

# Thick Evaluative Terms and Gradability

Pekka Väyrynen

June 7, 2008

## 1 Introduction

The distinction between “thick” and “thin” evaluative and normative terms or concepts, and its importance to moral theory, has been an active topic in recent metaethics. But the reputation this debate has for obscurity is not, I think it fair to say, wholly unearned. We can get at the relevant class of concepts via examples. Thick concepts are those often expressed by words like *lewd*, *cruel*, *nasty*, *generous*, and *kind*, thin concepts those often expressed by words like *good*, *wrong*, and *ought*.<sup>1</sup> But it is even hard to find an adequate neutral description of what makes an evaluative concept count as thick and distinguishes these from concepts that are thin in some relevant contrasting sense. In a widely cited discussion, Bernard Williams takes thick concepts to express some kind of “union of fact and value” in that their applicability is both “action-guiding” (it typically indicates the presence of reasons for action) and “world-guided” (it depends on how the world is in certain non-evaluative respects).<sup>2</sup> Accordingly, people usually take a term to stand for a

---

<sup>1</sup>I use italics to indicate words or phrases mentioned, and boldface to indicate semantic values.

<sup>2</sup>See Williams (1985: 128, 140). I say “often” because it can be perfectly felicitous to use a term like *cruel* merely to descriptively categorize something. So the label “a thick evaluative term” functions here as shorthand for those occurrences of words from the relevant class which function evaluatively. Here I remain open as to what exactly *that* involves. It is one of the issues at stake in these debates.

thick concept if it expresses an evaluative/normative concept with significant non-evaluative/normative content. Thick concepts are then supposed to differ from thin concepts in the way they seemingly combine description and evaluation in some sense as a matter of their meaning. But it remains unclear whether this way of describing the distinction succeeds in drawing any significant or even accurate contrast.<sup>3</sup>

In this paper I'll bypass attempts to characterize the thick-thin distinction. Instead I'll try to use only some general tools from philosophy of language to make more tractable at least some of the issues about the semantics of evaluative and normative terms which are at stake in debates about thick evaluative concepts. I'll advance three claims in particular. (1) Adjectival forms of thick evaluative terms fall into the class of gradable adjectives. (2) The standard sort of treatment of gradable adjectives in formal semantics casts doubt on a family of arguments concerning thick evaluative terms. (3) This semantics can be used to make fruitful sense of a number of philosophical disputes in the literature on thick concepts as reflecting (implicit) disagreements about the "metasemantics" rather than the basic semantics of thick evaluative terms. This result will hopefully help us see where to look for solutions to the deeper issues at stake in debates about thick evaluative concepts.

## 2 The semantics of gradable adjectives

Adjectival forms of thick terms are gradable because they admit of degree constructions. First, they admit comparatives: Things can be *more generous* or *more nasty* than others, *less cruel* or *less courageous* than others. Second, they take degree modifiers: Things can be *very generous*, *somewhat nasty*, *extremely greedy*, etc.<sup>4</sup>

---

<sup>3</sup>Eklund (unpublished) levies serious criticisms of the existing accounts of thick-thin distinction.

<sup>4</sup>Of course, the adjectival forms of at least some thin terms are also gradable; consider *good*. It is controversial whether such non-adjectival epithets as *schmuck*, *jerk*, and *wop* are relevantly like thick evaluative terms. But it may not be a problem if there are thick terms that have no adjectival form. Even a term that has no gradable adjectival form might still be semantically linked to a scale of some sort. This is all that the discussion to come requires in order to generalize.

One standard sort of semantic analysis is that gradable adjectives, like *tall*, locate objects on a scale, like the scale of tallness. In explaining this analysis I'll draw on an account due to Kennedy (2007) and its simplified presentation in Glanzberg (2007). On this analysis, the basic meaning of a gradable adjective is a *function* from objects to degrees (understood as abstract representations of measurement), ordered by some (total) ordering with respect to some property (such as height in the case of *tall*). Call this property a *dimension*, and call a set of degrees (totally) ordered with respect to some dimension a *scale*. So the semantic value of a gradable adjective is a function from objects to degrees of (or, values on) a scale.

