).

618 This type of approach clearly provides an explanation for the truth conditional
619 variability of vague predicates. (29) is true just in case the cost of the coffee in
620 Rome exceeds the value of the standard of comparison, whatever that is, and false
621 if it is exceeded by the standard. Since different contexts of utterance will invoke
622 different standards (e.g., one based on the price of coffee in Italian cities vs. one
623 based on the price of coffee in Rome and Chicago), the truth of (29) may shift. For
624 the very same reason, this approach provides a partial answer to Fara's Semantic
625 Question: characterizing the meaning of (the positive form of) *expensive* in terms

of a relation between two degrees amounts to accepting the sharp boundaries claim, since the truth or falsity of a sentence like (29) is simply a function of the relation between these two degrees.

This looks like a good result at first: the Sorites Paradox disappears, because the second premise is guaranteed to be false. However, we have no obvious account for our judgments about borderline cases, and certainly no explanation for why we might have thought the second premise to be true. That is, we have no answers to the borderline case subpart of the Semantic Question, nor do we have answers to the Epistemological or Psychological Questions. We might appeal to some sort of indeterminacy in, or incomplete knowledge of, the assignment function involved in fixing standards of comparison in order to gain some traction on the Epistemological Question and the status of borderline cases, but this move will not help us with the Psychological Question. If knowing the meaning of a vague predicate means knowing that it requires its argument to have a degree of a scalar property whose value exceeds a standard that gets *fixed by the context*, then all other things being equal, we ought to be willing to reject the inductive premise of the Sorites Paradox. We should know that at some point along the line this relation must fail to hold, even if we don't know exactly where it is.

This is not to say that something like the traditional linguistic analysis of gradable predicates couldn't be augmented or supplemented with some other principles

that would allow for an answer to all of Fara's questions. Such principles could be semantic, but they could also be pragmatic or even cognitive; the analyses I will discuss in more detail below are differentiated roughly along these lines. The importance of taking this extra step must be emphasized, however. Semantic theories (such as the approach to gradable predicates outlined above) are typically designed in such a way that lexical and compositional meaning together result in expressions that support clear judgments of truth or falsity, possibly in a context dependent way, given a certain set of facts. The Epistemological and Psychological Questions highlight the fact that even when a set of crucial facts is known — the actual distribution of costs of coffee in various cities, for example, or even just the knowledge that there is a distribution of costs — judgments of truth and falsity can remain unclear (with borderline cases) or can even be wrong (the inductive premise of the Sorites, if the sharp boundaries claim is in fact correct). But this then calls into question the initial step of characterizing meanings in terms of truth functions: if we want to maintain this aspect of semantic theory, then we need to have answers to all three questions about vagueness.

Before moving to a discussion of particular approaches to vagueness, I also want to point out that vagueness is by no means restricted to gradable adjectives, even though the majority of examples discussed in both the linguistic and philosophical literature involve expressions from this class. Using the three characteristics of

truth conditional variability, borderline cases, and the Sorites Paradox as a guide, we can find vague terms in all grammatical classes: nouns (like *heap*, which gives the Sorites Paradox its name), verbs (such as *like*, or more significantly, *know*), determiners (such as *many* and *few*), prepositions (such as *near*) and even locative adverbials, as in the Fry and Laurie dialogue in (33).

(33) FRY: There are six million people out there....

LAURIE: Really? What do they want?

Here the humor of Laurie's response comes from the vagueness of *out there*: whether it extends to cover a broad region beyond the location of utterance (Fry's intention) or whether it picks out a more local region (Laurie's understanding). Vagueness is thus pervasive, and its implications for the analysis of linguistic meaning extend to all parts of the lexicon.

3.3 Approaches to vagueness

It is impossible to summarize all analysis of vagueness in the literature, so I will focus here on an overview of four major approaches, based on supervaluations, epistemic uncertainty, contextualism, and interest relativity. For a larger survey of approaches, see Williamson (1994); Keefe & Smith (1997); and Fara & Williamson (2002).

3.3.1 Supervaluations

Let's return to one of the fundamental properties of vague predicates: the existence of borderline cases. (34a) is clearly true; (34b) is clearly false; (34c) is (at least potentially) borderline.

(34) a. Mercury is close to the sun.

b. Pluto is close to the sun.

c. The earth is close to the sun.

However, even if we are uncertain about the truth of (34c), we seem to have clear intuitions about (35a-b): the first is a logical truth, and the second is a contradiction.

(35) a. The earth is or isn't close to the sun.

b. The earth is and isn't close to the sun.

This is not particularly surprising, as these sentences are instances of (36a-b):

(36) a. $p \vee \neg p$

b. $p \wedge \neg p$

As noted by Fine (1975), these judgments show that logical relations (such as the Law of the Excluded Middle in (36a)) can hold between sentences which do not themselves have clear truth values in a particular context of utterance. Fine

accepts the position that sentences involving borderline cases, such as (34c), can fail to have truth values in particular contexts (this is what it means to be a borderline case), and accounts for judgments like (35a) by basing truth valuations for propositions built out of logical connectives not on facts about specific contexts of evaluation, but rather on a space of interpretations in which all "gaps" in truth values have been filled. In other words, the truth of examples like (35a-b) is based on SUPERVALUATIONS (van Fraassen 1968, 1969) rather than simple valuations.

