

8

Thick Concepts: Where's Evaluation?¹

Pekka Väyrynen

1. THE ISSUE

A distinction between “thick” and “thin” terms and concepts is nowadays commonplace in ethics, aesthetics, and epistemology. In ethics the distinction is often introduced by pointing to concepts like CRUEL, SELFISH, COURAGEOUS, GRATEFUL, and KIND or the corresponding predicates, and contrasting them with the sorts of concepts that are often expressed in evaluative contexts by predicates like *good*, *bad*, *wrong*, and *ought*.² The intuitive contrast is that thick terms and concepts somehow “hold together” evaluation and non-evaluative description, whereas thin terms and concepts are somehow more purely evaluative or normative.³ Consider, for instance, a bully who enjoys making other kids suffer. I can convey a negative evaluation of his bullying ways by calling them cruel and I can do so by calling them bad. But even if taking enjoyment in causing others to suffer

¹ An earlier version of this chapter was presented as “Thick Concepts and Presupposition” to audiences at: University of Stirling; University of Reading; 11th Bellingham Summer Philosophy Conference at Western Washington University; 3rd Rocky Mountain Ethics Congress at University of Colorado, Boulder; a joint workshop of Centre for Metaphysics and Mind (University of Leeds) and the Centre for the Study of Mind and Nature (University of Oslo) at University of Leeds; and 7th Wisconsin Metaethics Workshop at University of Wisconsin, Madison. I am grateful to all these audiences for helpful and enjoyable discussions of this material in sessions and after-hours. Thanks also to Andrew Alwood, Jamin Asay, Matti Eklund, Daniel Elstein, Simon Kirchin, Brent Kyle, Angela Mendelovici, Debbie Roberts, Dennis Whitcomb, and anonymous referees for useful written comments. Reasons for the change of title will hopefully be clear from the essay. Work on the essay was supported by a leave award from the Faculty of Arts at University of Leeds, the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 231016, and the Arts and Humanities Research Council (UK) under grant agreement no. AH/H038035.

² I'll use small caps to denote concepts and italics to denote linguistic expressions.

³ Attempts to explicate this intuitive contrast raise subtle issues that I'll mostly set aside in this paper. For discussion, see Scheffler (1987) and Eklund (2011).

counts as both cruel and bad, only *cruel* seems to require as a matter of meaning that things falling under it involve taking enjoyment in causing others to suffer. The meaning of *bad* imposes no such constraint, since many bad things have nothing to do with it. So what is supposed to be distinctive of thick terms and concepts relative to the thin is that they bear some substantive conceptual connections to non-evaluative descriptions.⁴

In what sense exactly are thick terms and concepts evaluative? No doubt there is *some* important sense in which I would manifest a defective moral sensibility if I didn't think that cruel things are bad. More to the issue at hand, no doubt there is *some* important sense in which I wouldn't fully understand what concerned parents try to communicate when they utter sentences like *Bullying is cruel*, if I failed to grasp the sort of negative evaluation that they mean to convey. But just what kind of failure would be involved depends on how utterances involving thick terms and concepts convey evaluation. The question is this:

Evaluation Question (EQ): What is the relationship between thick terms and concepts and the evaluations that they may be used to convey?

I use "convey" to cover different means (such as content, presupposition, or implicature) by which utterances may transfer information. By "evaluations" I mean not mental acts but information or contents that are evaluative.

According to the standard answer to EQ, the evaluations that thick terms and concepts may be used to convey are built into their sense or semantic meaning.⁵ The aim of this essay is to challenge this view and present an alternative.⁶ The challenge is based on linguistic data that are well explained by the alternative view that the evaluations that thick terms and concepts may be used to convey are a certain kind of implications of their utterances which have a conversational explanation.⁷ As I'll explain later, much of these data can also be explained in a way that is consistent with the standard

⁴ This contrast remains even if the meanings of thick terms and concepts merely restrict their application to things of some general non-evaluative type without specifying which specific features of that type count. For instance, the meaning of *selfish* doesn't specify just which preferences for one's own happiness over others' count as selfish. See Elstein and Hurka (2009).

⁵ See Foot (1958), Platts (1979, ch. 10), McDowell (1981), Williams (1985), Hurley (1989), Dancy (1995), Elstein and Hurka (2009), Kirchin (2010), Kyle (2011), and Roberts (2011).

⁶ I challenge this view also in Väyrynen (2009). The present essay advances the earlier one by offering a rival hypothesis about the relationship between thick terms and concepts and evaluation as an explanation of the linguistic evidence discussed below.

⁷ Earlier examples of this kind of view (albeit with an unnecessary commitment to expressivism) include Hare (1952), pp. 121–2; (1981), pp. 17–18, 73–5, and Blackburn (1992).

semantic view. I believe that the alternative I am offering provides a better overall account of the relationship between thick terms and concepts and evaluation, but I won't be able to argue the point fully here; I develop a fuller argument in other work.⁸ Taken by itself this chapter aims to show only that this alternative view is a serious rival to the standard semantic view.

The relationship between thick terms and concepts and evaluation is important in its own right. It bears, among other things, on what notions count as evaluative in the first place. But thick terms and concepts are also often claimed to have deep and distinctive significance to moral philosophy. Moral theory in the traditional style of Kant, Mill, Sidgwick, and Ross is charged with oversimplifying our evaluative thought because of its exclusive focus on thin concepts. Attention to thick concepts is supposed to show evaluative thought and judgment, distinctions between "fact" and "value," the relationship between evaluative and non-evaluative distinctions, the objectivity of value, and more, in a new light.⁹ The view on the relationship between thick terms and concepts and evaluation that this essay makes a start in defending has largely deflationary consequences for claims on behalf of the deep and distinctive significance of thick terms and concepts. I'll explain some of these consequences in closing.