The interpretation of the comparative form is simple. (1a) is analyzed as (1b-c):

- (1) a. Stan is taller than Buster
- b. The value Stan takes on a scale of tallness is greater than the value Buster takes on a scale of tallness.
- c.  $\mathbf{tall(Stan)} > \mathbf{tall(Buster)}$  [where **tall** indicates a semantic value, interpreted here as a degree on scale]

Regarding the positive form, we want to capture the context-sensitivity of gradable adjectives. For instance, someone can count as tall in a context where silent film comedians are under discussion but not tall in a context where basketball players are under discussion. According to Kennedy, the interpretation of the positive form is that the degree to which something is tall is greater than some particular degree value  $d$  determined by context. So (2a) gets analyzed as in (2b-c):

- (2) a. Stan is tall.
- b. The value Stan takes on a scale of tallness is greater than (or, at least as great as) the minimal value required for counting as satisfying *tall* in context  $c$ .
- c.  $\mathbf{tall(Stan)} > d$
- d.  $\mathbf{tall(Stan)} > s(\mathbf{tall})$

The value of  $d$  (determined for the adjective by context) is the *standard* for that adjective and context. (Kennedy instructs us to prefer (2d) to (2c). He thinks we must take  $d$  to be determined by context from, so far as possible, the meaning of the given adjective. In (2d),  $s$  is a contextually given function that determines a contextually significant degree  $d$ , based on the interpretation of the given adjective.)

Basic applications of this analysis to thick evaluative terms give intuitively the right results. Someone can count as courageous relative to a context where school-yard bullies or jumping off a three-meter diving board are under discussion, but not count as courageous when what is under discussion is getting across Omaha Beach, defusing bombs, or smoke jumping. Here, too, whether something counts as satisfying *courageous* is a matter of whether the value it takes on an appropriate scale is greater than what is required for counting as courageous in the context.

### 3 Some arguments concerning thick evaluative terms

I will now argue that this analysis of gradable adjectives casts doubt on various arguments concerning the nature of the relationship between the evaluative and non-evaluative “sides” or “aspects” of thick evaluative terms or the concepts they express. Here are three common but controversial theses about that relationship:

*Multivalence thesis* (MT): For any thick evaluative term  $T$ , utterances of  $T$  may express something with positive valence in some contexts but negative (or neutral) in others.

*Pragmatic Thesis* (PT): For any thick evaluative term  $T$ , the evaluative content  $E$  (including its valence) expressed by utterances of  $T$  is part of the pragmatic rather than semantic content of utterances of  $T$ .

*Disentanglement thesis* (DT): For any thick evaluative term  $T$ , the evaluative content  $E$  and the non-evaluative, descriptive content  $D$  expressed by utterances of  $T$  are distinct, separable components.

These formulations are deliberately generic and admit of various precisifications.<sup>5</sup>

(MT) has often been used to argue for (DT) and (PT).<sup>6</sup> The idea is that it seems difficult to explain how the valence of evaluation could vary if the meaning of a thick evaluative term were some inextricable fusion of evaluation whose on account of certain description. (PT) seems capable of explaining such variability while indicating a sense in which thick evaluative claims have distinct, separable evaluative and non-evaluative aspects, such that we can describe something in purely non-evaluative terms while withholding any evaluation that may normally be even tightly associated with the corresponding thick term (cf. Hare 1981: 73-74).

I take no stand on the truth of (MT), (PT), or (DT) here. I lack the space to assess their different versions carefully. But the extant arguments for them (or against them, for that matter) are mostly flawed. What I wish to suggest is that scale semantics for gradable adjectives can help us explain the flaws in these arguments. Here I focus mainly on arguments for (MT) that proceed from examples.

**3.1 Argument from *too F*.** R. M. Hare (1952: 121) offers the following example in support of (MT):

‘[T]idy’ and ‘industrious’ . . . are normally used to commend; but we can say, without any hint of irony, ‘too tidy’ or ‘too industrious.’