The crucial components of Fine's theory are stated in (37)-(39); a similar set of proposals (and a more comprehensive linguistic analysis) can be found in Kamp (1975) and especially Klein (1980).

(37) *Specification space*

A partially ordered set of points corresponding to different ways of specifying the predicates in the language; at each point, every proposition is assigned true, false or nothing according to an "intuitive" valuation. This valuation must obey certain crucial constraints, such as Fine's *Penumbral Connections* which ensure e.g., that if x is taller than y, it can never be the case that *x is tall* is false while *y is tall* is true (cf. the *Consistency Postulate* of Klein 1980).

(38) *Completability*

Any point can be extended to a point at which every proposition is assigned a truth value, subject to the following constraints:

a. FIDELITY: Truth values at complete points are 1 or 0.

b. STABILITY: Definite truth values are preserved under extension.

(39) *Supertruth*

A proposition is supertrue (or superfalse) at a partial specification iff it is true (false) at all complete extensions.

According to this approach, the reason that we have clear intuitions about (35a) and (35b) is because for any ways of making things more precise, we're always going to end up with a situation where the former holds for any proposition and the latter fails to hold, regardless of whether the proposition has a truth value at the beginning. In particular, given (38), (35a) is supertrue and (35b) is superfalse.

This theory provides an answer to Fara's Semantic Question about vagueness. According to this theory, any complete and admissible specification will entail a sharp boundary between the things that a vague predicate is true and false of. This renders inductive statements like the second premise of (40) superfalse, even when the argument as a whole is evaluated relative to a valuation that does not assign truth values to some propositions of the form *x is heavy*, i.e., one that allows for

borderline cases.

(40) a. A 100 kilogram stone is heavy.

b. Any stone that weighs 1 gram less than a heavy one is heavy.

c. #A 1 gram stone is heavy.

The supervaluationist account thus gives up bivalence (the position that all propositions are either true or false relative to particular assignments of semantic values), but still manages to retain important generalizations from classical logic (such as the Law of the Excluded Middle) and assign a definitive value of FALSE to the inductive premise of the Sorites through the concept of supertruth.

However, although supervaluation accounts address the semantic question, on their own they have little to say about the Epistemological and Psychological Questions, as pointed out by Fara (2000). Fine (1975) attempts to answer the former by arguing that the extension boundaries for vague predicates are both arbitrary and infinitely variable. We are unable to identify a cutoff point, according to Fine, because it could in principle be in an infinite number of different places, if we allow an infinite domain (though it must always respect "admissibility"). There need be no determinate fact about where it is; and in particular, there need be no *linguistic* fact (one rooted in our knowledge of meaning) about where it is.

A number of objections can be raised to Fine's response to the Epistemological Question. (For example, it is is not clear that the boundaries for a vague predicate are entirely arbitrary: intuitively, an object counts as *expensive* or *heavy* only if it has an appropriately "high" degree of the relevant property.) But even if it is accepted, the Psychological Question still remains unanswered. If knowing the meaning of the universal statement in (40b) means knowing that it invokes supertruth, and if knowing the meaning of a vague predicate means knowing how it could be made precise (as claimed in Fine 1975, p. 277), then it is unclear why we are unwilling to assign a judgment of false when we are confronted with such statements.

Finally, supervaluationist accounts have been criticized for not even providing a satisfactory answer to the Semantic Question. The problem is that even though a supervaluation analysis predicts the existence of borderline cases by allowing for incomplete models, in any particular incomplete model, the boundary between the things that a vague predicate is definitely true of and those things for which it is indeterminate is crisp. But our judgments about "borderline borderlines" are no more clear than our judgments about "central borderlines", suggesting that the boundaries aren't so crisp after all. If we now need to invoke some other mechanism to explain such cases of HIGHER ORDER VAGUENESS, then we can legitimately ask whether supervaluations provide the right starting point for the core cases.

3.3.2 Epistemic uncertainty

The epistemic analysis of vagueness, developed most extensively in the work of Timothy Williamson (1992, 1994, 1997), starts from the assumption that vague predicates (and in fact all predicates in language) sharply define a positive and negative extension: there are no extension gaps, and there is no denial of bivalence, as in supervaluation accounts. Vagueness arises because the exact boundaries of these sets are not known; in fact, they are *unknowable*. Vagueness thus reflects an underlying ignorance about a fundamental feature of meaning: the precise factors that determine the extension of a predicate.

It should be clear that this approach provides a straightforward answer to the first part of the Semantic Question: it begins from an assumption of sharp boundaries, so the second premise of the Sorites is false. In order to see how it handles borderline cases and the Epistemological and Psychological Questions, we need to take a closer look at its answer to the core question of why we are ignorant about extension boundaries. Why can't we figure out what the sharp boundaries of a vague predicate are, and in so doing eliminate borderline cases and identify where the second premise of the Sorites Paradox fails?