2. THE METHODOLOGY

A project like mine requires fixing some common ground among different answers to EQ regarding the notion of meaning, the notion of evaluation, and appropriate ways of approaching their relationship in the case of thick terms and concepts.

What count as the evaluations whose relationship to thick terms and concepts is supposed to settle EQ? Care is due here because all sorts of terms can be used to convey all sorts of evaluations in suitable circumstances through various mechanisms of information transfer. Since most of us believe that pleasure is good, calling something *pleasant* typically implicates that it is good, but this shows only that the term *pleasant* can be used for evaluative purposes, not that it is evaluative in meaning. That is one possibility about the thick as well. But picking the right evaluations is important to assessing the matter. For instance, calling something *cowardly*

⁸ See Väyrynen (Unpublished).

⁹ See Foot (1958), Platts (1979, ch. 10), McDowell (1981), Putnam (1981; 2002, ch. 2), Williams (1985), Hurley (1989), Dancy (1995), Kirchin (2010), and Roberts (2011).

often conveys moral condemnation, but showing that moral condemnation isn't built into the meaning of *cowardly* wouldn't show that no other sort of evaluation is either. Care is due also because there might be no uniform specification of the relevant evaluations and no uniform account of their relationship to thick terms and concepts. Such uniformity assumptions are common in the literature on thick concepts but by no means obviously true.¹⁰

The way I propose around these complications is twofold. First, define the focus of EQ as evaluations that are *most closely connected* to thick terms and concepts, whatever those evaluations may be like in content and however close that relationship turns out to be in the end. (I'll use the locution "the evaluations that thick terms and concepts may be used to convey" to refer specifically to these evaluations.) Second, refer to these evaluations with a placeholder that covers evaluations with many different kinds of content.

My proposal for a placeholder expression assumes that evaluation is somehow positive or negative in flavor. More precisely, I'll understand evaluation as information to the effect that something has (or lacks) merit, worth, or significance (that is, has a positive or a negative standing) relative to a certain kind of standard, namely a standard that grounds claims of merit, worth, or significance.¹¹ What I'll assume is that whatever the evaluations that are most closely connected to thick terms and concepts may be like otherwise, they will be positive or negative in this sort of way.¹²

My placeholder expressions for these evaluations are *good in a way* and *bad in a way*.¹³ These expressions are broad in scope. Being morally, aesthetically, or prudentially good (bad) are ways of being good (bad). Being intrinsically, finally, or instrumentally good (bad) are ways of being

¹⁰ Many further complications remain. For instance, recent treatments of evaluative language often draw a four-part classification of evaluative expressions as *thin*, *thick*, *pejorative*, and *expressive*. It is unclear whether any single notion of evaluation captures them all.

¹¹ This idea derives in part from Williams (1985), p. 125. To anticipate the discussion of "objectionable" thick concepts below, this notion of evaluation allows that the application of a term might imply a positive (negative) standing relative to a standard that might be thought to ground a claim of merit (demerit) but in fact doesn't do so.

¹² I'll focus exclusively on what Daniel Elstein and Thomas Hurka call "global" evaluations rather than what they call "embedded" evaluations (Elstein and Hurka 2009). The issue at stake is, roughly, not whether the meaning of *courageous*, for instance, builds in some such evaluative condition as "acting for the sake of good goals" (where this embedded evaluation is independent of whether pursuing those goals is courageous), but whether it builds in the condition that accepting certain risks of harm for the sake of good goals whose value is greater than the badness of the harms risked is good in a certain way. For more, see Väyrynen (Unpublished).

¹³ A similar strategy can be found in Kyle (2011), pp. 7–8.

good (bad). Being admirable, praiseworthy, or condemnable (perhaps in a certain respect) are ways of being good or bad. Being good or bad in some way related to sexual display, or in some way related to taking enjoyment in causing others to suffer, are also ways of being good or bad. Being courageous, lazy, or cruel might be thought to be ways of being good or bad. Some of these ways of being good or bad mark thin or general evaluations, whereas others are restricted, relativized, or more specific in various ways. Any of them might in principle be claimed to belong to the meanings of thick terms and concepts.

Because these placeholders are as flexible as they are, more or less everyone can adopt the terminological convention that the evaluations that are most closely connected to thick terms or concepts are evaluations as good or bad in some way.¹⁴ For instance, the issue whether thick terms and concepts bear conceptual connections not only to some non-evaluative descriptions but also to some information that is positive or negative in flavor is neutral with respect to whether thick terms and concepts mark inseparable fusions of description and evaluation. The same is true of whether the evaluative aspect of the thick is reducible to the thin (whatever the answer to EQ). Words like *good* or *bad* may stand for all manner of evaluations that aren't automatically thin just because these words make an appearance; consider *a good thief* or *good with children*. It therefore seems legitimate to adopt expressions like *good in a way* and *bad in a way* as placeholders for the sorts of evaluations that are to be used in settling the relationship between thick terms and concepts and evaluation.¹⁵

The methods with which I'll approach this relationship are standard in semantics and the philosophy of language. I'll consider how well different answers to EQ can explain various judgments from speakers. Of special importance to semantics are judgments that reveal competence regarding meaning, including entailment and contradiction judgments, synonymy and non-synonymy judgments, and judgments about truth conditions.¹⁶ These contrast with judgments that are of special importance to pragmatics, including judgments about the wider import of what speakers meant, implied, suggested, or presupposed by their utterances and judgments about whether an utterance was misleading or infelicitous in some way

¹⁴ The one exception may be those who think that thick terms and concepts are evaluative in some distinctive sense that doesn't involve positive or negative flavor. I think it is fair to say that this notion of evaluation has yet to be satisfactorily explicated.