The context of the example indicates that it is meant to support a thesis like (MT). The argument it is meant to illustrate must go something like this:

(P1) For any thick evaluative term  $T$ , if utterances of  $T$  normally express something with positive (negative) valence but *too T* expresses something with negative (positive) valence, then  $T$  is multivalent.

---

<sup>5</sup>For example, at this point (MT) and (DT) are meant to be non-committal as to whether what is expressed includes only “what is said,” or everything that is communicated by an utterance, or something in between. Likewise, the relationship asserted in (DT) can be interpreted as analytic or not, *a priori* or not, and so on.

<sup>6</sup>I mean ‘often’ as a sociological point. *Direct* arguments of this kind seem invalid.

- (P2) For many  $T$ , utterances of  $T$  normally express something with positive (negative) valence but utterances of *too*  $T$  express something with negative (positive) valence.
- (C1) So, many thick evaluative terms are multivalent.
- (P3) Thick terms are, in this regard, homogeneous.<sup>7</sup>
- (C2) So, (MT).

It should be clear that this argument is bad. Either (P1) is false or, at minimum, Hare's example cannot support (P1). We cannot conclude anything much about what some word  $F$  means from what *too*  $F$  means any more than from what *not*  $F$  means. But scale semantics can explain where Hare goes wrong. Consider:

- (3) a. Michael is too industrious.  
 b. Michael is too industrious to have time to visit Bordeaux wineries.  
 c. Michael is too industrious for me to keep up with his work.

*Too*-constructions are comparative (Meier 2003). To determine what (3a) says, we need a standard of comparison from context. We can fix it explicitly with a *to-* or *for-*phrase as in (3b-c); otherwise context must set the parameter. The standard determines upper and lower bounds of admissibility – that is, an *interval* – on a scale associated with the adjective (Schwarzschild and Wilkinson 2002). Depending on the adjective's polarity, a *too*-construction places the object above the upper bound (*too expensive*) or below the lower bound (*too young*) of the interval.

In a context where (3a) means (3b), we can analyze (3a) as in (3d-e):

- (3) a. Michael is too industrious.

---

<sup>7</sup>Discussions of thick evaluative often assume implicitly that such terms are sufficiently uniform for examples to generalize straightforwardly. Such assumptions can certainly be challenged, but I set that issue aside here.

- d. The value Michael takes on a scale of industriousness is greater than the maximum admissible value determined by context  $c$ .
- e. **industrious(Michael)** >  $max_{d^*}\{[\mathbf{industrious(Michael)} = d^*] \rightarrow \text{Michael has time to visit Bordeaux wineries}\}$ <sup>8</sup>

What this argument from *too F* neglects is that, in a given context, the standard for counting as satisfying *too F* is typically not determined by the same set of contextual inputs as the standard for counting as satisfying  $F$ . Those inputs are typically richer in the former case. So we shouldn't expect to be able to conclude anything much about the meaning or evaluative valence of  $T$  from the valence of *too T*.

**3.2 Argument from *not F enough*.** Simon Blackburn (1992: 296) offers the following example in support of (MT):

[We may] worry that this year's Carnival was not lewd enough.

The background assumption here is that calling something *lewd* usually expresses negative evaluation. If the example is to support (MT), it must be that if a greater degree of lewdness would have been welcome, then what lewdness there was, doesn't, as such, count as bad or objectionable.

The argument from this claim to (MT) is parallel to the argument from *too F*. So, the crucial premise is that if utterances of  $T$  normally express something with positive (negative) valence but utterances of *not T enough* express something with negative (positive) valence, then  $T$  is multivalent. The flaw is also the same. The standard for counting as satisfying *not F enough* is typically not determined by the same contextual inputs as the standard for counting as satisfying  $F$ .

Also (4) suggests that *enough*-constructions provide no support for (MT):

- (4) The carnival was lewd. It made me squirm. And yet it was not lewd enough.

---

<sup>8</sup>(3e) is one plausible gloss of (3d). Other ways of glossing it are available.

