Two Kinds of Subjectivity

Christopher Kennedy
University of Chicago

May 25, 2012

1 Introduction

What makes a predicate subjective? For example, why is it that (bare) predicates of personal taste like *tasty, disgusting* and *fun* are subjective according to certain tests (which we will examine in more detail below), but "phenomenological category" predicates like *semantic(al), metaphysical* and *epistemological* are not? To a certain extent at least, this question is independent of the question of how (or whether) subjectivity should be captured by semantic theory, in the sense that any theory that recognizes a subjective/nonsubjective distinction in the first place should be able to say what kinds of features constitute the difference between the subjective and the nonsubjective terms of the language. However, it is often the case that a better understanding of fine-grained details of the lexical semantic properties of particular classes of terms can inform our theoretical understanding and analysis of the categories and constructions that those terms enter into. The aim of this paper is to begin to develop such an understanding. My strategy will be to focus on one empirical domain, scalar predicates, and to look for patterns of distribution and interpretation of expressions that track the subjective/nonsubjective distinction. The conclusion will be that there are (at least) two kinds of subjectivity which are distinguished in that one affects distribution and one does not: subjectivity that is associated with vagueness generally, and subjectivity that is associated with predicates that involve qualitative assessment, independent of vagueness. A caveat: focusing on scalar predicates allows me to carry out a fine-grained examination of lexical semantic details, but it remains to be seen whether the conclusions that I reach about subjectivity in this domain can also help us understand subjectivity in modals, conditionals, and other kinds of constructions.

2 Subjectivity and selection

In order to investigate subjectivity, we need to know how to identify it. The most common diagnostic for subjectivity involves truth assessment: if two opposing sentences can be used to contradict each other without implying that one of the speakers has said something false, they are thought to involve some form of subjectivity. This kind of test is exemplified by the "faultless disagreement" pattern in (1) (Kölbel 2002).

(1) a. Anna: "Trippa alla romana is tasty."
    b. Beatrice: "Trippa alla romana is not tasty."

Beatrice's utterance in (1b) is understood as contradicting Anna's utterance in (1a), and so represents a kind of disagreement, yet we have the clear sense that both Anna and Beatrice are in some sense right, and so the disagreement is "faultless." This situation contrasts with the one in (2), in which we have a case of disagreement in which one of the speakers (Anna) is without question wrong, because trippa alla romana is made from the stomach lining of a cow or other animal.

(2) a. Anna: "Trippa alla romana is vegetarian."
    b. Beatrice: "Trippa alla romana is not vegetarian."

The scenario in (1) also contrasts with the one in (3), where we imagine a context in which Anna and Beatrice are looking at menus that, unknown to each other, differ on the dish that will follow the one they are currently eating: for Anna, it is trippa alla romana; for Beatrice, it is saltimbocca.

(3) a. Anna: "Trippa alla romana is next."
    b. Beatrice: "Trippa alla romana is not next."

As in (1) and unlike (2), neither Anna nor Beatrice make a false assertion, but in (3) this is because they are simply saying different things: there is no real disagreement in the first place.

Faultless disagreement and other phenomena related to truth assessment sensitivity have played a central role in theorizing about subjectivity, and in particular in deciding between contextualist and relativist accounts of the phenomenon. According to both accounts, subjectivity — and corresponding faultlessness in disagreements like (1) — arises when a predicate is relativized to a "judge", which is typically the speaker in simple assertions like the ones under discussion here. In a relativist analysis, the character and content of a subjective predicate are fixed, but its extension is judge-dependent (see e.g. Lasersohn 2005, 2009; Stephenson 2007; Kölbel 2002, 2009). On this view, Anna and Beatrice truly disagree about

whether the same property can be predicated of trippa alla romana; at the same time, their two utterances can both be true in virtue of the fact that the extension of *tasty* relativized to Anna may include trippa alla romana, while the extension of *tasty* relativized to Beatrice may not.

In a contextualist account, in contrast, the character of a subjective predicate is fixed, but its content and extension are judge-dependent (see e.g. Glanzberg 2007; Stojanovic 2007). Faultlessness in (1) is due to the fact that the property claimed to hold of trippa alla romana in Anna's utterance in (1a) is different from the property denied to hold of trippa alla romana in Beatrice's utterance in (1b): roughly *tasty for Anna* vs. *tasty for Beatrice*. The problem for this view is that it does not provide a clear explanation of why Anna and Beatrice appear to be in disagreement. In particular, without saying more, it fails to explain the difference between the discourse in (1) and the one in (3), in which a standard contextualist analysis of *next* derives exactly the right result: *next* in Anna's mouth means *next on Anna's menu* and *next* in Beatrice's mouth means *next on Beatrice's menu*, so there is no disagreement (and no fault).

While patterns of truth assessment like the faultless disagreement paradigm provide us with a diagnostic for subjectivity, and may also help distinguish competing theoretical analyses, they do not help us decide whether subjectivity correlates with some feature of the linguistic representation (cf. Stojanovic 2007). To answer this question, we need to ask whether there are patterns of interaction between subjective predicates and other expressions. The one I want to focus on here involves subjective attitude verbs such as English *find* in the construction *find x pred*, which is discussed in detail by Sæbø (2009). As shown by the pattern of acceptability in (4), this construction requires the predicate that heads the small clause complement of *find* to be subjective: replacing the subjective predicate *tasty* with a non-subjective predicate like *vegetarian* or *tasty for Beatrice* (in which *for Beatrice* makes the judge explicit, and renders the predicate non-subjective) results in unacceptability.

(4) a. Anna finds trippa alla romana tasty.
b. ?? Anna finds trippa alla romana to be vegetarian.
c. ?? Anna finds trippa alla romana tasty for Beatrice.

Speaking descriptively, these examples show that subject of *find* must be understood as the judge of the embedded predicate — (4a) means roughly the same thing as *Trippa alla romana is tasty to Anna* — and when this condition is not met, the resulting sentence is unacceptable.