Williamson's response comes in several parts. First, he assumes that meaning supervenes on use; as such, a difference in meaning entails a difference in use,

but not vice-versa. Second, he points out that the meanings of some terms may be stabilized by natural divisions (cf. Putnam's 1975 distinction between H_2O and XYZ), while the meanings of others (the vague ones) cannot be so stabilized: a slight shift in our disposition to say that the earth is close to the sun would slightly shift the meaning of *close to the sun*. The boundary is sharp, but not fixed. But this in turn means that an object around the borderline of a vague predicate P could easily have been (or not been) P had the facts (in particular, the linguistic facts) been slightly different — different in ways that are too complex for us to even fully catalogue, let alone compute. Given this instability, we can never really know about a borderline case whether it is or is not P.

This last point leads to the principle in (41), which is another way of saying that vague knowledge requires a margin for error.

(41) *The Margin for Error Principle*

For a given way of measuring differences in measurements relevant to the application of property P, there will be a small but non-zero constant c such that if x and y differ in those measurements by less than c and x is known to be P, then y is known to be P.

The upshot of this reasoning is that it is impossible to know whether *x is P* is true or false when x and y differ by less than c. That's why we fail to reject the second

premise of the Sorites, and also why "big" changes make a difference. (If we replace *1 cent* with *1 dollar* in (30), or *1 gram* with *10 kilograms* in (40), the paradox disappears.)

There are a number of challenges to this account, most of which focus on the central hypothesis that we can be ignorant about core aspects of meaning and still somehow manage to have knowledge of meaning at all. Williamson (1992) lists these challenges and provides responses to them; here I focus on the question of whether this theory is adequate as an account of vagueness. According to Fara (2000), it is not, because although it addresses the Semantic and Epistemological Questions, it does not address the psychological one. In particular, there is no account of why we don't have the following reaction to the inductive premise: "That's false! I don't know where the shift from P to $\neg P$ is, so there are cases that I'm not willing to make a decision about, but I know it's in there somewhere, so the premise must be false."

Williamson (1997) suggests the answer has to do with the relation between imagination and experience. The argument runs as follows:

(42) i. It is impossible to gain information through imagination that cannot be gained through experience.

ii. It is impossible to recognize the experience of the boundary transition in a sorites sequence because the transition lacks a distinctive appear-

ance.

iii. Therefore, it is impossible to imagine the transition.

This failure of imagination then makes it impossible to reject the inductive premise, since doing so precisely requires imagination of the crucial boundary transition. However, according to Fara, this response doesn't help us with the trickier question of why we *believe* of every pair in a sorites sequence that the boundary is NOT there. In order to answer the psychological question, she says, we need an account that is more directly *psychological*.

We will examine two such accounts in the next section, but before moving to this discussion, I want to point out a more purely empirical problem for the epistemic analysis of vagueness, which comes from the phenomenon of CRISP JUDGMENTS, discussed in Kennedy (2007). For an illustration of the phenomenon, consider a context in which we are deciding who should review various papers for a semantics journal. Our two reviewers are Professors Jones and Smith. We are considering pairs of papers, which are similar in content but distinguished by their length, as described in (43).

(43) SCENARIO A: a 15-page paper and a 25-page paper

SCENARIO B: a 25-page paper and a 40-page paper

SCENARIO C: a 24-page paper and a 25-page paper

In scenarios A and B, we could felicitously use (44) to issue instructions about which reviewer should get which paper.

(44) Let Jones review the long paper and let Smith review the short one.

In each scenario, *the long paper* refers to the longer of the two papers and *the short paper* refers to the shorter of the two. Focusing on the former case (the latter is the same), the existence and uniqueness presuppositions of the definite description require that there be one and only one object in each scenario that satisfies the predicate *long* (since both satisfy *paper*); this means that a length of 25 pages counts as long in scenario A but does not count as long in scenario B. That this is so is not surprising given what we already know about the context-dependence of standards of comparison: in this kind of example, the presuppositions of the definite determiner cause us to accommodate a standard that makes *long* uniquely true of one of a pair of objects of different lengths.

What is surprising is that (44) cannot be felicitously used in scenario C, where the length difference between the two papers is small; here only a variant using the comparative form of the adjective (*longer*) is acceptable (the comparative form is also acceptable in scenarios A and B, of course):

(45) Let Jones review the longer paper and let Smith review the shorter one.

The contrast between (44) and (45) in scenario C is important because it shows that even under pressure from the presuppositions of the definite determiner, we cannot accommodate a standard of comparison for *long* that makes it true of a 25-page paper and false of a 24-page paper: we cannot use the positive form of the adjective to make what Kennedy (2007) calls "crisp judgments" to distinguish between a pair of objects that differ in length by only a small degree. This kind of judgment is precisely analogous to the kind of judgment that would be involved in rejecting the inductive premise of the Sorites, but there is a crucial difference: in this case, we know exactly where the cutoff point for *long* would have to be, namely somewhere between 24 and 25 pages in length. The epistemic account of vagueness provides no account of this fact (nor does an unaugmented supervaluationist account, since it too needs to allow for contextual shifting of standards of comparison). If the impossibility of crisp judgments in examples like these involving definite descriptions and our judgments about the second premise of the Sorites Paradox are instances of the same basic phenomenon, then the failure of the epistemic account of vagueness to explain the former raises questions about its applicability to the latter.

