



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/76775/>

Version: Published Version

---

**Article:**

Vayrynen, PP (2011) Thick Concepts and Variability. *Philosophers' Imprint*, 11 (1). 1 - 17.  
ISSN: 1533-628X

---

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# Thick Concepts and Variability

Pekka Väyrynen

*University of Leeds*

© 2011 Pekka Väyrynen

<[www.philosophersimprint.org/011001/](http://www.philosophersimprint.org/011001/)>

## 1. Introduction

The distinction between “thick” and “thin” terms and concepts and its philosophical significance has been an active topic in ethics, and increasingly in aesthetics and epistemology. Concepts commonly regarded as thick in ethics include *GENEROUS*, *COURAGEOUS*, *CRUEL*, and *GREEDY*, those commonly regarded as thin include many concepts expressed by terms like *good* and *bad* or *right* and *wrong*. What is supposed to be distinctive about thick terms and concepts relative to the thin is that they somehow “hold together” evaluation and non-evaluative description, whereas thin terms and concepts are somehow more purely evaluative.<sup>1</sup> I can convey a negative evaluation of a bully’s actions by calling them cruel or by calling them bad. Such judgments differ with respect to description. Even if taking enjoyment in causing others to suffer is both cruel and bad, the term *cruel* seems to entail by its very meaning that things falling under it to have some non-evaluative features of this general sort, whereas *bad* doesn’t. Some bad things involve no enjoyment in the suffering of others. This contrast will remain even if these non-evaluative features aren’t fully specified in an analysis of a thick term or concept (it might not be a purely conceptual matter just which preferences for one’s own happiness over a greater happiness for others count as selfish) but the meaning of the term or concept merely restricts them to a certain general type (*kind* doesn’t admit of people who are disposed to respond to others in cold or exploitative ways).

The distinction between thick and thin terms and concepts raises important issues of formulation.<sup>2</sup> This paper largely brackets those is-

1. Or normative, insofar as there is any clear distinction. It will be harmless here to include the normative under the evaluative. I’ll use small caps to denote concepts and italics to denote words and sentences.
2. One issue is how to explicate the contrast. Existing accounts, such as Williams (1985), Gibbard (1992), or Dancy (1995), don’t adequately distinguish the thick from the thin in anything like a theoretically non-committal way (see Eklund forthcoming). Another issue is the status of the contrast. The distinction may mark some kind of difference of degree along a spectrum of concepts rather than a binary distinction in kind (Scheffler 1987:

sues. I'll simply assume that thick terms and concepts bear conceptual, analytic, or semantic connections to non-evaluative features of at least a certain general type: giving of oneself in the case of *generous*, taking enjoyment in causing others to suffer in the case of *cruel*, and so on. My interest in this paper lies rather in the relationship between thick terms and concepts and evaluation. No doubt there is an important sense in which the bully wouldn't understand what the angry parents of his victims try to communicate when they say that bullying is cruel, if he failed to grasp the kind of negative evaluation which they mean to convey. Just what kind of failure this would manifest depends on how thick terms and concepts are related to evaluation.

This question about the relationship between thick concepts and evaluation may be formulated as follows:

**Location of Evaluation (LE):** How are thick terms and concepts related to the evaluations they may be used to convey?

By 'evaluations' I mean information or contents that are evaluative, not mental acts of evaluation. LE thus asks how calling something generous or cruel, for instance, conveys evaluative information when it does so. I also assume that evaluation is somehow positive or negative in flavor. One, possibly controversial, way to explicate this is to say that evaluative information concerns the merit or worth of something or someone against some kind of standards.<sup>3</sup> Things that compare favorably relative to a particular standard are good in a certain way, and bad in a certain way if they compare unfavorably. I'll use 'convey' as an umbrella term such different modes of information transfer as semantic content and entailment, implicature, and presupposition.

The standard answer to LE is that thick terms and concepts are

inherently evaluative in their very meaning. Their relationship to evaluation is conceptual, analytic, or semantic in the same sort of way as their relationship to description.<sup>4</sup> Some evaluation, such as that things falling under *generous* are good in some way related to giving of oneself, belongs to the very meaning of *generous* as much as do such general type descriptions as that those things involve giving of oneself.<sup>5</sup> The natural alternative is to say that evaluation isn't a feature of what thick terms and concepts mean, but rather a feature of what thinkers or speaker-hearers may use them to mean, may associate with their meanings, glean from their uses, and so on.<sup>6</sup> (When 'meaning' occurs by itself in what follows, it refers specifically to semantic meaning, not speaker meaning.) On the standard view the bully we imagined is defective already in his competence with respect to meaning or conceptual content. (Perhaps what happens is that he fails to grasp the sense under which more virtuous people see cruel things or what the truth of the sentence *Bullying is cruel* requires.) On the alternative view, the bully is morally flawed but his mistake needn't involve conceptual confusion or semantic error.

The aim of this paper is not to answer LE either way. Its aim is rather to show that, contrary to what many philosophers have thought, LE cannot be settled on the basis of whether thick terms and concepts are contextually variable with respect to the valence of the evaluations they may be used to convey. (For brevity, I'll write about variability 'in evaluative valence' or 'in the direction of evaluation' from now on.) In what follows, I'll first describe the relevant sort of variability more clearly to indicate how it has been thought to bear on LE. I'll then consider two kinds of examples which have been used to support an argument from contextual variability against semantic views about the

417). Even such paradigmatic thin concepts as those expressed by *ought* in moral contexts may be thick to some degree, insofar as *S ought to φ* entails the non-evaluative modal claim *S can φ*.

3. Part of this thought derives from Williams (1985: 125). For a more careful stage setting in this and other respects, see Väyrynen (ms1).

4. See e.g. Foot (1958), Platts (1979: Ch. 10), McDowell (1979; 1981), Williams (1985), Dancy (1995), Putnam (2002: Ch. 2), and Elstein and Hurka (2009).

5. A related view is that evaluation is some such further semantic property of thick terms and concepts as conventional implicature.

6. See e.g. Hare (1952; 1981), Blackburn (1992; 1998), and Väyrynen (2009).

relationship between thick terms and concepts and evaluation. I'll argue that the first kind of examples fail to support the sort of contextual variability which this argument requires. And I'll argue that the second kind of examples fail to settle LE even if they support the relevant sort of variability, because several different hypotheses regarding LE can account for them. The conclusion I draw is therefore a kind of null result: attempts to settle LE on the basis of considerations of contextual variability in evaluative valence are bound to fail and LE must therefore be settled on some other grounds instead. I'll close by suggesting a more promising line of future inquiry.

## 2. The Variability Argument

What does it mean to say that thick terms and concepts are contextually variable in evaluative valence? A theoretically fairly neutral but (therefore?) rough formulation is that *brutal*, for instance, is variable in valence if it can be used to evaluate things negatively in some contexts but positively in others – if utterances of sentences of the form *X is brutal* convey sometimes that *X* is bad in some way but at other times that *X* is good in some way.<sup>7</sup> This formulation allows that thick terms and concepts may typically convey evaluations of a particular valence even if their valence is contextually variable. Calling something *cruel* or *insensitive* typically conveys something negative, calling something *generous* or *wise* typically conveys something positive.