Find is not the only verb whose subject can be understood as the judge of an embedded subjective term; the subjects of doxastic attitude verbs such as *think* and

believe, discussed by Stephenson (2007), can also be understood in this way. (5a), for example, is more or less synonymous with (4a).

(5) a. Anna believes trippa alla romana to be tasty.
b. Anna believes trippa alla romana to be vegetarian.
c. Anna believes trippa alla romana to be tasty for Beatrice.

However, unlike *find*, *believe* can also embed non-subjective predicates, as shown by the acceptability of (5b-c). Similarly, the contrast between (6a) and (6b) shows that only the subject of *find* is obligatorily understood to be the judge of the embedded predicate. (6a) is acceptable because a contextual individual (the cat) can be understood to be the judge of what counts as tasty, rather than the individual denoted by the subject of the attitude verb (Anna); (6b) is odd because Anna must be understood as the judge, which implies that she has been eating the cat food.

(6) a. Anna thinks/believes the cat food is tasty (because the cat ate it all up).
b. ?? Anna finds the cat food tasty (because the cat ate it all up).

Based on the contrast in (6), Stephenson (2007) proposes that *find* means the same thing as *think* or *believe*, but has an extra requirement that the doxastic anchor have direct experience of the embedded proposition. However, as Sæbø (2009) points out, this analysis doesn't explain the unacceptability of (7b), given the assumption that Homer has direct experience of his sexual orientation.

(7) a. Homer thinks/believes he is gay.
b. ?? Homer finds himself (to be) gay.

Sæbø argues that (7b) shows that *find* actually selects for a subjective predicate, and (in accord with the intuition stated above) fixes the judge of the embedded predicate to the semantic value of its subject. On this view, *find* does not itself introduce any truth-conditional content; it is instead a "radical judge-shifter."

Sæbø provides both a relativist and a contextualist implementation of the analysis. In the relativist variant, *find* causes the extension of the embedded predicate to be determined relative to its subject. On this view, the problem with (7b) is that the contribution of *find* is completely vacuous: since *gay* is non-subjective, its extension is judge-independent, and fixing its judge parameter to Homer makes no difference in meaning. In the contextualist version of Sæbø's analysis, subjective predicates are type-wise distinct from non-subjective predicates, in having an extra judge argument: *tasty* is type $\langle e, \langle e, t\rangle\rangle$, where the most external argument is the judge; *gay* is type $\langle e, t\rangle$. The function of *find*, on this view, is to bind the value of the judge argument of the embedded predicate to the matrix subject, and (7b) is

unacceptable the non-subjective predicate *gay* lacks a judge argument position: this is a case of type-mismatch.

In his paper, Sæbø discusses several patterns of data involving *find* and related verbs in other languages which support the hypothesis that there is a type-theoretic difference between the predicates that are acceptable under *find* and those that are not, and uses this result to draw conclusions about the right theory of subjectivity, arguing for a contextualist analysis of the sort proposed in Stojanovic 2007. In this paper, I want to accept Sæbø's conclusions about semantic type, but step back from the contextualism/relativism debate, and instead just use the subjective attitude verb construction as a way of probing the linguistic representation of subjectivity. In particular, I want to ask whether all predicates that are subjective in virtue of their behavior in faultless disagreement contexts are also acceptable in the *x find y pred* construction. If so, we may conclude that subjectivity uniformly has type-theoretic consequences. If not — as I believe the facts suggest — then we must recognize more than one form of subjectivity: one that involves a type-theoretic aspect of meaning, and one that does not.

3 Subjective standards vs. subjective evaluations

Richard (2004) observes that predicates of personal taste are not the only kinds of scalar predicates that display faultless disagreement effects. Instead, most (and maybe all) vague scalar predicates can give rise to this pattern:

(8) a. Anna: "Carla is rich/heavy/old/tall."
b. Beatrice: "No she's not!"

Richard points out that are two ways of understanding the dialogues in (8). On one interpretation, Anna and Beatrice appear to disagree because they are each implicitly comparing Carla to different groups of people or "comparison classes," e.g. for *rich*, the bottom 99% (for Anna) vs. the top 1% (for Beatrice) of the American population. But this is not a real disagreement: this situation is analogous to Anna and Beatrice's discussion about what is next on the menu in (3).

The other way of understanding these dialogues is one in which Anna and Beatrice are both implicitly comparing Carla to the same comparison classes but nevertheless have different assessments about whether she counts as rich, thin, heavy, and so forth. This constitutes a true case of faultless disagreement, so we must conclude that vague predicates in general are subjective, not just vague predicates of personal taste. In particular, even vague predicates which are used to describe (in principle) externally observable and measurable features of an object — such as its wealth (*rich*), its weight (*heavy*), its age (*old*), its height (*tall*), and so forth — can be subjective.

However, when we look at acceptability under *find*, we see that there is a subtle contrast between "dimensional" vague predicates such as the ones in (8a) and predicates of personal taste and other "evaluative" vague predicates. (These categories come from Bierwisch 1989.) Consider the latter group first. (9a) illustrates the acceptability of taste predicates under *find*, which we have already discussed; (9b) shows that other kinds of evaluative predicates are also acceptable here.

(9) a. Anna finds her bowl of pasta tasty/delicious/disgusting.
b. Anna finds Carla stimulating/annoying/boring/tedious.

In contrast, similar examples with dimensional vague predicates are odd:[1]

(10) a. ?? Anna finds her bowl of pasta big/large/small/cold.
b. ?? Anna finds Carla rich/heavy/old/tall.

This contrast is a subtle one, but a systematic examination of a range of data supports the conclusion that there is a real distinction here.

First, these examples can all be improved by adding an adverb that is derived from an unquestionably subjective adjective:

(11) a. Anna finds her bowl of pasta {surprisingly, remarkably, unusually} big/large/small/cold.
b. Anna finds Carla {annoyingly, disgustingly, irritatingly} rich/heavy/old/tall.