3.3.3 Contextualism and interest relativity

Raffman (1996) observes three facts about vague predicates and Sorites sequences. First, as we have already discussed, vague predicates have context dependent exten-

sions. Second, when presented with a sorites sequence based on a vague predicate P, a competent speaker will at some point stop (or start, depending on direction) judging P to be true of objects in the sequence. Third, even if we fix the (external) context, the shift can vary from speaker to speaker and from run to run. This is also a part of competence with P.

These observations lead Raffman (1994, 1996) to a different perspective on the problem of vagueness: reconciling tolerance (insensitivity to marginal changes) with categorization (the difference between being red and orange, tall and not tall, etc.). She frames the question in the following way: how can we simultaneously explain the fact that a competent speaker seems to be able to apply incompatible predicates (e.g., *red* vs. *orange*; *tall* vs. *not tall*) to marginally different (adjacent) items in the sequence and the fact that people are unwilling to reject the inductive premise of the Paradox?

Her answer involves recognizing the fact that evaluation of a sorites sequence trigger a context shift, which in turn triggers a shift in the extension of the predicate in such a way as to ensure that incompatible predicates are never applied to adjacent pairs, and to make (what looks like) a sequence of inductive premises all true. (For similar approaches, see Kamp 1981; Bosch 1983; Soames 1999.) This gives the illusion of validity, but since there is an extension shift, the predicate at the end of the series is not the same as the one at the beginning, so the argument is invalid.

There are three pieces to her account. The first comes from work in cognitive psychology, which distinguishes between two kinds of judgments involved in a sorites sequence. The first is categorization, which involves judgments of similarity to a prototype/standard; the second is discrimination, which involves judgments of sameness/difference between pairs. Singular judgments about items involve categorization, and it is relative to such judgments that a cutoff point is established. Discrimination, on the other hand, doesn't care where cutoff points fall, but it imposes a different kind of constraint: adjacent pairs must be categorized in the same way (Tversky & Kahneman 1974).

At first glance, it appears that the categorization/discrimination distinction just restates the problem of vagueness in different terms: if for any run of a sorites sequence, a competent speaker will at some point make a category shift, how do we reconcile such shifts with the fact that we resist discrimination between adjacent pairs? Note that the problem is not the fact that a speaker might eventually say of an object o_i in a sorites sequence based on P that it is not P, even if she judged o_{i+1} to be P, because this is a singular judgment about o_i. The problem is that given the pair $\langle o_i, o_{i+1} \rangle$, the speaker will refuse to treat them differently. This is what underlies judgments about the inductive premise of the Sorites Paradox and possibly the crisp judgment effects discussed above as well, though this is less clear (see below).

The second part of Raffman's proposal is designed to address this problem, by positing that a category shift necessarily involves a change in perspective such that the new category instantaneously absorbs the preceding objects in the sequence. Such BACKWARDS SPREAD is the result of entering a new psychological state, a Gestalt shift that triggers a move from one 'category anchor' or prototype to another, e.g., from the influence of the *red* anchor to the influence of the *orange* one. This gives rise to the apparent boundlessness of vague predicates: a shift in category triggers a shift in the border away from the edge, giving the impression that it never was there in the first place.

And in fact, as far as the semantics is concerned, when it comes to making judgments about pairs of objects, it never is. This is the third part of the analysis, which makes crucial appeal to the context dependence of vague predicates. Raffman proposes that the meaning of a vague predicate P is determined by two contextual factors. The EXTERNAL CONTEXT includes discourse factors that fix domain, comparison class, dimension, etc. of P. The INTERNAL CONTEXT includes the properties of an individual's psychological state that determine dispositions to make judgments of P relative to some external context. Crucially, a category shift causes a change in internal context e.g., (from a state in which the *red* anchor dominates to one in which the *orange* anchor does), which in turn results in a change in the extension of the predicate in the way described above, resulting in backwards

spread.

Taken together, these assumptions provide answers to each of Fara's questions. The answer to the semantic question is clearly positive, since the commitment to category shifts involves a commitment to the position that a vague predicate can be true of one member of an adjacent pair in a sorites sequence o_i and false of o_{i+1}. The reason we cannot say which $\langle o_i, o_{i+1} \rangle$ has this property, however, is that the act of judging o_i to be not P (or P) causes a shift in contextual meaning of P to P', which, given backwards spread, treats o_i and o_{i+1} the same. This answer to the epistemological question also underlies the answer to the psychological question: even though the inductive premise of the Sorites Paradox is false for any fixed meaning of a vague predicate, we think that it is true because it is possible to construct a sequence of true statements that look like (instantiations of) the inductive premise, but which in fact do not represent valid reasoning because they involve different contextual valuations of the vague predicate. For example, if we are considering a sequence of 100 color patches $\{p_1, p_2, ...p_{100}\}$ ranging from 'pure' red to 'pure' orange, such that a category shift occurs upon encountering patch p_{47}, the successive conditional statements in (46a) and (46c) work out to be true thanks to backwards spread (because their subconstituents are both true and both false in their contexts of evaluation, respectively), even though *red* means something different in each case.