¹⁵ This is of course far from the whole story about what is involved in evaluative uses of language. One big debate concerns whether assertive uses of these expressions express cognitive or non-cognitive mental states, and whether the other sort is conveyed by some other means.

¹⁶ On the data for semantics, see Chierchia and McConnell-Ginet (2000), ch. 1.

not related to its meaning or truth. It is fair to say that sensitivity to these distinctions isn't the strong suit of the existing literature on thick concepts.¹⁷ Attention to them may thus prove illuminating more generally. But here I'll be primarily interested in using various judgments from speakers to test the predictions borne by different answers to EQ.

These methods constrain the notion of meaning that they can be used to study. (Or perhaps they are better seen as based on a certain notion of meaning.) Entailment and contradiction judgments, for instance, work as clues to semantic meaning only insofar as information belonging to the meaning of an expression cannot be overtly denied or canceled without contradiction and only insofar as such judgments are stable across contexts. Accordingly I'll assume that the meaning of an expression or a sentence at least imposes a set of constraints on what its literal uses express which are stable across normal contexts.¹⁸ (Note that judgments about what count as such constraints may not be universally shared, since they may reflect idiolects or terms may be vague or unclear.)

This assumption about meaning distinguishes what linguistic expressions mean on a given occasion of use from what people mean or intend by deploying them on those occasions.¹⁹ It does this without prejudging general issues about the role of context in determining semantic content. It also leaves open the possibility that ordinary speakers may not draw psychologically significant distinctions between the meaning of an expression in this sense and other information they or their audiences may associate with that meaning. A speaker might link notions like cruelty or kindness firmly or even invariably to a certain kind of evaluation, but what this tells us regarding EQ is simply a further question. Even so, the assumption is substantive enough to give us a bead on the relationship

¹⁷ How the semantics/pragmatics distinction should be drawn is of course itself a controversial matter. (For different views, see e.g. the papers in Szabó 2005.) Throughout the essay I'll try to stay fairly neutral on these sorts of general issues in the philosophy of language.

¹⁸ Soames (2008, 2010, ch. 7) argues that meaning just *is* such a set of constraints; for a rival view that accepts my assumption, see e.g. King and Stanley (2005).

¹⁹ See Grice (1957; 1968) and Kripke (1977) for this contrast between semantic meaning and speaker meaning. At the level of thought we can distinguish between information that constitutes a concept and information one associates with it. Whether these distinctions would be isomorphic depends on the relationship between language and concepts.

Also note that the assumption I am making about meaning doesn't fully distinguish semantic meaning from certain further properties of conventional meaning. Conventional implicatures, for instance, are non-cancelable but truth-conditionally irrelevant aspects of the conventional profiles of the expressions that carry them. I'll discuss corresponding views about the relationship between thick terms and concepts and evaluation in detail in Väyrynen (Unpublished).

between thick terms and concepts and evaluation. For now we can ask whether the evaluations that thick terms and concepts may be used to convey are stable and indefeasible across their literal uses in normal contexts in the way meaning and semantic entailment are thought to be.

So far I have been talking about thick terms and concepts in one breath. Discussions of thick concepts rarely specify what they take concepts to be. What seems often to be assumed is that thick concepts are much the same as the meanings of thick terms. (The assumption might be that words *have* meanings whereas concepts *are* meanings.) But, depending on what concepts are, the mapping from thick terms to concepts might be one-to-many rather than one-to-one. Perhaps, for instance, concepts are something like Fregean senses.²⁰ This would make them more fine-grained than the entities, such as possible worlds or intensions, that semantic theories often assign to linguistic expressions as semantic values. Conclusions concerning the meanings of thick terms may thus not apply directly to thick concepts. Issues here include what sorts of information concepts may encode, how idiosyncratic or "cheap" they may be, and what kind of differences between thinkers generate differences in concepts.

I wish to remain fairly neutral on issues that determine how freely one may move between language and concepts. I won't therefore resist the possibility that some people might have concepts that build in evaluation even if thick terms weren't inherently evaluative in meaning. The flip side of this possibility is that there might be no such thing as *the* concept CRUEL (KIND, and so on) which would be a privileged candidate for the meaning of *cruel*. Particular concepts of cruelty (kindness, and so on) might also exhibit structures only some aspects of which correspond to stable aspects of public language. In that case there would be good grounds to draw distinctions between aspects of those concepts which track standard semantic and pragmatic distinctions but needn't be psychologically significant to the thinkers who wield those concepts. Linguistic evidence concerning the meanings of thick terms may in this way be relevant to the relationship between thick concepts and evaluation.

A more general point along these lines is that it would be odd if the contents of thick concepts floated completely free of the meanings of thick terms. One would expect the relationship between thick terms and evaluation at least to constrain the relationship between thick concepts and evaluation. Claims on behalf of the deep and distinctive significance of thick concepts to evaluative thought and judgment would seem to imply

²⁰ It is common in analytic philosophy to understand concepts in terms of their theoretical role in explaining phenomena of cognitive significance. This is to locate them in the realm of sense, as this is typically meant to contrast with reference.

that thick terms relate to evaluations in some stable way that isn't easily removed by context. Since the factors that determine how thick terms work are presumably the same sorts of factors as those that determine how the rest of language works, considerations of methodological unity and simplicity recommend seeing how far one can get with issues regarding thick concepts with general and independently plausible semantic and pragmatic principles. The proof of this approach will of course be in the pudding.