This makes sense if the adverb makes the whole predicate subjective (in the relevant way), which in turn implies that the predicate without the adverb is not subjective (in the relevant way).

Second, if an adjective has both a dimensional and an evaluative sense, *find* disambiguates to the latter. Consider adjectives like *heavy, light* and *dense*, used to describe a piece of cake. In the sentences in (12a), these adjectives can be understood in two distinct ways: as describing objective properties of the cake (its weight or density), or as describing a subjective assessment of the quality of the cake, made in virtue of the experience of tasting it.

(12) a. This piece of cake is heavy/light/dense.
b. I find this piece of cake heavy/light/dense.

[1]Note that there is a different sense of *find* meaning 'discover' which is acceptable when the embedded predicate is stage level: *When she returned to the table, Anna found her bowl of pasta cold/??big*. Because this sense of *find* is an achievement verb, while the one I am interested in is stative, I keep my example sentences in the simple present form in order to filter it out.

The examples in (12b), however, have only the latter reading: these sentences could be felicitously uttered after tasting the cake, but they would be very odd ways to report measurements of its physical properties. This is made particularly clear by the contrast in (13): the use of the measure phrase in (13b) blocks the subjective understanding that is available in (13a), and the resulting sentence is anomalous.

(13) a. I find this frosting thick.
b. ?? I find this frosting 2cm thick.

Finally, and most significantly, the examples in (10) can all be made perfectly acceptable by replacing the verb *find* with the verb *consider*, as in (14).

(14) a. Anna considers her bowl of pasta big/large/small/cold.
b. Anna considers Carla rich/thin/heavy/old/young/short.

The following pair illustrates this context in a particularly clear way. In a situation in which a passenger is checking his luggage for a flight, the attendant could report the fact that his bag exceeds the standard weight allowance by uttering (15a), but not by uttering (15b).

(15) a. I'm sorry, sir, but the airline considers this bag heavy. You will have to pay an extra baggage fee.
b. ?? I'm sorry, sir, but the airline finds this bag heavy. You will have to pay an extra baggage fee.

Similarly, in the discourse in (16), Anna can use *consider* to report her subjective assessment of whether Carla counts as tall in the context (one in which the vague standard is based on the needs of the team), but the corresponding sentence with *find* sounds strange.

(16) a. Anna: "We need a tall woman to play center on the basketball team."
b. Beatrice: "What about Carla? Is she tall?"
c. Anna: "I don't consider her tall."
d. ?? Anna: "I don't find her tall."

Syntactically, *consider* and *find* have almost identical distributions, and in constructions in which they are fully interchangeable, they appear to have quite similar semantic affects: (17a-b) sound synonymous, and in particular both are understood in a way that relativizes the embedded predicate to the surface subject.

(17) a. Anna finds the pasta tasty/beautifully presented.
b. Anna considers the pasta tasty/beautifully presented.

However, there is a crucial difference between *find* and *consider*: like *think* and *believe*, *consider* does not require its complement predicate to be subjective:

(18) a. Homer considers/??finds himself gay.
b. Homer considers/??finds trippa alla romana vegetarian.

These facts indicate that despite the similarity in meaning between the constructions in which *find* and *consider* occur, only the former is a "radical judge shifter" in Sæbø's sense. We may then conclude, based on the systematic difference in acceptability in dimensional and evaluative vague predicates under *find* and *consider*, that although the kind of subjectivity manifested by these predicates can give rise to faultless disagreement, it is not linguistically encoded in the same way.[2]

This conclusion can be both strengthened and refined by examining a second set of facts, which highlight another difference between subjectivity in dimensional predicates and subjectivity in evaluative predicates. We have already seen that both classes of predicates give rise to faultless disagreement effects in their vague, positive forms. When we turn to forms that are not vague, such as comparatives, we see that evaluative predicates still give rise to faultless disagreement:

(19) a. Anna: "The tripe is tastier than the haggis."
b. Beatrice: "No, the haggis is tastier than the tripe."

(20) a. Anna: "Skiing is the most fun!"

[2]There is, evidently, some degree of cross-linguistic variation here, or perhaps more likely, different subjective attitude verbs in different languages have slightly different properties. Sæbø (2009) shows that in Norwegian, the subjective attitude verb *synes* (which he glosses as 'seem') is incompatible with nonsubjective predicates, but is nevertheless acceptable with vague dimensional predicates, as shown in (i). (Here I follow Sæbø in just writing SUBJECTIVE ATTITUDE VERB in the free translation, to indicate the difficulty of providing true translations of these constructions.)

(i) a. ?? Mange forskere **synes** at dinosaurene ble utryddet av et voldsomt
Many researchers seem that dinosaurs were extinguished by a violent
kometnedslag for 65 millioner år siden.
comet.impact for 65 million years since
'Many scientists SUBJECTIVE ATTITUDE VERB that the dinosaurs were extinguished by a major comet impact 65 million years ago.'
b. De **synes** det er langt til lege.
they seem it is far to doctor
'They SUBJECTIVE ATTITUDE VERB it is far to the doctor.'

On the other hand, it may be the case that the dimensional predicates that Sæbø considers in his examples are like the ones discussed above in having both an objective and a subjective sense, with only the latter available in (ib). (The experience of distance can certainly be subjective; see the English examples with *long* below.) This should be a point for future investigation.

b. Beatrice: “No, skating is the most fun!”

(21) a. Anna: “Carla is more stimulating/annoying/boring/tedious than David.”

b. Beatrice: “No, David is more stimulating/annoying/boring/ tedious than Carla.”

Comparative forms of dimensional predicates, however, do not give rise to faultless disagreement: in the following dialogues, it is clear that one of Anna or Beatrice is right and the other is wrong.

(22) a. Anna: “The tripe is colder than the haggis.”

b. Beatrice: “No, the haggis is colder than the tripe.”