(46) a. If p_{45} is red, then p_{46} is red. $p_{45}, p_{46} \in [\![red]\!]^c$

TRUE $\rightarrow$ TRUE $\models$ TRUE

b. SHIFT at p_{47}: change from context c to context c'

c. If p_{46} is red, then p_{47} is red. $p_{46}, p_{47} \notin [\![red]\!]^{c'}$

FALSE $\rightarrow$ FALSE $\models$ TRUE

A variant of the contextualist analysis is provided by Fara (2000). Like the contextualist, Fara assumes that there is a fixed point (a 'standard') in any context that distinguishes the objects that a vague predicate is true of from those which it is false of. And like the contextualist, Fara's analysis entails that adjacent elements in a sorites sequence are always treated in the same way, an effect that she describes in terms of the constraint in (47).

(47) *Similarity Constraint*

Whatever standard is in use for a vague expression, anything that is saliently similar, in the relevant respect, to something that meets the standard itself meets the standard; anything saliently similar to something that fails to meet the standard itself fails to meet the standard.

With (47) in hand, Fara provides answers to the epistemological and psychological questions that are also quite similar to those provided in a contextualist analysis. We are unable to pinpoint the boundary between objects that a vague predicate is true and false of because in evaluating the predicate for any adjacent pair of objects

in a sorites sequence, we raise the similarity of the pair relative to the property that generates the sequence to salience, thereby rendering it true (or false) of both of the objects are considering. Since this further entails that any instance of the universal premise of the Sorites Paradox (expressed out as a conditional statement of the sort we saw in (46)) is true, it is no surprise that we are unwilling to judge the universal premise false.

Where Fara's account crucially differs from the contextualist approach is in the way that the Similarity Constraint is derived. In a Raffman-style contextualist account, (47) is a consequence of backward spread, which reflects a change in the content of a vague predicate at the moment of category shift. In contrast, the content of a vague predicate remains constant in Fara's account, but its extension can shift in a way that derives the Similarity Constraint. Specifically, Fara argues that vague predicates denote INTEREST RELATIVE properties, of the following sort: for any vague scalar predicate P, an object falls in its the positive extension of P just in case it has a degree of the scalar concept that P encodes that is *significant* given our interests (see also Bogusławski 1975). Interest relativity allows for shifts in the extension of a vague predicate without a corresponding shift in its content: whether an object counts as red or not might change as the interests of the individual evaluating the predicate changes, but the denotation of the predicate is the fixed property of having a significant degree of redness.

This proposal derives the Similarity Constraint in the following way. Among our interests is a standing interest in efficiency, which has the consequence that whenever two objects are saliently similar with respect to a vague scalar predicate and they are being actively considered, the cost of discriminating between them typically outweighs the benefit. As a result, they count as 'the same for present purposes', and one will have a degree of the relevant property that is significant relative to an evaluator's interests if and only if the other does. This result is the key to understanding how Fara reconciles her 'sharp boundaries' answer to the semantic question with the apparent 'shiftiness' entailed by her answer to the epistemological and psychological. In any context, there is a pair of objects in a sorites sequence o_i and o_{i+1} such that the predicate on which the sequence is based is true of one and false of the other. However, any attempt to evaluate the predicate for this particular pair will render them saliently similar, which, given interest relativity, will cause the extension of the predicate to shift in a way that ensures that they are evaluated in the same way. In Fara's words: "the boundary between the possessors and the lackers in a sorites series is not sharp in the sense that we can never bring it into focus; any attempt to bring it into focus causes it to shift somewhere else." (Fara 2000, pp. 75-76)

One of the reasons that the contextualist and interest relative analyses provide compelling answers to the psychological question is that they are inherently psy-

chological: the former in the role that psychological state plays in fixing context sensitive denotations; the latter in the role played by interest relativity. Moreover, in providing an explanation for judgments about pairs of objects, these analyses can support an explanation of the 'crisp judgment' effects discussed in the previous section, provided they can be linked to a semantics that appropriately distinguishes the positive and comparative forms of a scalar predicate. However, the very aspects of these analyses that are central to their successes also raise fundamental problems that question their ultimate status as comprehensive accounts of vagueness, according to Stanley (2003).

Focusing first on the contextualist analysis, Stanley claims that it makes incorrect predictions about versions of the Sorites that involve sequential conditionals and ellipsis. Stanley takes the contextualist to be committed to a view in which vague predicates are a type of indexical expression. Indexicals, he observes, have the property of remaining invariant under ellipsis: (48b), for example, cannot be used to convey the information expressed by (48a).