Apparent contextual variability in evaluative valence has mainly been used to argue against the view that thick terms and concepts are inherently evaluative in meaning. Such views require that sentences or propositions involving thick terms and concepts analytically entail evaluation. The main idea behind arguments from contextual variability is that one should find it hard to see how that could be so if the evaluations that thick terms and concepts may be used to convey were

7. The latter might be a natural interpretation of an utterance that pleads "Don't change this beautiful and brutal game." (Thanks to Remy Debes for this example from a discussion of concussion risks in American football.)

positive in some contexts but negative or even neutral in others.<sup>8</sup>

The main support for positing variability in evaluative valence comes from various sorts of examples. These examples, as we'll see, move at the level of language. Thus arguments built on them presume that linguistic evidence concerning thick terms bears on the nature of thick concepts. I'll understand concepts in the common way in terms of their theoretical function to explain phenomena of cognitive significance. (In one familiar terminology, this locates concepts closer to the realm of sense or modes of representation rather than reference. We also know that linguistic meaning isn't to be equated with reference.) It seems to be fairly common in the literature on thick concepts to assume that the meanings of thick terms are more or less the same as the concepts they express. This hides potential complications insofar as the relationship between language and concepts is complex and controversial.<sup>9</sup> But all that I'll need to assume about linguistic meaning is that aspects of the meaning of an expression must be stable at least across its literal uses in normal contexts. Such constraints on uses across contexts also seem relevant to individuating concepts; one would expect them at least to constrain the concepts that thick terms may be interpreted as expressing. I'll continue to speak of thick terms and concepts in one breath with this contact point in mind.

Variability arguments of the kind described above can be broken down to two main premises. One states the claim of contextual variability, typically supported by examples. The other makes a claim about how contextual variability bears on the relationship between thick terms and concepts and evaluation. The following "Variability Argument" is an argument of this general form:

(V1) *Variability Claim*: Thick terms and concepts are contextually

8. Compare standard examples of analytic truths: it cannot be that a bachelor sometimes is eligible to marry but sometimes isn't, or that being red all over sometimes excludes being green all over but sometimes doesn't.
9. The relationship between language and concepts may not be straightforward if, for instance, concepts are more fine-grained than the sorts of entities that semantics assigns as the meanings of expressions.

variable in evaluative valence.

- (V<sub>2</sub>) *Link Claim*: If thick terms and concepts are contextually variable in evaluative valence, then the evaluations that they may be used to convey don't belong to their meanings.
- (V<sub>3</sub>) Therefore, the evaluations that thick terms and concepts may be used to convey don't belong to their meanings; they are rather a feature of pragmatics.

Both premises of this argument require clarification.

The Link Claim in (V<sub>2</sub>) is to be understood as sufficiently generic with respect to 'meaning' to rule out any view according to which thick terms and concepts are inherently evaluative. Different versions of such views will render the notion more precise beyond the constraints already indicated. (One might claim that the relevant meanings are senses or some other (perhaps less fine-grained) entities that determine semantic value. Or one might claim that they are some such further semantic properties as conventional implicatures.) Thus one way to resist the Variability Argument is to agree that thick terms and concepts display genuine contextual variability in evaluative valence but argue that the relationship between thick terms and concepts and evaluation could be semantic and yet allow for such variability.<sup>10</sup>

The Variability Claim in (V<sub>1</sub>) is to be understood as allowing that many thick terms and concepts are typically used to convey evaluations of a particular valence. It is one thing to acknowledge that this is how thick terms and concepts typically work, quite another to infer that evaluations with that (or any) valence belong to its meaning. Such regularity might be enforced by something other than meaning.<sup>11</sup>

10. Dancy (1995) advances an inseparabilist (see below) version of such a view. But as we'll see in §5.3, this option is more general. As we'll see in §5.1, there also seem to be ways of arguing from (V<sub>1</sub>) to (V<sub>3</sub>) which don't require (V<sub>2</sub>). But at least Blackburn (1992) accepts (V<sub>2</sub>).

11. It is, similarly, one thing to acknowledge that *cruel* or *selfish*, for instance, are typically used to convey a negative evaluation even if the speaker doesn't intend this, quite another to infer that such evaluations belong to their meaning. Pragmatic meanings can be "generalized" to the extent that they are part of a defeasible default interpretation.

(V<sub>1</sub>) is also to be understood as saying that if *cruel*, for instance, is typically used to convey negative evaluations through whatever mechanism, then it can sometimes be used to convey positive evaluations through the same kind of mechanism. This qualification is meant to rule out cases where a thick term or concept is used to convey evaluations whose valence differs from the typical only thanks to non-literal use or conveys no evaluation because it is used non-evaluatively merely to classify things. (A historian might report that the cruel and unjust practices of a regime provoked rebellion without thereby evaluating the regime.)

These two qualifications to (V<sub>1</sub>) are related. Some evaluations whose valence differs from the typical seem to move at a different level from the typical. Suppose, for instance, that *deceitful* is typically used to convey that something is bad in some way but can be used to convey that it is good in some way. At least sometimes such a positive evaluation seems to be a mere conversational implicature to the effect that deceitfulness is good for helping secure what one wants, or winning a reality TV show, or the like. If this were the only sort of mechanism for using *deceitful* to convey positive evaluation and if *deceitful* conveyed negative evaluation through some different mechanism (perhaps in virtue of its meaning), then the evaluations would move at different levels in the relevant sense. The notion of contextual variability in (V<sub>1</sub>) is to be understood as ruling out such important asymmetries between negatively and positively evaluative uses of thick terms and concepts. Thus another way to resist the Variability Argument is to argue that putative examples of contextual variability in fact involve such an asymmetry or otherwise fail to speak to EQ. The bulk of this paper is devoted to discussing two sorts of examples of variability in valence which figure in the literature on thick concepts as support for the Variability Argument.

Variability in direction of evaluation is often invoked in debates between "separabilists" and "inseparabilists" which occupy much of the literature. This debate concerns whether thick terms and concepts are some kind of irreducible fusions of evaluation and description or

whether these two aspects can be separated or “disentangled” in some way.<sup>12</sup> Broadly pragmatic views on the relationship between thick terms and concepts and evaluation suggest that some kind of separability will be possible, since they treat evaluation and description as bearing different relations to thick terms and concepts. Thus it is no accident that the Variability Argument is meant to apply against the sorts of inseparabilist views which one typically finds in the literature. Presumably the thought is that it would be difficult to explain how thick terms and concepts could be contextually variable in evaluative valence if evaluation and description were inseparable.

The Variability Argument is more general, however. However well it works against inseparabilists, it should work equally well (or badly) against those separabilists who think that thick terms and concepts have a semantically fixed direction of evaluation which is separable from description. So, while the Variability Argument is geared at LE (since it targets the hypothesis that thick terms and concepts are inherently evaluative in meaning), the issue of separability lies downstream from LE (because such views could be either separabilist or inseparabilist).<sup>13</sup> The issue regarding (V2) whether thick terms and concepts could exhibit genuine contextual variability if they were inherently evaluative in meaning can similarly be pursued along either separabilist or inseparabilist lines.<sup>14</sup> Thus, although the issue of separability has some relevance to the Variability Argument, it doesn’t provide the most general conceptual frame for assessing how exactly that argument bears on LE.