(23) a. Anna: “Skiing is the most expensive!”

b. Beatrice: “No, skating is the most expensive!”

(24) a. Anna: “Carla is richer/taller/heavier/older than David.”

b. Beatrice: “No, David is richer/taller/heavier/older than Carla.”

The same pattern appears in interactions with subjective attitude verbs. Comparative forms of evaluative predicates can be embedded under *find*, but comparative forms of dimensional predicates cannot be (cf. Sæbø 2009):

(25) a. Anna finds the tripe tastier than the haggis.

b. Beatrice finds skating the most fun.

c. Anna finds Carla is more stimulating/annoying/boring/tedious than David.

(26) a. ?? Anna finds the tripe colder than the haggis.

b. ?? Beatrice finds skating the most expensive.

c. ?? Anna finds Carla richer/taller/heavier/older than David.

Turning to adjectives that have both evaluative and dimensional senses, we see that this polysemy is retained in the comparative form: the disagreement between Anna and Beatrice in (27) can be understood either as a disagreement about their quantitative measurements of the cakes’ weight or density (the dimensional senses of the adjectives), or as an argument about their qualitative assessments of the cakes’ imprint on their taste/digestion (the evaluative senses of the adjectives).

(27) a. Anna: “This cake heavier/lighter/denser than that one.”

b. Beatrice: “No, that cake is heavier/lighter/denser than this one.”

However, this disagreement is faultless only on the latter understanding, for example if Anna and Beatrice are food critics. If instead they are food scientists reporting on a set of culinary experiments, then their disagreement is not faultless: one of them is right and the other is wrong, because her measurements were off, or she misread her instruments, or whatever. Similarly, if we report Anna’s utterance in

(27a) with (28a), it is ambiguous whether we are reporting on her taste experience (as a food critic) or on her measurements (as a food scientist). (28b), on the other hand, can only be understood as a description of Anna's subjective taste experience.

(28) a. Anna thinks that this cake is heavier/lighter/denser than that one.
b. Anna finds this cake heavier/lighter/denser than that one.

This kind of polysemy between an evaluative/qualitative/subjective sense and a dimensional/quantitative/objective sense appears in other classes of predicates as well (i.e., not only predicates that can be used to describe aspects of taste experiences), with similar results. Consider, for example, the dialoge in (29), uttered in a context in which it is an objective fact that the flight from Chicago to Tokyo takes 13 hours and 5 minutes, while the flight from Chicago to Hong Kong takes 15 hours and 40 minutes.

(29) a. Anna: "The flight from Chicago to Hong Kong is longer than the one from Chicago to Tokyo."
b. Beatrice: "No, the flight from Chicago to Tokyo is longer than the one from Chicago to Hong Kong."

There is one reading of (29) in which Anna is right and Beatrice is wrong. There is, however, a second reading in which their disagreement is faultless, but the disagreement has to do with their subjective experiences of the flight time, rather than about the objective durations of the flights. On the latter, subjective reading, Beatrice could justify her claim in virtue of the fact that she has to fly coach from Chicago to Tokyo, but gets to fly first class from Chicago to Hong Kong, and we could report her view using either (30a) or (30b). On the former, objective reading, in which Beatrice is wrong about the objective difference in flight time, only (30a) is an appropriate description of her mental state.

(30) a. Beatrice thinks that the flight from Chicago to Hong Kong is longer than the flight from Chicago to Tokyo.
b. Beatrice finds the flight from Chicago to Hong Kong longer than the flight from Chicago to Tokyo.

The facts that we have examined so far support two generalizations. First, vague, positive form scalar predicates give rise to faultless disagreement effects across the board, but only evaluative positive form scalar predicates are acceptable under *find*. Second, only evaluative scalar predicates in the comparative form (which is not vague; see Kennedy 2011) show faultless disagreement effects and acceptability under *find*; comparative forms of dimensional scalar predicates do not. The conclusion to be drawn from these generalizations is that the kind of subjectivity associated with vagueness — subjectivity about whether or not an object meets

a standard for satisfaction of the predicate — must be distinguished from the kind of subjectivity associated with evaluativity — subjectivity about an assessment of an object's qualities — because only the latter licenses embedding under subjective attitude verbs. In particular, the facts suggest that the component of meaning that is responsible for making a predicate one that expresses a qualitative assessment has a representational, type-theoretic status that the component of meaning that is responsible for making a predicate one that expresses a relation to a vague standard does not. Let us now see how this observation fits in to current thinking about vagueness and evaluativity.

4 Vagueness

Richard's observation that vague predicates in general give rise to faultless disagreement effects is, from one perspective, unsurprising, as a number of researchers have argued for an essentially relativistic semantics for vague predicates. For example, Bogusławski (1975) proposes that a vague scalar predicate is true of an object just in case the value of the object on the relevant scalar continuum is "conspicuous," "noteworthy" or "sufficient to attract attention." Similarly, Fara 2000 claims that an object satisfies a vague scalar predicate just in case the degree to which it manifests the relevant scalar property is "significant," and uses this semantic analysis to explain the phenomenological properties of vague predicates. Both of these characterizations imply that the interpretation of a vague form is relativized to some agent: the entity relative to whom conspicuousness, noteworthiness, or significance is assessed.[3]

Moreover, given common assumptions about the compositional semantics of gradable predicates, the fact that the positive form of adjectives like *rich, tall* and *old* gives rise to faultless disagreement while the comparative form does not, is also unsurprising. This is because a core semantic difference between the positive and comparative forms — one that must be captured by any empirically adequate semantic analysis of gradable predicates — is that the latter lacks whatever semantic (or pragmatic) features give rise to the vagueness of the former, and simply expresses an asymmetric ordering relation. In a Fara-style analysis, for example, *tall* expresses the interest-relative, vague property of having a degree of height that significantly exceeds some threshold (for a comparison class); *taller than David*, on the other hand, expresses the non-relativistic, non-vague property of having a degree of height that exceeds David's degree of height.