(48) a. Kim voted for that$_A$ candidate because Lee voted for that$_B$ candidate.

b. Kim voted for that$_A$ candidate because Lee did ~~vote for that$_B$ candidate~~.

Given this, Stanley argues that if the contextualist account of vagueness entails that vague predicates are indexicals, then our judgments about sequences of conditionals

like (46) (keeping the context the same) should change when the predicates are elided. Specifically, since ellipsis requires indexical identity, it must be the case that the elided occurrences of *red* in (49) be assigned the same valuation as their antecedents, i.e. that $[\![red]\!]^c = [\![red]\!]^{c'}$.

(49) a. If p_{45} is red, then p_{46} is red too. $p_{45}, p_{46} \in [\![red]\!]^c$

TRUE $\rightarrow$ TRUE $\models$ TRUE

b. SHIFT at p_{47}: change from context c to context c'

c. If p_{46} is ~~red~~, then p_{47} is ~~red~~ too. $p_{46} \in [\![red]\!]^{c'}; p_{47} \notin [\![red]\!]^{c'}$

TRUE $\rightarrow$ FALSE $\models$ FALSE

But this is either in conflict with backwards spread, in which case ellipsis should be impossible, or it entails that (49b) should be judged false while (49a) is judged true. Neither of these predictions are borne out: the judgments about (49) are in all relevant respects identical to those about (46).

Raffman (2005) responds to this criticism by rejecting the view that the contextualist account necessarily treats vague predicates as indexicals, and suggests that the kind of 'shiftability' of vague predicates under ellipsis that is necessary to make the account work is analogous to what we see with comparison classes in examples like (50a), which has the meaning paraphrased in (50b) (Klein 1980).

(50) a. That elephant is large and that flea is too.

b. That elephant is large for an elephant and that flea is large for a flea.

This is probably not the best analogy, however: accounts of comparison class shift in examples like (50a) rely crucially on the presence of a binding relation between the subject and a component of the meaning of the predicate (see e.g., Ludlow 1989; Kennedy 2007), subsuming such cases under a general analysis of 'sloppy identity' in ellipsis. If a binding relation of this sort were at work in (49), the prediction would be that the predicate in the consequent of (49a) and the antecedent of (49b) should be valued in exactly the same way (since the subjects are the same), which, all else being equal, would result in exactly the problematic judgments about the truth and falsity of the two conditionals that Stanley discusses. In the absence of an alternative contextualist account of how vague predicates should behave in ellipsis, then, Stanley's objection remains in force.

This objection does not present a problem for Fara's analysis, which assumes that vague predicates have fixed denotations. However, the crucial hypothesis that these denotations are interest relative comes with its own problems, according to Stanley. In particular, he argues that this position leads to the implication that the meaning of a vague predicate is always relativized to some agent, namely the entity relative to whom significance is assessed. But this implication is inconsistent with the fact that we can have beliefs about the truth or falsity of a sentence like (51) without having beliefs about any agent relative to whom Mt. Everest's height is

supposed to be significant.

(51) Mt. Everest is tall.

Moreover, the truth of the proposition conveyed by an utterance of (51) by a particular individual can remain constant even in hypothetical worlds in which that individual doesn't exist, something that would seem to be impossible if the truth of proposition has something to do with the utterer's interests. Fara (2008) rejects Stanley's criticism on the grounds that it presumes that an "agent of interest" is a constituent of the proposition expressed by (51), something that is not the case given her particular assumptions about the compositional semantics of the vague scalar predicates she focuses on, which builds on the decompositional syntax and semantics of Kennedy (1999, 2007) (see also Bartsch & Vennemann 1972). However, to the extent that her analysis creates entailments about the existence of individuals with the relevant interests, it is not clear that Stanley's criticisms can be so easily set aside.

4. Conclusion

Ambiguity and vagueness are two forms of interpretive uncertainty, and as such, are often discussed in tandem. They are fundamentally different in their essential features, however, and in their significance for semantic theory and the philosophy

of language. Ambiguity is essentially a "mapping problem", and while there are significant analytical questions about how (and at what level) to best capture different varieties of ambiguity, the phenomenon per se does not represent a significant challenge to current conceptions of semantics. Vagueness, on the other hand, raises deeper questions about knowledge of meaning. The major approaches to vagueness that I have outlined here, all of which come from work in philosophy of language, provide different answers to these questions, but none is without its own set of challenges. Given this, as well as the fact that this phenomenon has seen relatively little in the way of close analysis by linguists, vagueness has the potential to be an important and rich domain of research for semantic theory.

References

Barker, Chris 2002. The Dynamics of Vagueness. *Linguistics and Philosophy* 25(1), 1–36.

Bartsch, Renate & Theo Vennemann 1972. The Grammar of Relative Adjectives and Comparison. *Linguistische Berichte* 20, 19–32.

Bierwisch, Manfred 1989. The Semantics of Gradation. In: Manfred Bierwisch & Ewald Lang (eds.) *Dimensional Adjectives*, Berlin: Springer-Verlag. 71–261.