3. OBJECTIONABLE THICK TERMS AND CONCEPTS

The linguistic data against which I'll address EQ concern the phenomenon of "objectionable" thick terms and concepts. This phenomenon arises when a thick term or concept (as used in a certain way) embodies values that really ought not to be endorsed.²¹ Similarly, someone regards a thick term or concept (as used in a certain way) as objectionable if it embodies values that she thinks ought not to be endorsed. For instance, those who use the words *lewd*, *chaste* and *unchaste*, and *lascivious* typically accept a prudish outlook on sexuality which many of us reject. To regard *lewd* as objectionable when so used is to regard the way of thinking conveyed by such uses as failing to draw genuine evaluative distinctions. Similarly, those who reject the Protestant work ethic tend to regard *industrious* and *lazy* as objectionable when used to evaluate positively those who exemplify hard work and negatively those who avoid hard work when able to do it.

What fundamentally is regarded as objectionable in these cases are certain ways of thinking; I'll address these indirectly through their reflections in language. The data I'll describe are meant to generalize to any thick terms or concepts that are in principle open to being regarded as objectionable in this sense. This is a very wide range of thick terms and concepts. A fundamentalist may find *tolerant* or *forgiving* objectionable when used to convey positive appraisal and an egoist may find *selfish* objectionable when used to convey negative appraisal. The fundamental difference between thick terms and concepts that are in fact objectionable and those that are open to being regarded as such but in fact aren't objectionable is a substantive evaluative matter. There is no reason to suppose that they will differ fundamentally in their logic or meaning. Throughout I'll use "objectionable thick terms and concepts" as shorthand for the broader notion of

²¹ See Blackburn (1992), Gibbard (1992), Richard (2008), Väyrynen (2009), and Eklund (2011).

thick terms and concepts that are in principle open to being regarded as objectionable, at least when used in certain ways.

In considering how the evaluations that objectionable thick terms and concepts may be used to convey behave, I'll adopt *lewd* as my main example. I'll draw occasional comparisons and contrasts with other terms primarily as a reminder that the data I'll use *lewd* to illustrate are supposed to generalize to the other thick terms and concepts that are in principle open to being regarded as objectionable. I'll take it as common ground that *x is lewd* conceptually entails some such description as *x is sexually explicit* or *x is an overt display of sexuality that transgresses conventional boundaries*. These might not be the strongest such descriptions, but they will do as proxies. Whatever the strongest conceptually entailed *lewd*-free description may be, I won't assume that it will be sufficient to satisfy *lewd*, let alone extensionally equivalent with it.

I'll also take it as common ground that people who use the term *lewd* in the first place normally call something lewd not just when they think that it involves overt and transgressive sexual display, but only if they also judge it to be worse for it in some way. So when I hear someone call something *lewd*, I have a pretty good (but defeasible) bead on what the speaker is trying to convey. Such speakers don't typically mean merely that the thing in question involves sexual display of a certain kind, but also that it is the worse for it in whatever way prudish people regard lewdness as bad. Those who don't go in for that kind of sexual morality regard *lewd* as objectionable when used in this way. This looks like a genuine disagreement of values in the sense in which disputes over any information that may be communicated by an utterance may count as disagreements over an utterance, including not just disputes over the truth of its content but also disputes over information conveyed through such other modes of information transfer as presupposition, implicature, and manner.²² The description of the case doesn't settle whether the evaluations on which the disagreement is focused belong to the literal meaning of *lewd* or to some other information that speakers associate closely with its meaning.

I should also note that in what follows I'll largely bracket comparisons between objectionable thick terms and concepts, on the one hand, and ethnic slurs and other pejorative expressions, on the other. The theoretical location of the derogatory or pejorative aspect of the latter expressions is itself a controversial matter. Moreover, however it is resolved, the same resolution cannot be assumed in advance to provide the best account of the

²² See Sundell (2011). If disagreement over an utterance had to concern the truth of its content, describing this kind of case as a disagreement would prejudice EQ.

evaluations that terms like *cruel*, *courageous*, and so on, may be used to convey.²³

4. EVALUATIONS AND PROJECTION

One way to test whether some piece of information that may be conveyed by an utterance belongs to its meaning is to consider its projection behavior, in the following sense of “projection”:

Projection An implication projects if and only if it survives as an utterance implication when the expression that triggers the implication occurs under the syntactic scope of an entailment-canceling operator.²⁴

Certain types of implications exhibit distinctive patterns of projection and non-projection that distinguish them from semantic entailments that arise in virtue of meaning in the sense laid down above. If an evaluation conveyed by literal uses of a thick term or concept in normal contexts satisfies Projection, then that evaluation isn’t built into the meaning of the thick term or concept in question. Thus Projection provides a way to test whether the meanings of thick terms and concepts build in evaluations as good or bad in some way.

The operators to which Projection refers include questions, negation, possibility modals and epistemic modals, antecedents of conditionals, and possibly more. To illustrate, Projection is satisfied, among many other cases, by certain implications of utterances containing factive verbs like *regret*. Consider that (1) has at least the two implications in (a) and (b):

- (1) Anna regrets drinking instant coffee in the morning.
 (a) Anna has negative feelings about drinking instant coffee in the morning.
 (b) Anna drank instant coffee in the morning.

Of these implications, (b) but not (a) is also conveyed by the question version of the sentence in (1), as in (2a), or when (1) is embedded under entailment-canceling sentential operators, such as negation, as in (2b), an epistemic modal, as in (2c), or the antecedent of a conditional, as in (2d).

²³ One difference concerns belief reports. We seem to find it acceptable to utter such reports as *The Pope believes that the Rio carnival is lewd* even if we find *lewd* as used by the Pope objectionable. Reports of analogous utterances involving ethnic slurs tend to be found much less acceptable, increasingly so as the slur in question becomes more explosive.