[3]This position has been called into question by Stanley (2003) on the grounds that we can have beliefs about the truth or falsity of a sentence like *Mt. Everest is tall* without having beliefs about any agent relative to whom Mt. Everest's height is supposed to be conspicuous, noteworthy or significant.

In many analyses of gradable adjectives and their various forms, this difference between positive and comparative is accounted for by hypothesizing that the adjectival root does not express a property on its own, but instead just introduces a mapping between individuals and scalar values, or degrees.[4] This basic idea is implemented in different ways in the literature; since these differences are not relevant to the questions I am asking here I will assume without argument the implementation in Kennedy 1999, in which gradable adjectives denote functions from objects to degrees: they combine with an individual and return a measure of how much that individual manifests a particular type of scalar value, such as height, age or wealth, as shown in (31).

(31) a. $[\![tall]\!] = \textbf{height}_{\langle e,d\rangle}$
b. $[\![old]\!] = \textbf{age}_{\langle e,d\rangle}$
c. $[\![rich]\!] = \textbf{wealth}_{\langle e,d\rangle}$

Gradable adjectives are converted into properties through composition with degree morphology, which in turn give rise to the observed differences in meaning between different forms of the adjective. Among the inventory of degree morphology in English is the comparative morpheme COMP (realized as the degree word *more* or the suffix *-er*) and, by hypothesis, a phonologically null positive form morpheme POS. It is the difference in meaning between these two morphemes that derives the differences in meaning between the positive and comparative forms, and which can in turn be used to explain their difference in faultless disagreement contexts.

The comparative morpheme maps the adjective onto a property that is true of an object iff its degree on the relevant scale exceeds the degree introduced by the *than*-constituent, as shown in given in (32).

(32) $[\![\text{COMP}]\!] = \lambda g_{\langle e,d\rangle}\lambda d_{than}\lambda x.g(x) \succ d_{than}$

The positive morpheme, on the other hand, maps the adjective onto a property that is true of an object iff its degree on the relevant scale is above a contextually appropriate threshold or "standard," e.g. one that represents the minimum value required to have a significant or noteworthy degree of the relevant property in the context of utterance. Its denotation is given in (33), where $\textbf{stnd}(g)$ represents "the standard

[4]Every semantic analysis of positive and comparative gradable predicates must include some way of deriving the difference in vagueness between the two forms. The approach that I present here is one way of doing so, but there are others which maintain the view that adjectives (both positive and comparative) denote properties (see e.g. Wheeler 1972; Kamp 1975; Klein 1980; van Rooij 2011). I adopt a degree-based, decompositional semantics of gradable adjectives here mainly because it allows for a transparent characterization of the difference between positive and comparative adjectival predicates, but I suspect that my central claims could be restated in terms of a different set of initial assumptions.

appropriate for the kind of measurement encoded by g," which, for the sake of this discussion, I am assuming to be computed in a Fara-style way.

(33) $[\![\text{POS}]\!] = \lambda g_{\langle e,d\rangle}\lambda x.g(x) \succeq \textbf{stnd}(g)$

The denotations of the comparative and positive forms of the adjectival root *tall*, according to this analysis, are as shown in (34).

(34) a. $[\![\text{COMP}]\!]([\![tall]\!]) = \lambda y\lambda x.\textbf{height}(x) \succ d_{than}$
b. $[\![\text{POS}]\!]([\![tall]\!]) = \lambda x.\textbf{height}(x) \succeq \textbf{stnd}(g)$

The comparative denotes a property that is true of an object just in case its height exceeds the degree introduced by the *than*-constituent. This is a precise property, and is moreover fully objective, since whether it holds of an object or not is fully determined by facts about that object's height. The positive, on the other hand, denotes the property of having a height that exceeds a standard of significance or noteworthiness, which is both vague (Fara 2000), and certainly subjective, since whether it holds of an object depends not only on that object's height, but also on some subjective assessment of significance (Richard 2004).

With these semantic assumptions in hand, we predict the following pattern: positive form adjectives should be subjective; comparative form adjectives should not be. If we restrict our empirical focus to dimensional adjectives and our diagnostics to faultless disagreement, this is indeed what we see, but once we broaden our empirical domain to include evaluative adjectives, and add subjective attitude verbs to our set of diagnostics, the picture becomes more complex. The following table summarizes the pattern we observed in the previous section:

(35)

	POSITIVE	COMPARATIVE
DIMENSIONAL	+FD, -FIND	-FD, -FIND
EVALUATIVE	+FD, +FIND	+FD, +FIND

What this table makes clear is that vagueness (i.e., positive form semantics) is a sufficient condition for faultless disagreement effects with scalar predicates, but not a necessary one. Vagueness is neither a necessary nor a sufficient condition for embedding under *find*, however; in fact, vagueness does not correlate with acceptability under *find* at all. Instead, whether a scalar predicate is acceptable under *find* correlates with the evaluative/dimensional distinction: based on the data we have seen, evaluativity is both a necessary and a sufficient condition for embedding under *find*. The behavior of comparative forms of evaluative adjectives shows moreover that evaluativity is a sufficient condition for faultless disagreement.