1121 Bogusławski, Andrzej 1975. Measures are Measures: In Defence of the Diversity
1122 of Comparatives and Positives. *Linguistiche Berichte* 36, 1–9.

1123 Bolinger, Dwight 1967. Adjectives in English: Attribution and Predication. *Lingua*
1124 18, 1–34.

1125 Bosch, Peter 1983. "Vagueness" Is Context-Dependence. A Solution to the Sorities
1126 Paradox. In: Thomas Ballmer & Manfred Pinkal (eds.) *Approaching Vagueness*,
1127 Amsterdam: North Holland. 189–210.

1128 Cinque, Gugliemo 1993. On the Evidence for Partial N Movement in the Romance
1129 DP. In: *University of Venice Working Papers, Vol. 3, No. 2*. Venice, Italy: Centro
1130 Linguistico Interfaculta, 21–40.

1131 Cresswell, M. J. 1977. The Semantics of Degree. In: Barbara Partee (ed.) *Montague*
1132 *Grammar*, New York: Academic Press. 261–292.

1133 Fara, Delia Graff 2000. Shifting Sands: An Interest-Relative Theory of Vagueness.
1134 *Philosophical Topics* 20, 45–81.

1135 Fara, Delia Graff 2008. Profiling Interest Relativity. *Analysis* 68(4), 326–335.

1136 Fara, Delia Graff & Timothy Williamson (eds.) 2002. *Vagueness*. Burlington, VT:
1137 Ashgate.

1138 Fine, Kit 1975. Vagueness, Truth and Logic. *Synthese* 30, 265–300.

van Fraassen, Bas C. 1968. Presupposition, Implication and Self-Reference. *Journal of Philosophy* 65, 136–152.

van Fraassen, Bas C. 1969. Presuppositions, Supervaluations and Free Logic. In: K. Lambert (ed.) *The Logical Way of Doing Things*, New Haven, CT: Yale University Press. 67–71.

Gillon, Brendan 1990. Ambiguity, Generality and Indeterminacy: Tests and Definitions. *Synthese* 85, 391–416.

Gillon, Brendan 2004. Ambiguity, Indeterminacy, Deixis and Vagueness. In: Steven Davis & Brendan Gillon (eds.) *Semantics*, Oxford, UK: Oxford University Press. 157–187.

Grice, H. Paul 1957. Meaning. *The Philosophical Review* 66(3), 377–388.

Heim, Irene 2000. Degree Operators and Scope. In: Brendan Jackson & Tanya Matthews (eds.) *Semantics and Linguistic Theory 10*. Ithaca, NY: CLC Publications, 40–64.

Heim, Irene & Angelika Kratzer 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Jacobson, Pauline 1999. Towards a Variable-Free Semantics. *Linguistics and Philosophy* 22, 117–184.

Jacobson, Pauline 2000. Paycheck Pronouns, Bach-Peters Sentences, and Variable-Free Semantics. *Natural Language Semantics* 8, 77–155.

Kamp, Hans 1975. Two Theories of Adjectives. In: Edward Keenan (ed.) *Formal Semantics of Natural Language*, Cambridge: Cambridge University Press. 123–155. Reprinted in 1984.

Kamp, Hans 1981. A theory of truth and semantic interpretation. In: J. Groenendijk & M. Stokhof (eds.) *Formal Methods in the Study of Language*, Amsterdam: Mathematical Centre. 277–322.

Kamp, Hans & Barbara Partee 1995. Prototype Theory and Compositionality. *Cognition* 57, 129–191.

Keefe, Rosanna & Peter Smith (eds.) 1997. *Vagueness: A Reader*. Cambridge, MA: MIT Press.

Kennedy, Christopher 1999. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland. (1997 UCSC Ph.D Dissertation).

Kennedy, Christopher 2001. Polar Opposition and the Ontology of 'Degrees'. *Linguistics and Philosophy* 24, 33–70.

1174 Kennedy, Christopher 2007. Vagueness and Grammar: The Semantics of Relative
1175 and Absolute Gradable Predicates. *Linguistics and Philosophy* 30(1), 1–45.

1176 Kennedy, Christopher & Louise McNally 2005. Scale Structure and the Semantic
1177 Typology of Gradable Predicates. *Language* 81(2), 345–381.

1178 Kennedy, Christopher & Louise McNally to appear. Color, Context and Composi-
1179 tionality. To appear in *Synthese*.

1180 Klein, Ewan 1980. A Semantics for Positive and Comparative Adjectives. *Linguis-*
1181 *tics and Philosophy* 4, 1–45.

1182 Lakoff, George 1970. A Note on Vagueness and Ambiguity. *Linguistic Inquiry*
1183 1(3), 357–359.

1184 Larson, Richard 1998. Events and Modification in Nominals. In: Devon Strolovitch
1185 & Aaron Lawson (eds.) *Proceedings from SALT VIII*. Ithaca, NY: CLC Publica-
1186 tions, 145–168.