²⁴ Simons et al. (2010), p. 309.

Thus, by Projection, the second but not the first implication of (1) projects. (The first implication is presumably a semantic or conceptual entailment.)

- (2) (a) Does Anna regret drinking instant coffee in the morning?
 (b) Anna doesn't regret drinking instant coffee in the morning.
 (c) Anna might regret drinking instant coffee in the morning.
 (d) If Anna regrets drinking instant coffee in the morning, then my dislike of instant coffee is vindicated.

Projection is widely used as a diagnostic for presupposition.²⁵ It occurs systematically also with various other truth-conditionally irrelevant implications that don't have certain standard properties of presupposition, including conventional (but not conversational) implicatures and more besides.²⁶

In this section I'll present data that are well explained by the hypothesis that the evaluations in virtue of conveying which a thick term or concept is in principle open to being regarded as objectionable satisfy Projection. The availability of such an explanation challenges the semantic view on EQ even if the latter can also explain the data. For if (as I argue in other work) the former provides a better explanation of the data, then we shouldn't take those evaluations to belong to the sense or semantic content of thick terms and concepts.

Determining whether the evaluations that objectionable thick terms and concepts may be used to convey satisfy Projection requires priming an appropriate context. Projection allows that utterance implications may be defeasible and thus may not always arise in the first place when the expressions that trigger them occur unembedded. One may not in general be able to "read off" utterance implications from the sentence uttered; such implications depend on what is consistent with semantic entailments, context, real world knowledge and other background assumptions, implicatures, and presuppositions. (In Section 5 below I'll suggest that even the evaluations that are most closely connected to thick terms and concepts are defeasible in certain ways.) It would be inappropriate to stipulate the context as one where *lewd* is used to make a negative evaluation. But what we do know is that typically only those with prudish views on sexual display go in for calling things *lewd*.

Now imagine a conversation between two persons we know to be prudes about an upcoming concert that is known to involve explicit sexual display.

²⁵ See e.g. Levinson (1983), Soames (1989), Chierchia and McConnell-Ginet (2000), Kadmon (2001), Simons (2006), and Huang (2007).

²⁶ See e.g. Karttunen and Peters (1979), Chierchia and McConnell-Ginet (2000), Barker (2003), Potts (2005), and Simons et al. (2010).

(The performer might be Madonna, Lady Gaga, Britney Spears, Rihanna, or Iggy Pop for that matter.) Suppose one utters (3) and the other utters (4):

- (3) Madonna's stage show is lewd.
- (4) Madonna's stage show isn't lewd.

In the context primed above the speaker of (4) is naturally heard as denying that Madonna's show involves sexual display that is explicit to a degree that it takes to count as lewd. Such an utterance might well be insufficiently informative, since typically we expect speakers who deny what others have said to offer reasons for denial. But if the speaker of (4) supplied a reason, one would expect her to say something in the spirit of (5):

- (5) Madonna's show is sexually insinuating alright, but it's not lewd because no private parts are exposed.

The standard for lewdness is a potential subject of disagreement among prudes.

What the speaker of (4) isn't naturally heard as denying are some such evaluations as the generic claim in (6) or the singular conditional claim in (7):

- (6) Overt displays of sexuality that transgress conventional boundaries are bad in some way.
- (7) If Madonna's show involves overt display of sexuality that transgresses conventional boundaries, then it is bad in some way.

It is reasonable to suppose that in the context primed above the speaker of (4) accepts something like (6) or (7).²⁷ Treating (6) or (7) as the principal implication of both (3) and (4) in such contexts would also explain why uttering (4) doesn't imply that Madonna's show *is* bad in some way but uttering (3) does. The inference from (6) or (7) to that singular evaluation is blocked in the case of (4) but not in the case of (3).

It seems clear that those who find *lewd* as used by prudes objectionable have a significantly different reaction to (4). These *lewd*-objectors (to use a nomenclature due to Brent Kyle) will typically have nothing at stake in the sort of disagreement concerning the extension of *lewd* that is expressed by (5) (Gibbard 1992). We saw that utterances of (4) in the context primed above are naturally read as using negation in the ordinary truth-conditional

²⁷ The relevant readings are again those that allow disputes among prudes over how explicit a sexual display must be to count as satisfying *lewd*.

way to express this sort of disagreement, not to question evaluations like (6) or (7).²⁸ The context is a conversation in which utterances of (3) are reasonably taken to imply negative evaluations whose truth isn't at issue among prudes but which *lewd*-objectors reject. In such contexts *lewd*-objectors should typically, other conversational purposes aside, not be prepared to utter (4), since such utterances won't typically convey the disagreement they have with those evaluations.

The data here concern the conversational acceptability of certain utterances to *lewd*-objectors in a certain kind of context. It is a further question how they would respond if, for instance, they were forced to render a verdict on (3) and (4) as true or false. Perhaps *lewd*-objectors should be willing to classify certain things as *lewd*, especially if they can do so in a way that is dissociated from the evaluations they reject. But this is a controversial matter (see also Section 6).

The differences in acceptability judgments between prudes and *lewd*-objectors can be explained by treating evaluations such as (6) or (7) as implications of utterances of (3) that satisfy Projection. What Projection essentially tests for is whether an implication of an utterance is in the background of a conversation or, instead, part of what is "at issue" among the speakers when the utterance is made.²⁹ (This distinction concerns the discourse status of a piece of information, which may vary across contexts.) The hypothesis that the evaluations that (objectionable) thick terms and concepts may be used to convey in the kind of context primed above are implications that satisfy Projection may therefore be understood as the hypothesis that those evaluations aren't in such contexts part of the main point of utterances involving thick terms and concepts.