12. Separabilists include Hare (1952: 121-2; 1981: 17-8, 73-5), Blackburn (1992; 1998: 101ff.), and Elstein and Hurka (2009). Inseparabilists include McDowell (1979; 1981), Platts (1979: Ch. 10), Williams (1985: 129-30, 140-1), Dancy (1995), Wiggins (1998), Putnam (2002: Ch. 2), and Kirchin (2010). As the metaphorical character of how these views are characterized perhaps suggests, the semantic, metaphysical, and epistemic properties of this relation of (in)separability are often unclear.

13. Whether pragmatic views entail separabilism depends on such general issues as whether there is “pragmatic intrusion” into semantic content.

14. See notes 25 and 27 below.

### 3. Variability and Comparative Constructions

One argument for the Variability Claim is based on examples like the following:

- (1) ‘[T]idy’ and ‘industrious’ . . . are normally used to commend; but we can say, without any hint of irony, ‘too tidy’ or ‘too industrious.’ (Hare 1952: 121.)
- (2) [We may] worry that this year’s Carnival was not lewd enough. (Blackburn 1992: 296.)

It is clear from context that these examples are meant to support something like the Variability Claim introduced above.<sup>15</sup> Hare explicitly takes his example to be a case of literal use to convey negative evaluation. Blackburn takes his example to be a case of literal use to convey that the carnival would have been better had it been more lewd, so that *lewd* conveys positive evaluation instead of the typical negative one. But what is the argument? I’ll now argue that two natural routes from examples involving such constructions as *too* and *not . . . enough* to the Variability Claim won’t work.

One option is to claim that the evaluative valence of *too tidy* and *not lewd enough* is determined in some way that requires *tidy* and *lewd* to have contextually variable valence as well. But in fact the way modifiers like *too* and *not . . . enough* work in general explains how the former constructions may flip the evaluative valence of their unmodified roots without the latter having to be variable in valence.

Constructions of the forms *too F* and *not F enough* are comparative. There is broad consensus that determining what they say requires some kind of standard of comparison. Such a standard can be either set implicitly by the context of utterances of sentences like (3) and (4) or fixed explicitly by devices which can serve to supply a standard. Examples include the *to-* and *for-*phrases in (3a-b) and (4a-b):

15. Thick terms and concepts are often assumed to be sufficiently uniform for examples to generalize across the board. Such uniformity assumptions are dubious, but I’ll bracket the issue here.

- (3) Michael is too industrious.
- a. Michael is too industrious to have time to visit Bordeaux wineries.
  - b. Michael is too industrious for me to read all his work.
- (4) The party isn't loud enough.
- a. The party isn't loud enough to keep the neighbors awake.
  - b. The party isn't loud enough for Simon.

Some people think that utterances of sentences like (3) and (4) can express complete propositions without any implicit or explicit standard being supplied. This would effectively be to claim that one can be too industrious, period, or not loud enough, period. But if there is anything to say about what kind of criticism such an utterance is meant to express, then some kind of standard will be at least implicit in the answer. In the case of *industrious*, the standard might concern, for instance, the shape of a balanced life, the time and attention due to one's family, or the like. The standard needn't be a moral or prudential one, though. Semantics will admit any respect in which degrees of a quality such as industriousness may be compared. (It will also admit different views on what sort of property is thereby measured.)

How these comparative constructions work could be analyzed in great detail by drawing on the standard degree-based semantics for gradable expressions.<sup>16</sup> But some basic points will suffice here. The standard of comparison determines upper and lower bounds of admissibility on a scale associated with the word, such as the scale of industriousness in (3) or the scale of loudness in (4). Depending on the polarity of the modified expression, *too F* places an F either above the upper bound of the interval (e.g., *too expensive* on the dimension of cost) or below its lower bound (e.g., *too young on the dimension of age*), and analogously for *not F enough* (consider *not cheap enough* and *not old*

*enough*). (3) and (4) can thus be analyzed as follows:

- (3') The value Michael takes on a scale of industriousness is greater than the maximum admissible value determined by context *c*.
- (4') The value the party takes on a scale of loudness is lower than the minimum admissible value determined by context *c*.

So when (3) means (3a), it says something like this: Michael's degree of industriousness is greater than the maximum degree *d* such that if Michael's degree of industriousness is (no greater than) *d*, then he has time to visit Bordeaux wineries.

What the examples under (3) and (4) show is that the standard for counting as satisfying *too F* or *not F enough* is typically neither the same as the standard for satisfying *F* nor determined by the same factors. A party might count as satisfying *loud*, but be less noisy than would be required to keep the neighbors awake or to keep Simon around (*not loud enough*), and yet be more noisy than is safe for hearing or comfortable for infants (*too loud*). This is a perfectly general point about how these modifiers work. There is no reason why the standards of comparison they invoke should have any implications for whether thick terms and concepts have variable valence in unmodified form.

Another option is to claim that if *tidy* were inherently a term of positive evaluation, or *lewd* inherently a negative one, criticizing something as *too tidy* or *not lewd enough* shouldn't be semantically permissible. More of a positive quality, or less of a negative one, should hardly be a bad thing. Since it seems perfectly coherent to criticize something by calling it *too tidy* or *not lewd enough*, *tidy* and *lewd* should be regarded as contextually variable rather than semantically fixed in their evaluative valence.<sup>17</sup> Another way of putting this idea is that if counting as *tidy* were inherently good in some way and counting as *lewd* were inherently negative in some way, then it would be difficult to explain

16. See e.g. Creswell (1977), Klein (1980), von Stechow (1984), Schwarzschild and Wilkinson (2002), Meier (2003), Kennedy (2007), and Glanzberg (2007).

17. This argument was suggested to me (but not endorsed) by Daniel Elstein (p.c.). Blackburn (1992: 286) may have something similar in mind.

how *too tidy* and *not lewd enough* could invert the evaluative valence of their unmodified roots as they do.

This argument also fails to appreciate how modifiers like *too* and *not ... enough* work. Calling something *too tidy* or *not lewd enough* can be semantically permissible even if tidiness is an inherently positive quality or lewdness an inherently negative one, because *too F* and *not F enough* take a standard of comparison which, as we saw, is typically neither the same as the standard for satisfying *F* nor determined by the same factors. Something can count as *too F* relative to some contextually determined purpose or standard even if more of *Fness* is better so far as the standard for counting as *F* is concerned. A military commander can perfectly well count a soldier as *too courageous* for a camouflage mission for which cowards are suitable and select him for the real mission because he is very courageous indeed.