(4) can be fine if context shifts between the two occurrences of *lewd*. The carnival's degree of lewdness may have been insufficient to count – and this may have been salient in context – as one of those “necessary eruptions of the Dionysiac into the fragile Apollonian order” which even prudes can tolerate (Blackburn 1992: 296).

**3.3 Argument from the positive form.** Blackburn offers another example in support of (MT):

A previous Master of my Oxford College [call him Dr. Tweedy] was . . . condemned as frugal, a fatal flaw in someone whose main job is dispensing hospitality (Blackburn 1992: 286; cf. Hare 1981: 73-74).

Assume Dr. Tweedy counts as satisfying *frugal* in the context. (We may also assume that ‘Dr. Tweedy is frugal’ means he is frugal for a college master.) The argument must go something like this if the example is to support (MT):

- (P4) For any thick evaluative term  $T$ , if utterances of  $T$  normally express something with positive (negative) valence but some utterances of  $T$  express something with negative (positive) valence (or something evaluatively neutral), then  $T$  is multivalent.
- (P5) For many  $T$ , utterances of  $T$  normally express something with positive (negative) valence but some utterances of  $T$  express something with negative (positive) valence (or something evaluatively neutral).
- (C1) So, many thick evaluative terms are multivalent. [Plus (P3) and (C2) as before.]

Unlike the arguments from *too F* and *not F enough*, here the crucial premise (P4) at least states a plausible sufficient condition for multivalence.

Blackburn's example supports (MT) only if, in a context where a college master's job performance is under discussion, to say that Dr. Tweedy is frugal *is to say* that his job performance is deficient. But if this were true, then (5) ought to sound bad or at best awkward in the relevant context.

(5) Dr. Tweedy is frugal. By the way, the claret is abundant and fine.

(5) can be just fine, however. So Blackburn's example fails to support (MT).

There are at least three alternative analyses of how (5) can be fine which are consistent with scale semantics and don't support (MT):

*Implicature analysis:* (5) can be fine if an utterance of 'Dr. Tweedy is frugal' generates a conversational implicature that his performance as college master is deficient. Uttering (5) can cancel this implicature.

*Multiple scales analysis:* (5) can be fine if context sometimes determines the relevant dimension as college expenditure more generally and not as hospitality expenditure in particular, much as *smart* can mean book smart or street smart. (These two scales may be too close for context to select between them, however.)

*Contextual enrichment analysis:* (5) can be fine if Dr. Tweedy's college expenditure is low except on hospitality and context adds this information to the dimension of a single scale of frugality (which can be an ordering of something like money spent). E.g., perhaps Dr. Tweedy travels considerable distance to get a good price on the claret.

**3.4 Summing up.** Scale semantics helps explain why certain data that show up in the literature show nothing much about the semantics of thick evaluative terms. In particular, arguments from examples to (MT) found in the literature fail to support (MT). These arguments show nothing much about whether the valence of thick evaluative terms is semantically determined or constrained.

To be clear, it is easy to find contexts where a sentence using the positive form, such as 'He is brave,' can be used to express an evaluation whose valence is atypically negative. All we require is a context where what we need are cowards so that being brave is not a desirable quality. But this is a general linguistic phenomenon that

has nothing in particular to do with evaluative terms. It alone supports no version of (MT) that concerns the semantics of thick evaluative terms.

All of the above bears not only on (MT) but also on (PT). Blackburn (1992), for example, argues that because thick evaluative terms can be used to make evaluations of different valences in different contexts, their evaluative content makes no contribution to the truth-conditions of thick evaluative claims. It is part of the pragmatics of thick evaluative terms, not part of their semantics.<sup>9</sup>

But if the semantic value of a thick evaluative term is relative to a contextually determined standard of comparison, then variation in valence can have a semantic explanation at least when it is due to variation in the standard value or the scale. But such variation is not automatically due to such factors. So the scale semantics makes the above alternative analyses available but not obligatory. Whether thick evaluative terms are multivalent by their semantics remains so far open.