This hypothesis has several attractions. If such evaluations as (6) and (7) were background assumptions in typical conversations among prudes, then the speaker of (4) wouldn't be naturally heard as denying them in such contexts. The apparent survival of such evaluations under ordinary truth-conditional negation in (4) therefore has a straightforward explanation if those evaluations aren't part of the main point of utterances of sentences like (3) in typical conversations among prudes. That explanation relies on

²⁸ "Metalinguistic" uses of negation may be different; see Väyrynen (2009; Unpublished) and Section 5 below.

²⁹ See e.g. Stalnaker (1974), Soames (1989), p. 353, Chierchia and McConnell-Ginet (2000), p. 350, Potts (2005), pp. 31, 58, and Simons et al. (2010). An implication that is backgrounded in this sense may introduce new or even controversial information. It needn't belong to the "common ground": the information state that contains all the propositions (including general world knowledge and discourse knowledge) shared by the discourse participants.

the hypothesis that the evaluations that thick terms and concepts may be used to convey aren't built into the meanings of thick terms and concepts but are rather implications of their utterances which are typically not at issue in normal contexts. (They may, of course, become part of what is at issue if, say, a *lewd*-objector enters the conversation.)

The evaluations that those thick terms and concepts that are in principle open to being regarded as objectionable may be used to convey seem to satisfy Projection also with respect to other entailment-canceling operators besides negation. Consider first the question version of the sentence in (3):

(8) Is Madonna's show lewd?

Uttering (8) would be a natural way for a prude to inquire whether to protest against Madonna's show. It is naturally heard as asking whether Madonna's show is sexually explicit enough to make the grade for lewdness but not as subjecting evaluations like (6) and (7) to question. This would explain *lewd*-objectors' unwillingness to utter (8).

Next consider embedding in possibility and epistemic modals:

- (9) (a) Madonna's show might be lewd.
 (b) Maybe/it is possible that Madonna's show will be lewd.
 (c) Presumably/probably Madonna's show will be lewd.

Uttering the sorts of sentences found in (9) would be a natural way for prudes to compile a list of potential events to protest against. They are naturally heard as modalizing the claim that Madonna's show is sexually explicit enough to make the grade for lewdness but not as modalizing evaluations like (6) and (7). This would explain *lewd*-objectors' unwillingness to utter the sentences in (9).

Finally consider embedding in the antecedent of a conditional:

(10) If Madonna's show is lewd, the tabloid press will go nuts.

Uttering (10) would be a natural way for a prude to state that the sensationalist press will dwell lavishly in pictures and text on Madonna's show if it is sexually explicit enough to make the grade for lewdness. (10) is naturally heard as conditionalizing on this non-evaluative aspect of (3) but not on evaluations like (6) or (7). This would explain *lewd*-objectors' unwillingness to utter (10).

Each of (8)–(10) therefore resembles (4) in that their utterances in the kind of context primed above seem to cancel the entailments of (3) regarding the involvement of sexual display explicit enough to make the grade for lewdness but to imply some such evaluation as (6) or (7). This is to say that the evaluations that those thick terms and concepts that are in principle open to being regarded as objectionable may be used to convey

seem to satisfy Projection. This would make good sense. Since implications that satisfy Projection are attached to an utterance of the given sentence not only when it is asserted but also when it is denied, questioned, or offered as a possibility or a hypothetical assumption, it is unsurprising that *lewd*-objectors would be unprepared to assert any of (3)–(4) and (8)–(10) in the sort of context primed above.

A related phenomenon worth noting is that certain implications tend not to fall into the scope of certain attitudes when embedded in the complement clauses of the corresponding attitude verbs. Examples of such verbs are *hope* and *regret*.³⁰ Compare (11) and (12):

- (11) I regret that Israel and Palestine aren't managing to form a two-state solution.
 (12) I regret that Madonna's show is lewd.

(11) carries the (defeasible) implication that Israel and Palestine are trying to form a two-state solution. But this isn't part of what the speaker regrets; she regrets that their attempts are failing. It is rather something she assumes (truly or not) as background.

The evaluations that thick terms and concepts may be used to convey seem to behave similarly. Evaluations like (6) or (7) wouldn't seem to be part of the regret in (12) as uttered in the sort of context primed above. Such utterances are natural to prudes but would typically be unacceptable to *lewd*-objectors. So evaluation (if present in the first place) seems to escape from the scope of regret when embedded in the complement of *regret*. That isn't what one would expect if (6) or (7) were semantic entailments of (3), which is the embedded sentence in (12).³¹ So the way these evaluations behave when embedded in the complements of certain attitude verbs is further evidence that they are typically not part of the main point of utterances of atomic sentences involving thick terms and concepts.

³⁰ The implications in question include at least presuppositions. See von Fintel (2004), Yablo (2006), and, for some complications, Beaver and Geurts (2011, Section 5.3).

³¹ One might worry about this test. One might think that the hope that would typically be expressed by an utterance of *I hope that Sam is a bachelor* by an appropriately specified female is that Sam is unmarried and that the proposition that Sam is male is exempted from the content of the hope, although the sentence entails rather than implies it. I am not convinced. The hope that one would typically express by uttering this sentence would be crushed if Sam turned out to be female. That suggests that the content of the hope is that Sam is an eligible unmarried male; it is just that some aspects of this content may sometimes be more conversationally salient than others.