It is also worth noting that if the argument works, it may apply more broadly than one might have thought. Thin terms and concepts seem to work in the same way in the relevant comparative constructions so far as semantic permissibility goes. A wine can be said to be *too good* to be used for cooking (or dyeing clothes, or whatever) even if the degree of goodness which makes it so relative to that standard also makes it better as a wine or makes drinking it a better state of affairs. Satan in a career slump can regard his activities as *not (morally) bad enough* for a satisfactorily Satanic job performance even if those activities still clearly qualify as (morally) bad.<sup>18</sup> In general, the standard for counting as satisfying *too good* or *not bad enough* is also typically neither the same as the standard for satisfying *good* or *bad* nor determined by the same factors. But if the above argument then generalized to show that *good* and *bad* also lack a semantically fixed direction of evaluation, it couldn't be used to show that thick terms and concepts

18. Whether the distinction between "attributive" and "predicative" uses of *good* complicates examples in this vein is unclear. Some argue that (surface grammar aside) *good* is uniformly predicative but context-sensitive with respect to a standard (Szabó 2001). Nor does that distinction undermine the similarity noted next in the text between thick and thin terms.

in particular are variable in valence. The possibility that thick and thin terms and concepts might *both* be contextually variable in valence is tangential to my present aims. Such parity claims would in any case require independent support beyond the examples of apparent variability presently on the table.

It is also no accident that those who defend the Variability Claim with examples like (1) and (2) never bring up thin counterparts of such examples. Many philosophers wouldn't or (given their other views) couldn't welcome the result that thin terms and concepts are contextually variable in their evaluative valence.<sup>19</sup>

This explanation of why the two arguments from *too* and *not ... enough* to the Variability Claim fail also handles examples like (5):

- (5) I found this year's carnival lewd – just think of all those people, almost naked, dancing in that sexually suggestive way on the main street. And yet the carnival wasn't lewd enough.

An utterance of (5) needn't be contradictory. Even prudes can accept that an occasional opportunity for a socially permissible transgression of conventional boundaries on sexual display has a valuable social function to serve. As Blackburn puts it, carnivals can serve this kind of social function as "necessary eruptions of the Dionysiac into the fragile Apollonian order" (1992: 296). An utterance of (5) can be perfectly appropriate in certain contexts of utterance. A carnival can count as lewd, as the first half of (5) registers, but not sufficiently so for it to satisfy this social function, which is what the second half of (5) can express in a suitable context. This is fully explicable in terms of how the relevant modifiers work in general.

I conclude that the Variability Claim cannot be established on the basis of how thick terms and concepts behave under modifiers like *too*

19. In particular, Hare cannot allow that *good* is contextually variable in its use as what he calls a "primarily evaluative" term (Hare 1952: Ch. 7). And Blackburn cannot allow that *good* varies with respect to whether it expresses pro-attitudes or con-attitudes, since these play different roles in guiding desires and choices (Blackburn 1993: 137, 168).

and *not ... enough*. But it would be natural to wonder whether these comparative constructions are required to get examples of apparent variability in valence going. I'll now turn to discuss whether the Variability Claim could be supported by examples involving thick terms and concepts in unmodified positive form.

#### 4. Variability in the Positive Form?

Another argument for the Variability Claim is based on examples where a simple predication of a thick term or concept appears to be used to convey an evaluation whose valence differs from those it is typically used to convey:

- (6) The carnival was a lot of fun. But something was missing. It just wasn't lewd. I hope it'll be lewd next year. (Eklund, p.c.)
- (7) A previous Master of my Oxford College was condemned as frugal, a fatal flaw in someone whose main job is dispensing hospitality. (Blackburn 1992: 286.)
- (8) Wiping a sweaty torturer's brow can be considerate. But it might not be the better for it, and hardly functions as a reason to wipe. (Dancy 2009: §6.)<sup>20</sup>
- (9) We can say 'Yes, cruel certainly, but that's just what made it such fun.' (Hare 1981: 73.)

These examples raise two sorts of issues. One is whether the evaluations that the utterances they report are used to convey are correct. The second is whether those utterances are semantically permissible. These issues are distinct. It might, for instance, be mistaken to deny that considerate things are always better for it in some way, at least to some small degree. Such denials might nonetheless be semantically permissible. What matters to LE are the latter sort of facts.

There is no obvious need for anyone to deny that thick terms and concepts can sometimes be used in unmodified form somehow to convey evaluations whose valence differs from the typical. What matters

20. For other examples in this vein, see Swanton (2001: 48).

to LE are the mechanisms by which those evaluations would be conveyed. Thus uttering (6) can somehow convey that next year's carnival would be better in some way if it were lewd. (7) reports that sometimes – as in the case of Blackburn's college master, call him 'Tweedy' – *frugal* can be used somehow to convey criticism rather than praise. (8) claims that *considerate* can be used without conveying anything positive or even to convey negative evaluation (e.g., to criticize someone for facilitating a torturer's evil purpose). And (9) can sometimes be read as agreeing that an action counts as satisfying cruel but take that as recommending it, perhaps even without any obvious tint of sadism. These examples support the Variability Argument only insofar as they involve using thick terms or concept to convey one direction of evaluation by the same sort of mechanisms as they are (typically) used to convey a different direction of evaluation.

A satisfactory account of examples like these should account also for certain variants, such as the following sorts of variants of (7):

- (10) A: Tweedy is frugal.  
B: Yes, but the claret is abundant and fine.
- (11) C: Tweedy is frugal.  
D: Yes, the claret is worse than before and there isn't enough of it for us to stop caring.

Sometimes when A's utterance in (10) occurs in a discussion of Tweedy's job performance, B's reply should be able to sound fine. If A is conveying a negative evaluation of Tweedy's hospitality in calling him frugal, sometimes B's reply might be a polite expression of a wholesale disagreement. (This is one possible conversational effect of *Yes, but...*) But sometimes B might be objecting to A's criticism of hospitality while granting that A is saying something true. (This is another possible conversational effect of *Yes, but...*) For instance, A and B might agree that Tweedy generally manifests frugality in his duties as college master but disagree on whether he shows inadequate hospitality. (11) doesn't allow the same range of possibilities. When C's utterance occurs in a discussion of Tweedy's job performance, D's re-

ply can also sound fine. But D's reply expresses a negative evaluation of Tweedy's dispensation of hospitality and thus agrees with C's criticism of Tweedy's job performance. Certain instances of (10) and (11) may thus exhibit a certain kind of asymmetry.

## 5. Explanations of Contextual Variability

I'll now argue that there are many available explanations of the apparent contextual variability in evaluative valence which thick terms and concepts can manifest in their positive form. Some of these views have so far been represented in the literature on thick concepts inadequately or not at all. My discussion will therefore introduce as a side effect some new theoretical options into the literature. These explanations have different implications regarding LE. No doubt some will turn out to be more plausible than others. But the relative merits of these explanations will be largely independent of their implications regarding contextual variability. Nothing much about LE will be settled by those considerations. Or so I'll argue.