## 4 Thick evaluative terms and “metasemantics”

It is by no means straightforward to determine when variation in the valence of a thick evaluative term has a semantic explanation in virtue of being due to variation in the standard value or the scale. Scale semantics takes no stand on the content of the scale associated with a gradable adjective. Semanticists are typically uninterested and reluctant to hazard suggestions about what properties in the world provide the dimension for the scale. The semantics requires just that *there be* a total ordering along some dimension associated with a gradable adjective in a context.

To determine the interpretation of a gradable adjective we must typically use context to fix a scale. A scale is set by a computation based on a wide range of inputs from the context. Even fixing a dimension may be part of this computation. The relevant dimension won't in general be as easy as to determine, and as fully

---

<sup>9</sup>Blackburn (1992) may be led to this view because of a crude semantics/pragmatics distinction, which counts pretty much any content as pragmatic unless it is determined independently of context.

determined by the meaning of the adjective, as in cases like *tall* (height), *young* (age), and *cheap* (cost). There may also be no uniform answer as to what fixes the scale. What factors qualify as inputs, what computational rules must be taken into account, and how these must be weighed in fixing the values of contextual parameters like *d* in particular contexts are messy and complex issues in *metasemantics*.

Whereas semantics tells us what semantic properties expressions have, metasemantics tells us how expressions got their semantic properties and why they have them. For example, pronoun resolution rules, which help determine to *which* individual ‘he’ or ‘she’ refers to in a particular context, are metasemantic in this sense. I have no interest in terminological disputes, however. If one has a broader conception of what the semantics of gradable adjectives covers, then issues which I am labeling “metasemantic” may instead be seen as semantic in some broader sense.

What little systematic work there is on metasemantics suggests that potentially relevant factors in determining semantic interpretation include: salient properties of the context; the denotations of the expressions involved; the intentions and interests of speakers and hearers, plus coordinating intentions; the structure of the discourse in which the utterance appears; and what Kennedy (2007: 36) calls the “principle of interpretive economy,” which says that to compute what a sentence expresses in context, make as much use as you can of the (conventional) meanings of expressions.<sup>10</sup>

My proposal is that many disputes about the meaning of thick evaluative terms can be explained as reflecting different metasemantic views about what contextual inputs may enter into computations of standard or scale which determine semantic interpretation of thick evaluative terms. I only have space for two examples.

**4.1 Cognitivism vs. expressivism.** Scale semantics can explain how expressivists and cognitivists can agree on the basic interpretation of words like *generous* in a context. They can agree that it has the sort of meaning set out in (2). They can

---

<sup>10</sup>This list relies largely on Kennedy (2007) and Glanzberg (2007).

also agree that even if we make as much use as we can of its conventional meaning in computing what a thick evaluative term expresses in a context, that meaning may contribute only various non-evaluative constraints on the relevant scale and not fully determine the relevant dimension.<sup>11</sup> Their views diverge in metasemantics.

One way to capture expressivism in this framework is that among the contextual inputs that help determine the relevant scale and standard are *speakers' conative attitudes* regarding the various intervals on the relevant dimension: approval concerning one, disapproval concerning another, indifference concerning something in between, and so on. After all, according to expressivism, the fundamental function of moral judgment is to express conative attitudes towards things on the basis of some features one takes them to have. So understood, expressivism needn't collapse into any sort of subjectivism. Attitudes can help determine semantic values without going *into* semantic values. One way to capture cognitivism is that computations which fix the scale don't take the speaker's conative attitudes as inputs.<sup>12</sup>

This way of capturing cognitivism and expressivism enables us to see how their debate might proceed. Consider a multi-dimensional gradable like *courageous*. Cognitivists might claim that a scale can be computed from a set of dimensions determined by the adjective's meaning and context plus a rule of interpretation that links changes of location in this multidimensional space to an ordering. Perhaps there is a function that takes us from inputs such as credence that one will sustain an injury

---

<sup>11</sup>Note here a related issue that we can now construe as having a metasemantic dimension. Some writers claim that thick evaluative terms are associated with fairly sharp non-evaluative properties (see e.g. Blackburn 1992). Others claim that at least the sort of characterizations that we can usually give merely specify a domain of application for a thick term without supplying a "self-standing descriptive meaning" for it (see e.g. Williams 1985: 141, Gibbard 1992: 276-77, Dancy 1995: 277). So far as the semantics cares, the scale of frugality, the scale of lewdness, etc., in a context might be an ordering along a dimension which we can adequately represent in purely non-evaluative, descriptive terms. But, equally well, it might not.