5. EVALUATIONS AND DENIABILITY

Let's now turn from the projection behavior of the evaluations that thick terms and concepts may be used to convey to their deniability. We saw that *lewd*-objectors find *lewd* as used by prudes objectionable because utterances of it convey evaluations they reject. We also saw that uttering (4), read as using negation truth-conditionally, isn't typically a good way to express this disagreement with prudes. So can *lewd*-objectors reply to prudes' utterances of sentences like (3) (and their embeddings under entailment-canceling operators) in ways that target specifically the evaluations they reject for denial? And what would this tell us about the relationship between thick terms and concepts? Insofar as the evaluations that thick terms and concepts may be used to convey are plausibly treated as defeasible implications of utterances involving thick terms or concepts, the less plausibly those evaluations are treated as belonging to their meanings.³²

One observation in favor of the view that evaluation is built into the meanings of thick terms and concepts is that conjoining an unembedded occurrence of a thick term or concept with an overt denial of the evaluation that it is typically used to convey makes for an infelicitous utterance (marked by a "#"):

- (13) #Madonna's show is lewd, but it is in no way bad for involving explicit sexual display.
 (14) #Madonna's show is lewd and not bad in any way.

The apparent contradiction in (13) and (14) might be regarded as evidence that an evaluation like (6) or (7) is a semantic entailment of (3).³³

This is too quick. Certain implications that satisfy Projection are such that the conditions for their felicitous denial exhibit an important asymmetry: the outright denial of such an implication in a co-ordinate clause of a positive sentence without apparent contradiction is normally not possible, but it can often be overtly denied in a co-ordinate clause of a negative sentence without any apparent contradiction. A standard example is

³² Such defeasibility would count also against views according to which the evaluations that thick terms and concepts may be used to convey are truth-conditionally irrelevant properties of the conventional meanings of thick terms and concepts, such as conventional implicatures.

³³ This move is made most explicitly by Kyle (2011), p. 10.

presupposition.³⁴ Overt presupposition denial is infelicitous in (15) but perfectly felicitous in (16):

- (15) #Anna regrets drinking instant coffee because in fact she never did so!
 (16) Anna [doesn't/can't] REGRET drinking instant coffee; she never drank instant coffee in the first place.
Doesn't presuppose: Anna drank instant coffee.

The hypothesis that the evaluations that thick terms and concepts may be used to convey satisfy Projection can thus explain why the denials in (13)–(14) are infelicitous. The apparently contradictory character of those denials thus doesn't support analyzing evaluations like (6) or (7) as semantic entailments of (3). But are there reasons to think that those evaluations are defeasible implications whose denials exhibit the kind of asymmetry found in (15)–(16)?

One linguistic device for expressing disagreement are linguistic denials marked by expressions like *nope*, *nuh uh*, and *no it isn't*. Such denials can be equally felicitous in disagreements over the truth of the content of an utterance and over various sorts of implications (such as presuppositions and implicatures) triggered by utterances (Sundell 2011). For instance, the factive presupposition of an utterance of *Anna regrets drinking instant coffee* can be felicitously denied by uttering *No she doesn't, she never drank it*. Would similar denials of evaluations like (6) or (7) be felicitous in the sort of context primed above?

- (17) (a) Madonna's show is lewd.
 (b) Nuh uh, things are in no way bad just for involving overt sexual display that transgresses conventional boundaries.
 (b') No it isn't, its involving explicit sexual display doesn't mean that it's bad in any way.

My own judgment on (17) isn't firm. The rectifying clauses in responses (b–b') strike me as somewhat forced, but this doesn't mean that they are infelicitous. No matter. Either the exchange is felicitous or it isn't. If it is, (17) doesn't support the claim that the evaluation denied in (b–b') is a feature of the content of (a). And if it isn't, (17) only shows that a certain pattern of denial tends to be unavailable to *lewd*-objectors however the content of (a) is analyzed.

Next let's consider how *lewd*-objectors might respond to uses of *lewd* that seem to preserve evaluations like (6) and (7) under operators other

³⁴ See Gazdar (1979), pp. 119–23, Huang (2007), pp. 71–2, and Beaver and Geurts (2011), Section 3. Presuppositions may not be the only implications exhibiting this asymmetry which satisfy Projection.

than negation. Responses that focus denial on those evaluations seem to be available. We saw that *lewd*-objectors typically have no stake in disputes about how explicit a sexual display must be to count as lewd. It is then unclear why they couldn't simply bracket the issue whether Madonna's show counts as lewd when responding to (3) by denying the negative evaluation. Compare (18) and (19):

- (18) Whether or not Madonna's show is lewd, it's in no way bad for being sexually explicit.³⁵
 (19) Whether or not Madonna's show involves explicit sexual display, it would be in no way bad for that.

(18)–(19) are aimed to deny the sorts of evaluations that *lewd*-objectors reject but which utterances of both (3) and (4) typically imply in the mouths of prudes. I find no great difference in their acceptability to *lewd*-objectors, at least when (19) is read with a degree of explicit sexual display which prudes would typically consider lewd. Because uttering (18) would sidestep the issue whether Madonna's show is or might be lewd to focus on what concerns *lewd*-objectors in utterances that imply the evaluations they reject, it would seem to be a suitable way for them to deny the negative evaluations conveyed by questions like (8) and modalized claims like those in (9).³⁶ (Sidestepping an issue in this way can be a polite way to convey that one regards a question that exercises another speaker as irrelevant.)

One might worry how well data concerning terms like *lewd* will generalize. Many people report to find such analogues of (18) as (20)–(21) at least odd:

- (20) ?Whether or not risking one's life to protect the civilization is courageous, it's in no way better for accepting such a risk of harm to oneself for the sake of such goods.
 (21) ?Whether or not bullying is cruel, it's in no way worse for involving enjoyment in causing other kids to suffer.