### 5.1 Pragmatic Views

The explanations most commonly favored by those who offer examples like (6)-(9) locate the evaluations that thick terms and concepts may be used to convey outside their conventional meaning as some aspect of the information that speaker-hearers may use them to express, associate with their meanings, and glean from their uses, even in literal uses in normal contexts. That is no surprise: the Variability Argument is meant to support the conclusion that the relationship between thick terms and concepts and evaluation is broadly pragmatic. This would be to say that thick terms and concepts are evaluative only in the weak sense that they can be used evaluatively.<sup>21</sup>

21. This is true of many terms that aren't evaluative in meaning. Calling something *physically strenuous*, for instance, can convey a positive evaluation in some contexts but a negative evaluation in others by means of conversational implicature. Terms such as *selfish* and *generous* might be mere special cases where evaluations of a particular direction of valence are sufficiently

Regarding (6), a pragmatic hypothesis has resources to explain how general conversational principles can in some contexts suggest interpreting the speaker as implying or suggesting that an overt display of sexuality would have been the better for it, not the worse for it, had it transgressed conventional boundaries on sexual display. Regarding (7), it can allow that whether *frugal* is used to convey positive or negative evaluation can depend on context in subtle ways. For instance, it doesn't require that any literal use of *frugal* within any discussion of college masters or their job performance will convey negative evaluation, but only that some such contexts work this way, these perhaps being contexts where commonly shared background assumptions favoring generosity in hospitality are salient. (10) and (11) might be thought to illustrate one potential difference between contexts. If it is manifest in the context that Tweedy's performance as master might be deficient in some way related to low expenditure, this could explain how A's utterance in (10) or C's utterance in (11) can suggest that this is so, and why B's reply could be heard as denying this suggestion but D's reply be heard as reinforcing it. Regarding (8) and (9), if *considerate* and *cruel* aren't inherently evaluative in meaning, then there is no deep theoretical problem with understanding how *considerate* could sometimes be used literally in a way that suspends the positive evaluations which are commonly associated with it or even conveys negative evaluation, or how *cruel* could in some suitable contexts be used literally to convey positive evaluation. All that is required is that a suitable discourse situation fall into place.

Pragmatic views of this sort support the Variability Claim. They are consistent with the Link Claim, but in the foregoing a pragmatic explanation of what is going on in examples like (6)-(11) was run without the Link Claim.<sup>22</sup> Thus pragmatic views on the relationship between thick terms and concepts and evaluation seem to require only some

deeply entrenched in social norms and cultural context as to be part of their default interpretation. See Väyrynen (ms1) for discussion.

22. Blackburn (1992) endorses the Link Claim. I would argue that his case for it relies on a crude view of the semantics/pragmatics interface.

weaker assumptions about how contextual variability in evaluative valence bears on that relationship.

### 5.2 Semantic Invariantism

At the other end of the spectrum lie explanations of examples like (6)-(11) which claim both that the relationship between thick terms and concepts and evaluation is semantic and that this relationship involves a fixed direction of evaluation in the literal uses of thick terms and concepts in normal contexts (Link Claim). This kind of semantic invariantism explains away the apparent variability of thick terms and concepts in evaluative valence as non-literal use or abnormal context. Its strategy is to challenge the Variability Claim.

The natural way to explain away the appearance of contextual variability is to invoke the distinction between semantic meaning and speaker meaning, note that these may come apart, and argue that they come apart in our examples in a different way than they do according to pragmatic views on the relationship between thick terms and concepts and evaluation. According to pragmatic views, the evaluations that thick terms and concepts may be used to convey are always some aspect of speaker meaning. According to semantic invariantism, one direction of evaluative valence is semantically fixed but thick terms and concepts may be used to convey evaluations of different valence as a matter of speaker meaning.

Regarding (6), the semantic invariantist could say that uttering (6) isn't a case of literal use but a way of mocking the sorts of prudish evaluations which *lewd* conveys as a matter of meaning.<sup>23</sup> Regarding (7), the explanation of how *frugal* might be used to convey negative evaluation would be that when Blackburn condemns Tweedy as frugal, the word *frugal* in Blackburn's mouth conveys something positive, but it is clear in context that he as speaker means something negative.

23. Blackburn suggests that *lewd* is often used in such cases to mock prudish sensibilities, although he denies that negative evaluation is semantically fixed (Blackburn 1992; 1998: 103).

Thus the exchange in (10) could sometimes be read as denying a negative implicature regarding Tweedy's hospitality, whereas the exchange in (11) could sometimes be read as a comment on the negative effect that the way in which frugality is good has on Tweedy's dispensation of hospitality. Similarly, regarding (9), the invariantist could say that *cruel* always means something negative as a semantic matter, but it is clear in context that the speaker means something positive. This might be because the speaker is a sadist, or because the cruel action was a response to an unjustified slight which the speaker enjoyed retaliating against, and so on. (What sort of positive evaluation would be conveyed depends on whether *fun* is used literally or to mean *good*.) Regarding (8), one might say that when a considerate action contingently manifests the negative quality of assisting an evil purpose, this fact might be so salient, and considerateness as such so minimally good, that it would be true but misleading to say that wiping the torturer's sweaty brow is good, and instead the salient negative evaluation is conveyed as a matter of conversational implicature. This would be to explain (8) by appeal to general conversational principles, such as that one should make one's conversational contributions informative and relevant (Grice 1975).

It is a further question whether the apparent variability in valence that is exhibited by examples like (6)-(9) is *always* best explained away as a pragmatic spin on a semantically fixed direction of evaluation, although no doubt they sometimes are. Perhaps sometimes considerateness is good in no way, or lewdness bad in no way. Is there a way to increase the general plausibility of semantic invariantism? (7) and its variants can be used to suggest one option.

A speaker who condemns Tweedy as frugal needn't think that he is a worse person for his frugality, but only a worse college master. One way to explain this is to say that what is semantically fixed in *frugal* is only a positive evaluation of people *qua* persons. In that case it would be consistent with the meaning of *frugal* as a term of positive evaluation that Tweedy is bad in a role such as college master. That might even be because of the very features which make him frugal, and thereby good

in some way, as a person. Apparent contextual variability could then be explained away by saying that features which make people frugal may also ground negative evaluations which calling someone frugal can convey by pragmatic means.<sup>24</sup>

Semantic invariantism can also allow context to play a role in determining which evaluations are salient: perhaps *frugal* is semantically fixed as a term of positive evaluation in contexts where someone is assessed as a person but has no semantically fixed direction of valence in contexts where someone is assessed in a specific role such as college master. Thus, if the relevant comparison class in contexts where (7) and (10)-(11) are uttered is college masters rather than persons, that might explain why what is salient in such contexts isn't the semantically fixed positive evaluation of Tweedy as a person but rather some evaluation of Tweedy as master which may (according to this view) be positive or negative so far as the meaning of *frugal* goes. And perhaps other thick terms and concepts exhibit similar behavior.

Semantic invariantist explanations of the apparent contextual variability of thick terms and concepts in evaluative valence can agree with the Link Claim. But they support no version of the Variability Claim which bears on LE in the way that the Variability Argument requires.<sup>25</sup> Essentially similar explanations of apparent variability in valence might be available to views according to which thick terms and concepts have a fixed direction of evaluation because evaluation is related to thick terms and concepts as some such further semantic property as conventional implicature or semantic presupposition.

### 5.3 Semantic Contextualism

A different kind of semantic explanation of examples like (6)-(11) agrees with semantic invariantism that the relationship between thick

terms and concepts and evaluation is semantic but claims that the valence of evaluation is context-sensitive rather than invariable. This kind of semantic contextualism has so far gone largely unnoted in the literature on thick concepts. But what it allows is precisely that the evaluations that thick terms and concepts may be used to convey can be positive in some contexts but negative in others, and yet be a feature of their meaning all the same.<sup>26</sup> We know that what degree of height counts as tall or what intensity and duration of a certain sort of qualitative experience it takes for a sensation to count as painful may vary with context. This sort of semantic contextualism claims that what may vary with context is not only how overt and transgressive of conventional boundaries a sexual display must be to count as lewd but also whether such display would be bad in some way. It caters similar options regarding what sort of expenditure counts as frugal and whether the relevant degree of frugality would be good in some way, what sort of responsiveness to others counts as considerate and whether it would be good in some way, what degree of suffering caused to others and enjoyment taken from causing it counts as cruel and whether that degree of cruelty would be bad in some way, and so on. Thus it would seem to have resources to explain (6)-(9).