<sup>12</sup>It may be helpful to note that the "standard" for the adjective and context need not be any subject's moral standard. Semantics requires only that it be some sort of standard of comparison.

if one does something, magnitude of the said injury, credence that one will achieve a gain if one does it, and magnitude of the said gain, to an ordering which gives us something we can treat as a scale of courageousness.

Cognitivists could here argue that general requirements on semantic interpretation favor a computational rule of their preferred kind. One such requirement is that the standard of comparison should always provide a non-trivial partitioning of the domain (Klein 1980: 23). A rule that ignores speakers' cognitive attitudes might well ensure this result. Ensuring it might be harder for expressivists. Some speakers' conative attitudes might well be such that in some contexts all (or no) options count as satisfying *courageous*. Disputes about such general interpretive requirements and their implications are metasemantic.

**4.2 Role of evaluative practice.** Some writers claim that one cannot “master the extension” of a thick evaluative term, and reliably apply it to new cases in the same way as the community in which the term is used, unless one grasps their evaluative outlook or the term's evaluative point.<sup>13</sup> One example given by these writers is the word *chaste* as used by Victorians. How could it be that grasping what they would count as satisfying *chaste* requires grasping their evaluative outlook?

One hypothesis is that what Victorians would count as satisfying *chaste* is determined partly by *presuppositions that are widely shared* in discourse among Victorians about sexual behavior.<sup>14</sup> The idea would be that when a speaker utters a sentence like ‘Emma is chaste’ in a discourse so structured, intentions that conflict with such widely shared presuppositions might easily fail to set values to contextual parameters like *d* (see Glanzberg 2007: 26 on *rich*). This hypothesis is metasemantic.

We can again see how the debate might proceed. One might claim that the hypothesis that widely shared presuppositions can help set contextual parameters like the degree value helps explain other important data, such as the possibility of

---

<sup>13</sup>See e.g. McDowell (1981), Williams (1985), and Gibbard (1992).

<sup>14</sup>But only partly. The conventional meaning of the term, even when enriched by this factor, may not fully determine the relevant scale or even dimension. Cf. Gibbard (1992: 274-77).

refusing to use certain thick terms because one finds them objectionable. One way in which a term can be objectionable is that claims made using it carry presuppositions which one rejects. Perhaps one thinks that *chaste* as used by Victorians reflects outdated or otherwise objectionable views about appropriate sexual behavior. Such a speaker will think that their use of *chaste* fails to pick out a genuine normative kind. But if the degree value which determines what would count as satisfying *chaste* in a context is set partly by presuppositions the speaker rejects, then she wouldn't appropriately assert either 'Emma is chaste' or its negation. So this hypothesis can seemingly explain why one may refuse to call anything chaste or not. Of course, many fascinating but complex issues arise here, such as the complications discussed by Gibbard (2003: 142-46, 168-78) and whether context can fail to determine a scale altogether when pragmatic presuppositions fail. But we can already see that casting the debate as involving metasemantic issues promises to be fruitful.

## 5 Conclusion

I first argued that recent work on the semantics of gradable adjectives helps explain why certain arguments concerning the meaning of thick evaluative terms are highly questionable. I then proposed that we can make sense of certain disputes concerning their meaning if we construe them as reflecting disagreements concerning what sort of contextual factors can contribute to setting the values of such contextual parameters as the standard of comparison for adjective in a context. So far this proposal remains highly preliminary. The various metasemantic suggestions introduced in section 4 to begin making certain disputes about thick evaluative terms more tractable probably raise more questions than they settle. But I hope that the proposal may already now be seen to have various advantages. Here are a few.