1187 Ludlow, Peter 1989. Implicit Comparison Classes. *Linguistics and Philosophy* 12,
1188 519–533.

1189 McConnell-Ginet, Sally 1982. Adverbs and Logical Form. *Language* 58, 144–184.

1190 Nunberg, Geoffrey 1995. Transfers of Meaning. *Journal of Semantics* 12, 109–132.

Parsons, Kathryn 1973. Ambiguity and the Truth Definition. *Noûs* 7, 379–394.

Partee, Barbara H. & Mats Rooth 1983. Generalized Conjunction and Type Ambiguity. In: R. Bäuerle et al. (eds.) *Meaning, Use, and Interpretation of Language*, Berlin: de Gruyter. 361–383.

Pinkal, Manfred 1999. On Semantic Underspecification. In: H. Bunt & R. Muskens (eds.) *Computing Meaning*, Dordrecht: Kluwer Academic Press. 33–55.

Pinker, Steven 1994. *The Language Instinct: How the Mind Creates Language*. New York: Willian Morrow.

Postal, Paul 1969. Anaphoric Islands. In: Robert I. Binnick (ed.) *Proceedings of CLS 5*. Chicago, IL: Chicago Linguistic Society, 205–239.

Pustejovsky, James 1995. *The Generative Lexicon*. Cambridge: MIT Press.

Putnam, Hilary 1975. The Meaning of 'Meaning'. In: Keith Gunderson (ed.) *Language, Mind, and Knowledge*, University of Minnesota Press. 131–193.

Pylkkänen, Liina & Brian McElree 2006. The Syntax-Semantics Interface: On-line Composition of Sentence Meaning. In: M.A. Gernsbacher & M. Traxler (eds.) *Handbook of Psycholinguistics (Second Edition)*, New York: Elsevier. 539–580.

Raffman, Diana 1994. Vagueness without Paradox. *The Philosophical Review* 103, 41–74.

Raffman, Diana 1996. Vagueness and Context Relativity. *Philosophical Studies* 81, 175–192.

Raffman, Diana 2005. How to Understand Contextualism about Vagueness: Reply to Stanley. *Analysis* 65.3, 244–248.

Rips, L. & W. Turnbull 1980. How Big Is Big? Relative and Absolute Properties in Memory. *Cognition* 8, 145–174.

Rotstein, Carmen & Yoad Winter 2004. Total Adjectives vs. Partial Adjectives: Scale Structure and Higher-Order Modifiers. *Natural Language Semantics* 12, 259–288.

Saka, Paul 2007. *How to Think about Meaning*. Dordrecht: Springer.

Schwarzschild, Roger & Karina Wilkinson 2002. Quantifiers in Comparatives: A Semantics of Degree Based on Intervals. *Natural Language Semantics* 10, 1–41.

Seuren, Pieter A.M. 1973. The Comparative. In: Ferenc Kiefer & Nicolas Ruwet (eds.) *Generative Grammar in Europe*, Dordrecht: Riedel. 528–564.

Siegel, Muffy E.A. 1976. *Capturing the Adjective*. Ph.D. thesis, University of Massachusetts at Amherst.

Soames, Scott 1999. *Understanding Truth*. Oxford: Oxford University Press.

Stanley, Jason 2003. Context, Interest Relativity, and the Sorities. *Analysis* 63.4, 269–280.

von Stechow, Arnim 1984. Comparing Semantic Theories of Comparison. *Journal of Semantics* 3, 1–77.

Travis, Charles 1985. On What Is Strictly Speaking True. *Canadian Journal of Philosophy* 15(2), 187–229.

Travis, Charles 1994. On Constraints of Generality. *Proceedings of the Aristotelian Society* 94(2), 165–188.

Travis, Charles 1997. Pragmatics. In: Bob Hale & Crispin Wright (eds.) *A Companion to the Philosophy of Language*, Oxford: Blackwell. 87–106.

Tversky, Amos & Daniel Kahneman 1974. Judgment Under Uncertainty: Biases and Heuristics. *Science* 185, 1124–1131.

Ward, Gregory 2004. Equatives and Deferred Reference. *Language* 80, 262–289.

Ward, Gregory, Richard Sproat & Gail McKoon 1991. A Pragmatic Analysis of So-Called Anaphoric Islands. *Language* 67, 439–474.

Wheeler, Samuel 1972. Attributives and Their Modifiers. *Noûs* 6(4), 310–334.

Williamson, Timothy 1992. Vagueness and Ignorance. *Proceedings of the Aristotelian Society* 66, 145–162.

Williamson, Timothy 1994. *Vagueness*. London: Routledge.

Williamson, Timothy 1997. Imagination, Stipulation and Vagueness. *Philosophical Issues* 8, 215–228.

Zwicky, Arnold & Jerrold Sadock 1975. Ambiguity Tests and How to Fail Them. In: John P. Kimball (ed.) *Syntax and Semantics*, 4, Academic Press. 1–36.

Zwicky, Arnold & Jerrold Sadock 1984. A Reply to Martin on Ambiguity. *Journal of Semantics* 3, 249–256.

Keywords:

Ambiguity, vagueness, indeterminacy.

Christopher Kennedy, Chicago (USA)