This kind of semantic contextualism has resources to explain (10) and (11) as well. If A's utterance in (10) meant that Tweedy is good in some way related to low expenditure (or something along these lines), then B's reply would be understood as preempting the potential implicature that Tweedy's frugality has had a negative effect on the dispensation hospitality. If C's utterance in (11) meant the same, D's reply would be expected to begin *No, ...* and be understood as disagreeing

24. Thanks to Daniel Elstein (p.c.) for suggesting this kind of view to me.

25. Semantic invariantist explanations seem not to require a preference either way regarding the issue of (in)separabilism (recall the end of §2).

26. A natural version of this view is that the "character" of predicates expressing thick terms and concepts is a non-constant function from contexts to some kind of evaluative contents (cf. Kaplan 1989). The senses of expressions which return different contents in different contexts go with their character rather than content; standard examples include indexicals like *I* and *here*.

with the positive evaluation, given the duties of college masters to dispense hospitality. If A's utterance in (10) meant that Tweedy is frugal and bad for it, B's reply would be expected to begin *No, ...* or else be understood as agreeing with A's assessment of Tweedy as frugal but denying a negative evaluation regarding Tweedy's dispensation of hospitality. If C's utterance in (11) meant the same, then D's reply would be understood as agreeing with the negative evaluation. Presumably these also don't exhaust the range of semantic contents that an utterance of the sentence *Tweedy is frugal* could have under contextualism. It is also worth noting that contextualism would seem to have no particular problem explaining how comparative constructions like *too tidy* and *not lewd enough* work. These constructions may be understood as effecting the kind of context shift which can also occur in examples like (5) by supplying a new standard of comparison.

This kind of semantic contextualism portrays the relationship between thick terms and concepts and evaluation as semantic and yet as contextually variable in evaluative valence. (What is stable across literal uses in normal contexts is only the potential of thick terms and concepts to have evaluative content relative to context.) This is to agree with the Variability Claim but reject the Link Claim. The Variability Argument has no bite against this kind of semantic contextualism.

I know of one earlier view that holds that thick terms and concepts are inherently evaluative but accepts the Variability Claim. The trick is to reject the Link Claim. Jonathan Dancy argues that thick terms and concepts involve different directions of evaluation in different contexts and may even involve multiple evaluations of different valence in one and the same context. Calling something lewd, for instance, might typically express something negative, sometimes express that it is just what is called for, and sometimes express some mixture of the positive and the negative, such as that it would be appropriate to respond with some mixture of titillation and embarrassment (Dancy 1995: 265). Dancy thinks that such variability is built into the very meanings of thick terms and concepts. As he puts it, "competence with a thick concept [requires] a general understanding of the *range* of attitudes

associated with the concept" and its "meaning . . . is the range of differences that it can make" to how one should act; its semantic value relative to context is to be found within that range (Dancy 1995: 270; 2004: 107, 194).<sup>27</sup> These are controversial claims about meaning and conceptual competence, especially if distinctions between the meaning of an expression and other information associated with meaning needn't be psychologically significant to ordinary speakers.<sup>28</sup> But presumably they represent only one among many contextualist options. Much more thus remains to be said about contextualist accounts of the relationship between thick terms and concepts and evaluation.<sup>29</sup> Here I wish only to flag contextualism as another possible explanation of apparent contextual variability in evaluative valence.<sup>30</sup>

The form of semantic contextualism discussed above must be distinguished from an alternative according to which thick terms and concepts are context-sensitive with respect to extension without being context-sensitive with respect to evaluative valence.<sup>31</sup> This alternative treats their context-sensitivity only as an unsurprising special case of

27. Dancy understands his view as inseparabilist: there may be no way to determine just what non-evaluative features of a certain general type something must have to count as lewd, courageous, or the like, independently of determining what evaluations of it would be appropriate in the particular context. But semantic contextualism might also be understood along separabilist lines: thick terms and concepts have some content that isn't context-sensitive and some content that is, and their evaluative contents relative to contexts are among the latter.

28. For discussion, see e.g. Soames (2010: Ch. 7).

29. One issue here is how what Dancy calls "semantic particularism" relates to more standard forms of contextualism (Dancy 2004: 194). It could turn out that semantic particularism isn't best understood as a version of contextualism after all. (Thanks to Anna Bergqvist here.)

30. When semantic invariantism about some class of sentences has a contextualist rival, typically a relativist alternative is also available. (Examples include epistemic modals, knowledge attributions, future contingents, and predicates of personal taste.) I doubt that semantic relativism recommends itself as an explanation of contextual variability in evaluative valence, but I have no space to argue this here.

31. See Väyrynen (ms2) for a more extensive discussion of how context-sensitivity in general and gradability in particular bear on the relationship between thick terms and concepts and evaluation.

some already familiar sort of context-sensitivity exhibited by other expressions. It is familiar, for example, that gradable expressions like *tall* and *loud* are context-sensitive, and many thick terms and concepts are gradable.<sup>32</sup> For instance, *selfish* might be regarded as context-sensitive with respect to which preferences for one's own happiness over just how much greater a happiness for others counts as selfish. This hypothesis about context-sensitivity is compatible both with views according to which thick terms and concepts are inherently evaluative in meaning and with views that deny this. Variability in evaluative valence would presumably receive a semantic invariantist explanation in the former case and a broadly pragmatic one in the latter.

The latter form of contextualism might go as follows. First, thick terms and concepts denote different non-evaluative properties in different contexts. For instance, *frugal* can be used to measure different types of expenditure, much as *smart* can mean *book smart* or *street smart* or *large* as applied to cities can measure population or geography; all these expressions may of course denote different degrees of the relevant qualities in different contexts (cf. Glanzberg 2007: 10). Second, thick terms and concepts hold together evaluation and description in such a way that in some contexts *frugal* measures some type of expenditure which is commonly associated with a positive evaluation while in other contexts *frugal* measures some type of expenditure which is commonly associated with negative evaluation. Assuming that it is manifest in context that Tweedy's dispensation of hospitality should be generous, the exchange in (10) can be fine if context determines that *frugal* measures Tweedy's general college expenditure, whereas B's reply will be infelicitous if *frugal* measures Tweedy's hospitality expenditure in particular. The exchange in (11) sounds fine either way, but it would sound odd if *frugal* in that context measured expenditure in some particular dimension other than hospitality.

This sort of contextualism can go with either semantic or pragmatic

views on the relationship between thick terms and concepts and evaluation, depending on the mechanisms by which the contextually determined measures convey evaluation. It is therefore compatible with the Variability Claim and the Link Claim but entails neither and hence doesn't support the conclusion of the Variability Argument.