One conclusion to be drawn from these facts, if Sæbø (2009) is correct that a predicate is acceptable as the complement of *find* just in case it semantically selects for a judge (i.e., if its semantic type is such that it requires saturation by an

individual-denoting expression which is interpreted as the source of subjective assessment), is that the kind of subjectivity introduced by vagueness is not one that is reflected in semantic type. Putting things in terms of the compositional assumptions introduced above: the judge argument is not introduced by POS, but rather by the adjective, since the crucial factor determining acceptability under *find* is not positive vs. comparative (vague vs. not vague), a distinction that comes from the functional elements, but rather evaluative vs. dimensional, a distinction that comes from the lexical items.[5]

There are a number of plausible ways of explaining and analyzing the subjectivity of vagueness: relativistically, in virtue of a particular semantics of the positive form (Richard 2004); at a more discourse/pragmatic level, in virtue of the dynamics and uncertainty of vague standards (Barker 2002); or in some other way which does not have the same type-theoretic consequences as the kind of subjectivity that we see with (both positive and comparative) evaluative predicates. Identifying the best model of this kind of subjectivity is not a goal of this paper; instead, I want to focus on the question of how to model the kind of subjectivity manifested by evaluative predicates. We have seen that evaluative predicates must be distinguished from dimensional ones, in both their positive and comparative forms, in a way that licenses embedding under *find*. Following Sæbo, I take this distinction to be a type-theoretic one; the semantic question that now needs to be answered is what this distinction corresponds to in terms of meaning. In the next section, I will propose an answer to this question that builds on previous work by Kennedy and McNally (2010) on the semantics of color adjectives.

5 Evaluativity as qualitative assessment

One of the conclusions of the previous section, based on the interaction of adjective class and acceptability of embedding under the subjective attitude verb *find*, is that the dimensional/evaluative distinction is a type-theoretic one. The most straightforward way of capturing this distinction is to say that it is lexical, and indeed, this is the route taken in Bierwisch 1989, the most comprehensive discussion of the dimensional/evaluative distinction in the literature, but unfortunately not one that will help to explain the patterns we have seen in this paper. Specifically, Bierwisch proposes that the difference between dimensional and evaluative adjectives is that the former lexically encode degree functions, much in the way described in the previous section, while the latter are underlyingly "regular" (type $\langle e, t\rangle$) prop-

[5] Sæbo actually draws the opposite conclusion, based on facts like those discussed in footnote 2; what we need to do now is reassess the examples he discusses with the evaluative/qualitative vs. dimensional/quantitative distinction in mind.

erties. Their gradable meanings are derived using a special mapping function that turns them into degree predicates, and which accounts for various entailment patterns that Bierwisch wants to explain, but which does not introduce anything like a judge argument, and so does not give us a way of explaining the different pattern of acceptability of dimensional and evaluative adjectives under *find*.

Instead, what we would like to be able to say is something like the following. Dimensional adjectives are regular type $\langle e, d\rangle$ measure functions of the sort listed in (31), which combine with degree morphology to derive various kinds of properties as outlined above. Evaluative adjectives, on the other hand, denote judge-dependent, "subjective measure functions" and have the semantic type $\langle e, \langle e, d\rangle\rangle$. However, merely saying that the two classes of adjectives differ in semantic type in this way is not enough; we also need to say what the difference in actual meaning is. In other words, we need to make some sense of what it means to say that an evaluative adjective denotes a "subjective measure function."

As a starting point, let us take a look at adjectives like *salty, sugary, watery*, and so forth, which are similar to *heavy* and *dense* in that they have both quantitative senses and qualitative ones.[6] (36), for example, can be understood either as in (36a) or as in (36b). (I focus on comparatives here to eliminate any role of vagueness, but the same kinds of patterns hold for the non-comparative forms as well.)

(36) This dish is saltier than that one.
   a. This dish contains more salt than that one.
   b. This dish has a more (subjectively) salty quality than that one

These readings are truth-conditionally distinct. If the dishes in question contain different amounts of salt, but for some reason the salt taste is undetectable, then (36) is true on the (a) reading and false on the (b) reading. If the dishes actually contain no salt, but one tastes saltier than the other for some reason, then (36) can be true on the (b) reading and false on the (a) reading. Furthermore, disagreement dialogues and *find* distinguish the two readings in the expected way. Anna and Beatrice's disagreement in (37) is faultless only on the qualitative understanding of *saltier*, and only the qualitative understanding is present under *find* in (38).

(37) a. Anna: "This dish is saltier than that one."
   b. Beatrice: "No, that dish is saltier than this one."

(38) I find this dish saltier than that one (even though I know that it contains less salt).

[6] In the discussion of *heavy* and *dense* in Section 3, I referred to their dimensional senses as "quantitative" and their evaluative senses as "qualitative." I think that the quantitative/qualitative distinction is the more general one, of which the dimensional/evaluative distinction is a subtype, for reasons that will become clear as we move forward.

Adjectives of tastes (which are different from adjectives of personal taste), such as *salty, sugary*, and so forth, are derived from nouns, and in this respect are similar to another class of adjectives which show a quantitative/qualitative ambiguity: color words (Kennedy and McNally 2010). For example, (41) can be understood either quantitatively, as in (41a), or qualitatively, as in (41b), and the two readings can be teased apart through the use of different modifiers, as shown in (40).

(39) This leaf is greener than that one.
   a. More of this leaf is green than that one.
   b. This leaf is qualitatively closer to "pure green" than that one.

(40) a. The leaf is completely green. *quantitative*
   b. The leaf is perfectly green. *qualitative*

Adapting Bierwisch's (1989) analysis of evaluative adjectives, Kennedy and McNally account for this ambiguity by hypothesizing that the basic meaning of a color word is the nominal one, which they take to be the name of a color (a kind), and that there are two ways of deriving a meaning as a gradable adjective.[7] One way involves mapping the nominal meaning to an adjective meaning that measures the quantity of (the relevant) color manifested by the object; the other way is to map the nominal meaning to an adjective meaning that measures the quality of (the relevant) color manifested by the object. Kennedy and McNally give a decompositional semantics which makes use of two primitive functions **quant** and **qual** which, when combined with a color noun meaning, derive the respective qualitative and quantitative adjective meanings, as in (41)

(41) a. $[\![[_N \text{ green}]]\!]$ = **green**, the name of a kind
   b. $[\![[_{Aquant} \text{ green}]]\!] = \lambda x.\textbf{quant}(\textbf{green})(x)$: a function from an individual x to a degree which represents the quantity of **green** manifested by x
   c. $[\![[_{Aqual} \text{ green}]]\!] = \lambda x.\textbf{qual}(\textbf{green})(x)$: a function from an individual x to a degree which represents the quality of **green** manifested by x

I would like to suggest that predicates of tastes — and predicates that involve qualitative assessment more generally — should be analyzed in basically the same way, with one crucial modification: expressions whose meanings are defined in terms of **qual** have an extra argument that corresponds to the source of qualitative assessment, which I will represent as an index on the **qual** function. The two senses

[7] Kennedy and McNally also introduce a third, non-gradable adjective meaning for color words, which is not relevant for the current discussion. See Hansen (2011) for detailed assessment of Kennedy and McNally's proposals.

of *salty* can then be characterized formally as in (42), where **salt** is the denotation of the noun *salt* (a kind).