One advantage is that the proposal helps to identify constraints and force theoretical unity on the semantic interpretation of thick evaluative terms. In particular, proposed contextual inputs into their interpretation ought to be acceptable inso-

far as they are of the sort that are relevant to determining the semantic values of gradable adjectives in general; otherwise, we require some independent support for thinking that they can be parameter-setting contextual factors.

Another advantage is that the proposal points to ways to assess whether pejorative terms, which are increasingly studied by linguists and philosophers or language, offer a good model for thick evaluative terms. Even if they don't,<sup>15</sup> consideration of what makes pejorative terms objectionable might offer some help in identifying contextual inputs into the interpretation of objectionable thick terms.

One implication of the proposal is that it gives grounds for doubting that the truth or falsity of claims like (DT) has a purely semantic basis or can be settled by *a priori* reflection on the meanings of the expressions involved. It gives similar grounds for doubting those specific versions of claims like (DT) which assert that the relationships they concern hold analytically or in virtue of semantics. This is simply because those claims concern issues that are at least partly metasemantic.

A closer attention to metasemantic phenomena may also help determine, further, which phenomena associated with thick evaluative terms are features of their metaphysics or epistemology rather than semantics or metasemantics. For instance, I conjecture that (meta)semantic considerations won't help us determine whether the functions that are the semantic values of thick evaluative terms correspond with any (sparse) properties. So scale semantics may in this way be exploited to clarify various issues concerning thick evaluative terms.

One final point: scale semantics more generally implies that the basic semantics of gradable "thin" and "thick" evaluative terms will exhibit no deep difference in kind. For instance, the comparative and positive forms of *cruel* and *good* will receive the same basic semantic treatment, along the lines of (1) and (2). This will hold whether the evaluative terms in question are moral, epistemic, or aesthetic.

---

<sup>15</sup>Standard pejoratives (such as racial epithets and derogatory terms like *Boche*) are not gradable expressions and typically their derogatory content is part of their dictionary entry. Their interpretation may also make much lighter demands on context and metasemantics.

## References

- Blackburn, Simon (1992). “Through Thick and Thin.” *Proceedings of the Aristotelian Society* Supplementary Volume 66: 284-99.
- Dancy, Jonathan (1995). “In Defence of Thick Concepts.” *Midwest Studies in Philosophy* 20: 263-79.
- Eklund, Matti (unpublished). “Some Problems Regarding Thick Concepts.”
- Gibbard, Allan (1992). “Thick Concepts and Warrant for Feelings.” *Proceedings of the Aristotelian Society* Supplementary Volume 66: 267-83.
- Gibbard, Allan (2003). *Thinking How To Live*. Harvard UP.
- Glanzberg, Michael (2007). “Context, Content, and Relativism.” *Philosophical Studies* 136: 1-29.
- Hare, R. M. (1952). *The Language of Morals*. Clarendon Press.
- (1981). *Moral Thinking*. Clarendon Press.
- Kennedy, Christopher (2007). “Vagueness and Grammar: The Semantics of Relative and Absolute Adjectives.” *Linguistics and Philosophy* 30: 1-45.
- Klein, Ewan (1980). “A Semantics for Positive and Comparative Adjectives.” *Linguistics and Philosophy* 4: 1-45.
- McDowell, John (1981). “Non-Cognitivism and Rule-Following.” Reprinted in his *Meaning, Value, and Reality*. Harvard UP, 1998.
- Meier, Cécile (2003). “The Meaning of *Too*, *Enough*, and *So . . . That*.” *Natural Language Semantics* 11: 69-107.
- Schwarzschild, Roger and Wilkinson, Karina (2002). “Quantifiers in Comparatives: A Semantics of Degree Based on Intervals.” *Natural Language Semantics* 10: 1-41.
- Williams, Bernard (1985). *Ethics and the Limits of Philosophy*. Harvard UP.