#### 5.4 Contextual Enrichment

Another option worth mentioning which has so far gone unnoted in the literature is to appeal to mechanisms of "free" enrichment in explaining contextual variability.<sup>33</sup> The idea is that context might be relevant to the interpretation of utterances not just by helping to saturate a slot or variable in the meaning of a sentence (which is typically regarded as the site of contextual action for semantic contextualism), but also by providing additional information which isn't triggered by any component of sentence meaning. (12) and (13) are two putative examples of enriching the interpretation of an utterance in this way:

- (12) "I have had supper." (→ I have had supper tonight.)  
 (13) "I have nothing to wear." (→ I have nothing cool to wear for Emily's party.)

What happens in enrichment is that an expression is contextually given a more specific interpretation than it literally encodes.

Enrichment might play a role in some instances of (7) and (10)-(11). The hypothesis would be that the meaning of *frugal* builds in only some fairly general measure of frugality, such as one based on something like the money spent, but that it is semantically optional for context to add further information which in turn might matter to what sort of evaluation is conveyed by calling something frugal.' A context in which Tweedy can be condemned as frugal might be a context which specifies that Tweedy's expenditure is low with respect to hospitality and that this counts as a bad way for a college master to

32. The references in note 16 discuss the semantics of gradable expressions.

33. On free enrichment, see e.g. Recanati (2004: Ch. 2).

be frugal. But context might instead determine that Tweedy's expenditure is low except on some specific dimension such as hospitality. And even if context determines that his expenditure is low on hospitality, it might also add such information as that Tweedy travels considerable distance to get a good price on claret and so there is no skimping in the opportunities to imbibe.<sup>34</sup> In such contexts, B's reply in (10) would sound just fine, whereas D's reply in (11) would sound odd. Additional information might in general be relevant to what sort of evaluation is conveyed by using a thick term or concept. (Perhaps supplying more specific information about a cruel or considerate act sometimes leads to a valence switch.) Mechanisms of enrichment might therefore help to explain at least some cases of apparent contextual variability in evaluative valence. Just how enrichment bears on the relationship between thick terms and concepts and evaluation depends on what kind of role the kind of material which enrichment is supposed to provide plays in semantic interpretation. It is common to regard enrichment as a conversational mechanism that makes no contribution to semantic content relative to context, but the issue is a matter of controversy.

### 5.5 *What Does All This Show?*

This survey of explanations of the apparent contextual variability of thick terms and concepts in evaluative valence makes no claim to comprehensiveness. There may well be further options to those covered here. Nor is this survey meant to provide a critical assessment of the options that it does cover. No doubt some of them embody more plausible views on the relationship between thick terms and concepts and evaluation than others.<sup>35</sup> There are various questions and concerns to raise about each explanation which my discussion neither raises nor aims to raise. This paper is therefore not meant to favor any particular explanation of the apparent contextual variability of thick terms and

34. Thanks to Michael Glanzberg for this example.

35. I argue that thick terms and concepts are not inherently evaluative in meaning on the basis of different linguistic evidence in Väyrynen (ms1).

concepts with respect to their evaluative valence.

None of these limits in my discussion undermine its point, however. Its point is a modest one: there are many available explanations of this apparent variability in valence, but only some of these explanations support the Variability Argument for pragmatic views on the relationship between thick terms and concepts and evaluation. Some of these explanations, such as semantic invariantism, challenge the Variability Claim. Others, such as those contextualist explanations according to which the relationship between thick terms and concepts and evaluation is semantic and yet contextually variable, undermine the Link Claim by their mere availability. So long as at least some of these alternatives to pragmatic views remain in the running, considerations of variability in valence won't suffice to settle the relationship between thick terms and concepts and evaluation. Thus linguistic evidence concerning such variability seems not to take us very far in answering LE. These conclusions represent a null result that may feel disappointingly thin. But there you have it nonetheless: the relationship between thick terms and concepts and evaluation must be settled on grounds other than whether thick terms and concepts are contextually variable with respect to the valence of the evaluations they may be used to convey.

## 6. Conclusion

Let me close with a suggestion for a more promising line of research. Certain types of disagreement involving thick terms and concepts seem to me to provide a better guide than contextual variability to both the relationship between thick terms and concepts and evaluation in general and, as a consequence, the plausibility of the Variability Claim in particular.

One type of apparent disagreement arises when people who apply thick terms and concepts aren't disposed to make the same inferential transitions with respect, for instance, to how it would be appropriate to react. It can happen that two people call one and the same action brutal but only one of them infers from this that the action was thereby bad in some way whereas the other doesn't. One case of the latter sort of reac-

tion could be our earlier example (9) (*Yes, cruel certainly, but that's just what made it such fun*). What is going on in these cases is a substantive issue. One option is that the relevant term (*brutal, cruel, etc.*) is used evaluatively by some but descriptively by others. In that case there would be a ready explanation of the different inferential transitions, but the disagreement would concern at most what the meaning of the relevant term is, or should be, and might be merely apparent. Another option is that one person might fail to draw a negative inference because her grasp of the concept (*BRUTAL, CRUEL, or whatever*) is partial or defective, or even because he doesn't grasp the same concept as the more virtuous people. The moral blindness of a schoolyard bully who fails to see anything wrong with his behavior could be a case in point. (Especially in the latter scenario the extent of real disagreement might depend on the extent of similarity between the different concepts.) A further option is to explain the difference in inferential transitions in terms of differences in substantive conceptions of cruelty which the different speakers associate, perhaps very closely, with their terms or concepts but are consistent with a shared, univocal term or concept. Perhaps it is semantically permissible to resist a negative inference because it is mistaken but not necessarily confused to think that a cruel response might sometimes be permissible or even merited. In this case genuine moral disagreement concerning something that is agreed to be cruel would seem to be semantically possible.<sup>36</sup>

Only some accounts of this type of disagreement are compatible with the Variability Claim or, equally, with the claim that thick terms and concepts are inherently evaluative in meaning. How apparent disagreements of this kind are best explicated therefore bears directly on the relationship between thick terms and concepts and evaluation.

Another type of apparent disagreement arises when a thick term or concept is used to convey evaluations that some people regard as

objectionable in the sense that they really ought not to be endorsed. For instance, *lazy* and *industrious* are commonly used to convey, respectively, negative and positive evaluations. Those who don't value hard work might well resist applying either term to anything at all when context is one where an evaluative outlook, such as the Protestant work ethic, which entails those evaluations is widely shared or otherwise manifest.<sup>37</sup> For in such contexts utterances featuring these terms would convey evaluations rejected by those who find the Protestant work ethic objectionable and thereby mislead others about what evaluations they accept. Is such resistance best explained by analyzing thick terms and concepts as inherently evaluative in meaning? Others might happily apply *lazy* to more or less the same things as those who value doing hard work that one is able to do, but mean to convey something positive. Would they be in a genuine evaluative disagreement, and therefore share a concept to some significant degree, with those who endorse the sort of work ethic that *lazy* and *industrious* are typically used to convey? Or would they simply wield different concepts which happen largely to coincide in (intended) extension? Similarly, imagine someone who is happily using *lazy* and *industrious* but then stops to smell the flowers, changes one's mind about the value of hard work, and yet continues applying these terms to more or less the same things, but without the evaluations she used to use them to convey. It is a substantive issue whether such cases must involve a change in concepts or whether one and the same concept could persist through such changes in one's evaluative outlook. In the former case, evaluation could be built into the meanings of *lazy* and *industrious*; in the latter case, evaluation wouldn't seem to be essential to these concepts.