(42) a. $[\![[_{Aquant}\ \text{salty}]]\!] = \lambda x.\mathbf{quant}(\mathbf{salt})(x)$
b. $[\![[_{Aqual}\ \text{salty}]]\!] = \lambda x \lambda y.\mathbf{qual}_y(\mathbf{salt})(x)$

(42a) is a measure function that maps its argument x onto a degree that represents the quantity of **salt** manifested by x: how much salt it contains. (42b), on the other hand, is a "subjective measure function" which maps its argument x onto a degree that represents the quality of **salt** manifested by x according to y — it gives a measure of y's subjective assessment of x in relation to salt.

The intuition underlying this proposal is that qualitative measurements are distinct from quantitative ones in being relativized to a source of qualitative assessment, and that this difference is reflected in the argument structure of the predicates used to make the relevant kinds of measurements. In particular, this distinction is reflected type-theoretically in a way which, following Sæbø (2009), can explain the difference in acceptability of qualitative vs. quantitative meanings under *find*: qualitative meanings of adjectives like *salty* have a judge argument which can be saturated by the surface subject of *find*; quantitative meanings of the same adjectives do not. [8]

Adjectives of tastes like *salty, sugary, watery* have both qualitative and quantitative interpretations, as do adjectives based on dimensional concepts like weight (*heavy*), elapsed time (*long*) and density (*dense*), as we have seen. There are, I suspect, many other adjectives that pattern in the same way; adjectives based on

[8] An apparent problem for this proposal — or at least for the specific hypothesis that the qualitative/quantitative ambiguity that I have identified with the evaluative/dimensional distinction is the same as the qualitative/quantitative ambiguity that Kennedy and McNally (2010) posit for color words — is that color words are not embeddable under *find*, even when understood qualitatively:

(i) a. ?? I find this object green/blue/red/yellow.
b. I consider this object green/blue/red/yellow.

I would like to suggest that the problem with (ia) is not that the qualitative meaning of color words should be analyzed in some other way, but rather that color words are special in that their judge arguments are lexically saturated by an argument that has the semantic consequence of making the source of qualitative assessment the "normal observer;" in Lasersohn's (2005) terms, they are obligatorily EXOCENTRIC. This can be viewed as a consequence of the fact that although color can be manifested either quantitatively or qualitatively, the qualitative experience of color is something that is taken to be shared across individuals, rather than subject to individual variation in the way that experiences of taste, etc. are. Color adjectives do give rise to faultless disagreement patterns, as discussed in Hansen 2011, but this is arguably due to the fact that they are vague, albeit possibly in ways that involve more multidimensionality than the kinds of vague predicates discussed in Section 3.

substance terms, such as *plastic, metallic, wooden, icy*, and so forth provide one class of examples. At the same time, there are some adjectives which appear to have only quantitative readings (e.g., *tall* and *rich*), and many which appear to have only qualitative meanings (e.g., predicates of personal taste). This could indicate that these adjectives lexicalize one meaning but not the other, but it may also simply indicate that there are concepts which are naturally mapped to both quantitative or qualitative measurements, while others are naturally mapped only to one kind of measurement or the other. For example, it is natural to talk about a subjective experience of the salt manifested by a particular bowl of soup, but a bit odd (though maybe not impossible) to talk about the subjective experience of the height manifested by a particular individual. Likewise, it is natural to talk about an objective measurement of the quantity of salt in a bowl of soup, absent any subjective experience of it, but it is not so natural to talk about the objective quantity of taste in the soup, without any subjective experience of it.

In any case, I take the existence of a large class of adjectives that display a quantitative/qualitative ambiguity, and the fact that the two readings give rise to different patterns of distribution under *find* (as well as to other distributional differences of the sort observed by Kennedy and McNally 2010), as evidence for the core hypothesis that the analysis presented here aims to capture: that adjective meanings can be broken down into a basic concept term that provides some dimension of measurement, plus either a quantitative or (relativized) qualitative way of measuring an object relative to that dimension. Whether the particular implementation of this hypothesis presented above is the best way to capture the qualitative/quantitative distinction remains to be seen. I myself doubt that it is, and suspect that a more explanatory account will emerge out of careful thought about the basic elements of predicate meaning, and also by a close examination of languages in which predicate expressions show more grammatical complexity than they do in English (see, for example, Koontz-Garboden and Francez 2010). That said, the current proposal is at least based on a pattern of linguistic behavior that has some generality, and so hopefully represents a first step towards a more satisfying and explanatory analysis.

6 Conclusion

In this paper, I have compared different classes of scalar predicates relative to two diagnostics of subjectivity: faultless disagreement effects and acceptability under the subjective attitude verb *find*. We have seen that these diagnostics provide the basis for distinguishing two types of subjectivity: the kind that of subjectivity manifested by vague predicates, which triggers faultless disagreement but does not ensure embeddability under *find*, and the kind of subjectivity manifested by evaluative

predicates, which ensures embeddability under *find* and also triggers faultless disagreement, independent of vagueness. Following Sæbø (2009), I have assumed that acceptability under *find* is a function of semantic type, and I proposed a lexical semantic analysis of evaluativity as judge-dependent qualitative assessment that builds on Kennedy and McNally's (2010) analysis of quantitative vs. qualitative meanings of color adjectives. This means that the analysis is committed to a contextualist semantics for the class of predicates which involve qualitative assessment, which, according to the diagnostics discussed in this paper, include predicates of personal taste. This is a conclusion already reached by Sæbø (2009) (who advocates in particular the contextualist analysis proposed in Stojanovic 2007).