Only some accounts of this type of disagreement are compatible with the Variability Claim or, equally, with the claim that thick terms

36. This seems to be Hume's position on "monkish virtues" (Hume 1751/1998: IX.1). Hume agrees that some things exhibit self-denial and humility, but argues that the claim that these qualities are virtues is false.

37. See Hare (1981: 17) on *lazy* and *industrious*, Gibbard (1992) and Blackburn (1992; 1998: 101-2) on *lewd* and *chaste*, and (Väyrynen 2009; ms1) and Eklund (forthcoming) for discussion of the general phenomenon and its implications for the nature of thick terms and concepts and their relationship to evaluation.

and concepts are inherently evaluative in meaning. How apparent disagreements of this kind are best explicated therefore bears directly on the relationship between thick terms and concepts and evaluation.

I conclude that analyzing the various sorts of disagreements in which thick terms and concepts might be centrally featured holds a much better promise of helping to settle whether thick terms and concepts are inherently evaluative in meaning than considerations of contextual variability. Arguments from contextual variability of thick terms and concepts with respect to their evaluative valence to conclusions about the relationship between thick terms and concepts and evaluation are inconclusive and can be fully expected to remain so.<sup>38</sup>

### References

- Blackburn, Simon (1992). "Through Thick and Thin." *Proceedings of the Aristotelian Society* supplementary volume 66, 284-99.
- Blackburn, Simon (1993). *Essays in Quasi-Realism* (Oxford: Oxford University Press).
- Blackburn, Simon (1998). *Ruling Passions* (Oxford: Clarendon Press).
- Cresswell, Max (1977). "The Semantics of Degree." In *Montague Grammar*, ed. B. Partee (New York: Academic Press), 261-92.
- Dancy, Jonathan (1995). "In Defence of Thick Concepts." *Midwest Studies in Philosophy* 20, 263-79.
- Dancy, Jonathan (2004). *Ethics without Principles* (Oxford: Clarendon Press).
- Dancy, Jonathan (2009). "Moral Particularism." *The Stanford Encyclopedia of Philosophy* (Spring 2009 Edition), ed. E. N. Zalta.
38. Many thanks to audiences at Universities of Leeds, Manchester, and Sheffield, and at the British Society for Ethical Theory conference at University of Nottingham, for useful discussions of this paper. I am especially grateful for helpful comments from Matti Eklund, Daniel Elstein, Debbie Roberts, Adam Sennet, and two anonymous referees for this journal (one of whom has since identified himself as Simon Kirchin). Work on this paper was supported by a leave award from the Faculty of Arts at University of Leeds and by the European Community Seventh Framework Programme (FP7/2007-2013) grant no. 231016.
- <http://plato.stanford.edu/archives/spr2009/entries/moral-particularism/>
- Eklund, Matti (forthcoming). "What Are Thick Concepts?" *Canadian Journal of Philosophy*.
- Elstein, Daniel Y. and Hurka, Thomas (2009). "From Thick to Thin: Two Moral Reduction Plans." *Canadian Journal of Philosophy* 39, 515-36.
- Foot, Philippa (1958). "Moral Arguments." *Mind* 67, 502-13.
- Gibbard, Allan (1992). "Thick Concepts and Warrant for Feelings." *Proceedings of the Aristotelian Society* supplementary volume 66, 267-83.
- Glanzberg, Michael (2007). "Context, Content, and Relativism." *Philosophical Studies* 136, 1-29.
- Grice, H. P. (1975). "Logic and Conversation." In *Syntax and Semantics, Volume 3: Speech Acts*, ed. P. Cole and J. Morgan (New York: Academic Press), 41-58.
- Hare, R. M. (1952). *The Language of Morals* (Oxford: Clarendon Press).
- Hare, R. M. (1981). *Moral Thinking* (Oxford: Clarendon Press).
- Hume, David (1751/1998). *An Enquiry Concerning the Principles of Morals*, ed Tom L. Beauchamp (Oxford: Oxford University Press).
- Kaplan, David (1989). "Demonstratives." In *Themes from Kaplan*, ed. J. Almog, J. Perry, and H. Wettstein (Oxford: Oxford University Press), 481-563.
- Kennedy, Christopher (2007). "Vagueness and Grammar: The Semantics of Relative and Absolute Adjectives." *Linguistics and Philosophy* 30, 1-45.
- Kirchin, Simon (2010). "The Shapelessness Hypothesis." *Philosophers' Imprint* 10 (4), 1-28. <http://www.philosophersimprint.org/010004/>
- Klein, Ewan (1980). "A Semantics for Positive and Comparative Adjectives." *Linguistics and Philosophy* 4, 1-45.
- McDowell, John (1979). "Virtue and Reason." *The Monist* 62, 331-50.
- McDowell, John (1981). "Non-Cognitivism and Rule-Following." In *Wittgenstein: To Follow A Rule*, ed. S. Holtzman and C. Leich (London: Routledge & Kegan Paul), 141-62.
- Meier, Cécile (2003). "The Meaning of *Too*, *Enough*, and *So . . . That*."

- Natural Language Semantics* 11, 69-107.
- Miller, Alexander (2003). *An Introduction to Contemporary Metaethics* (Cambridge: Polity Press).
- Platts, Mark (1979). *Ways of Meaning* (London: Routledge, Kegan & Paul).
- Putnam, Hilary (2002). *The Collapse of the Fact-Value Dichotomy and Other Essays* (Cambridge, MA: Harvard University Press).
- Recanati, François (2004). *Literal Meaning* (Cambridge: Cambridge University Press).
- Scheffler, Samuel (1987). "Morality Through Thick and Thin: a Critical Notice of *Ethics and the Limits of Philosophy*." *Philosophical Review* 96, 411-34.
- Schwarzschild, Roger and Wilkinson, Karina (2002). "Quantifiers in Comparatives: A Semantics of Degree Based on Intervals." *Natural Language Semantics* 10, 1-41.
- Soames, Scott (2010). *Philosophy of Language* (Princeton: Princeton University Press).
- Stechow, Armin von (1984). "Comparing Semantic Theories of Comparison." *Journal of Semantics* 3, 1-77.
- Swanton, Christine (2001). "A Virtue Ethical Account of Right Action." *Ethics* 112, 32-52.
- Szabó, Zoltán Gendler (2001). "Adjectives in Context." In *Perspectives on Semantics, Pragmatics, and Discourse*, ed. I. Kenesei and R. M. Harnish (Amsterdam: John Benjamins), 119-46.
- Väyrynen, Pekka (2009). "Objectionable Thick Concepts in Denials." *Philosophical Perspectives* 23, 439-69.
- Väyrynen, Pekka (ms1). "Thick Concepts: Where's Evaluation?" Unpublished.
- Väyrynen, Pekka (ms2). "Thick Concepts and Underdetermination." Unpublished.
- Wiggins, David (1998). "A Sensible Subjectivism?" In his *Needs, Values, Truth*, Third Edition (Oxford: Oxford University Press), 185-214.
- Williams, Bernard (1985). *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press).