An open question, then, is how to reconcile a contextualist semantics for predicates of personal taste and other predicates of qualitative assessment, which is justified on a purely grammatical basis by the interaction with subjective *find*, with the fact that they can be used in dialogues that appear to express genuine disagreement. (Recall from Section 2 that the problem that the faultless disagreement pattern presents for a contextualist semantics is not faultlessness, but disagreement.) Unfortunately, this is not a question that I can satisfactorily answer at this time, but let me finish this paper with a few thoughts about it. The argument against contextualism based on disagreement crucially relies on the assumption that the presence of a genuine disagreement (faultless or otherwise) in a linguistic exchange entails that the speakers engaged in the exchange are using sentences that have the same (relevant) semantic content, something that is not the case if the properties expressed by qualitative expressions vary with the judge. However, this "shared content" assumption has been challenged in recent work by Sundell (2010) and Plunkett and Sundell (2012), who shows that genuine disagreement can arise even in exchanges that involve sentences with demonstrably distinct semantic content (e.g., disagreements that involve implicatures). If genuine disagreement does not in fact entail shared content, then one of the central arguments against a contextualist account of predicates of personal taste disappears. (That said, the contextualist analysis still needs to be supplemented with an account of e.g. the difference between "bare" taste predicates like *tasty*, which give rise to disagreement, and predicates with explicit judges, like *tasty to me*, which do not.)

A final word on faultless disagreement with vague dimensional adjectives: if this is due to some aspect of the meaning of the positive form, (e.g. its interest relative semantics, if Fara's (2000) analysis of vagueness is correct), then given that this aspect of meaning is absent from the comparative form of evaluative adjectives, but such adjectives still give rise to faultless disagreement, we must conclude that faultless disagreement can have more than one source. This would further entail that we cannot conclude from the observation that a particular expression E shows faultless disagreement effects that it has some semantic/pragmatic feature F (e.g., a

relativistic meaning), because the same effects could be due to some other feature F' (e.g., type-theoretic subjectivity). On the other hand, it may be the case that faultless disagreement in these two cases is due not to the semantics of the positive form of a dimensional adjective, nor to the type-theoretic subjectivity of a comparative evaluative adjective, but rather to some more general, non-representational property of meaning that both expressions share. In order to answer this question, we need to think more about the ways that the two classes of expressions are similar, and about the features they share with other kinds of expressions that give rise to faultless disagreement.

References

Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25:1–36.

Bierwisch, Manfred. 1989. The semantics of gradation. In *Dimensional adjectives*, ed. Manfred Bierwisch and Ewald Lang, 71–261. Berlin: Springer-Verlag.

Bogusławski, Andrzej. 1975. Measures are measures: In defence of the diversity of comparatives and positives. *Linguistiche Berichte* 36:1–9.

Fara, Delia Graff. 2000. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics* 20:45–81.

Glanzberg, Michael. 2007. Context, content, and relativism. *Philosophical Studies* 136:1–29.

Hansen, Nat. 2011. Color adjectives and radical contextualism. *Linguistics and Philosophy* 34:201–221.

Kamp, Hans. 1975. Two theories of adjectives. In *Formal semantics of natural language*, ed. Edward Keenan, 123–155. Cambridge: Cambridge University Press. Reprinted in 1984.

Kennedy, Christopher. 1999. *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland. (1997 UCSC Ph.D Dissertation).

Kennedy, Christopher. 2011. Vagueness and comparison. In *Vagueness and language use*, ed. Paul Egré and Nathan Klinedinst, 73–97. Houndmills, UK: Palgrave.

Kennedy, Christopher, and Louise McNally. 2010. Color, context and compositionality. *Synthese* 174:79–98.

Klein, Ewan. 1980. A semantics for positive and comparative adjectives. *Linguistics and Philosophy* 4:1–45.

Kölbel, Max. 2002. Faultless disagreement. *Proceedings of the Aristotelian Society* 104:55–73.

Kölbel, Max. 2009. The evidence for relativism. *Synthese* 166:375–395.

Koontz-Garboden, Andrew, and Itamar Francez. 2010. Possessed properties in Ulwa. *Natural Language Semantics* 2:197–240.

Lasersohn, Peter. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28:643–686.

Lasersohn, Peter. 2009. Relative truth, speaker commitment, and control of implicit arguments. *Synthese* 166:359–374.

Plunkett, David, and Timothy Sundell. 2012. Disagreement and the semantics of normative and ethical terms. Ms., UCLA and University of Kentucky.

Richard, Mark. 2004. Contextualism and relativism. *Philosophical Studies* 119:215–242.

van Rooij, Robert. 2011. Implicit versus explicit comparatives. In *Vagueness and language use*, ed. Paul Egré and Nathan Klinedinst, 51–72. Houndmills, UK: Palgrave MacMillan.

Sæbø, Kjell Johan. 2009. Judgment ascriptions. *Linguistics and Philosophy* 32:327–352.

Stanley, Jason. 2003. Context, interest relativity, and the sorities. *Analysis* 63.4:269–280.

Stephenson, Tamina. 2007. Towards a theory of subjective meaning. Doctoral Dissertation, Massachusetts Institute of Technology.

Stojanovic, Isadore. 2007. Talking about taste: Disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy* 30:691–706.

Sundell, Timothy. 2010. Disagreements about taste. *Philosophical Studies* 155:267–288.

Wheeler, Samuel. 1972. Attributives and their modifiers. *Noûs* 6:310–334.

Department of Linguistics
University of Chicago
1010 E. 59th St.
Chicago, IL 60637

ck@uchicago.edu