The oddness of (20)–(21) would constitute no objection to my argument if *courageous* and *cruel* weren't even in principle open to being regarded as objectionable when used to convey, respectively, positive and negative

³⁵ Not everyone agrees that (18) can be acceptable to *lewd*-objectors. If that is because they find corresponding utterances unacceptable in the case of racial slurs, the analogy is irrelevant for reasons mentioned at the end of Section 3. It might be, for instance, that a racial slur in the form of (18) is unacceptable simply because slurs are taboo words (Anderson and Lepore forthcoming).

³⁶ Analogous cases might include those where someone concerned with law rather than morality objects to ϕ -ing by saying that ϕ -ing is illegal whether or not the violated law is just.

evaluation. It would simply impose a principled limit on the scope of the argument.

More importantly, *courageous* and *cruel* might be in principle open to being regarded as objectionable in their typical evaluative uses even if we had trouble imagining objections to the values that those uses typically embody.³⁷ Imagining evaluative possibilities that one rejects is notoriously subject to imaginative resistance.³⁸ It is therefore wholly predictable that it may be difficult for *us* to get into the right head space to assess whether (20)–(21) would be acceptable to those who regard *courageous* and *cruel* as objectionable in their typical evaluative uses. (This doesn't, of course, settle whether all or only some thick terms and concepts are in principle open to being regarded as objectionable.)

Let's finally turn to overt denials of implications in the co-ordinate clauses of negative sentences, as in (16). Their felicitousness tends to be improved by additional intonational stress (marked with small caps):

- (16) Anna {doesn't/can't} REGRET drinking instant coffee; she never drank instant coffee in the first place.
 (22) Grandma didn't go PEE; she went to freshen up.
 (23) The show wasn't GOOD, it was great!

In each case, the material under the scope of the negation in the first half seems to be naturally heard as echoing an actual or potential positive utterance by another speaker, followed in the second half by a rectifying clause that expresses a comment on some aspect of that utterance. (The relevant aspect is presupposition in (16), register in (22), and a scalar implicature in (23).)

Let's replicate this pattern in denials of evaluations like (6) and (7). Consider utterances of (25) and (26) in reply to an utterance of (3), repeated here for convenience as (24), by a prude:

- (24) Madonna's show is lewd.
 (25) Madonna's show isn't LEWD; things are in no way bad just for overt sexual display that transgresses conventional boundaries.
 (26) Madonna's show isn't LEWD because it's not in any way bad for involving explicit sexual display.

³⁷ An example due to Tom Dougherty might help. Imagine Himmler briefing Hitler as follows: "Mein Führer, we have found an excellent new way to exterminate Jews. It is really cruel." This sounds coherent, something we understand, and not obviously parasitic (e.g. an "inverted-commas" use) on the typical use of *cruel* to convey negative evaluation.

³⁸ For discussion, see e.g. Gendler (2000) and Weatherson (2004).

My judgment is that (25)–(26) should typically be acceptable to *lewd*-objectors when uttered with additional intonational stress. What would explain this?

One option is to interpret (25)–(26) as utterances that deny the whole content of (3)/(24) in the way truth-conditional negation does but present its evaluative aspect as the salient ground for denial. According to this analysis, (25)–(26) don't instantiate the pattern in (16) and (22)–(23). For it implies that the positive claim denied in the first half of (25)–(26) and the rectification in the second half are in fact contradictory, which isn't the case in (16) and (22)–(23). The idea is that (25)–(26) are instead parallel to (27):

(27) Sam isn't a bachelor because he isn't eligible to marry.

But the putative parallel fails with respect to (18). (18) involves no contradiction, apparent or otherwise, whereas the corresponding example in (28) clearly does:

(28) #Whether or not Sam is a bachelor, he isn't eligible to marry.

I am therefore inclined to dismiss this rival explanation.³⁹

Another option is to explain the acceptability of (25)–(26) to *lewd*-objectors by interpreting the appearance of *lewd* in (25)–(26) as echoing its use in (3)/(24) and then focusing denial on the evaluations that *lewd*-objectors reject.⁴⁰ Especially insofar as (25)–(26) would strike us as odd if uttered on their own out of the blue, they are better heard as comments on an actual or potential utterance of a sentence like (3)/(24) by a prude, much in the way (16) and (22)–(23) are naturally heard as echoing another utterance.

Examples like (16) and (22)–(23) are commonly treated as instances of a "metalinguistic" use of negation. This is a device for objecting to a previous utterance on the basis of its presuppositions, implicatures, or linguistic form or manner, rather than its content. Implications that survive under truth-conditional negation and satisfy Projection can often be denied through not only the sorts of linguistic denials considered above (this they have in common with semantic entailments) but also through metalinguistic negation. If (25)–(26) were instances of metalinguistic negation, that would further strengthen the case for treating the evaluations they deny

³⁹ Another disanalogy is that the felicitousness of (27) isn't improved by intonational stress (on *bachelor*), whereas it seems to improve the acceptability of (25)–(26) to *lewd*-objectors.

⁴⁰ This is to interpret *lewd*-objectors as not using *lewd* in the way it is used literally in typical utterances of (4)–(5).

as implications of utterances involving *lewd* that satisfy Projection. For it would imply that evaluation isn't built into the sense or semantic content of thick terms and concepts.

Unfortunately the tests that have been proposed for metalinguistic negation are both too complex to be run properly here and controversial.⁴¹ I won't therefore here suppose that (25)–(26) in fact are instances of metalinguistic negation. I'll only suppose that (25)–(26), when uttered in the sort of context primed above, are naturally read as attributing certain aspects of utterances of (3)/(24) to someone other than the speaker herself and expressing a negative comment on those aspects—in this case, an evaluation like (6) or (7).

The question that now remains is whether the negative comment is a comment on the truth of (3)/(24) or on some other aspect of its utterance that may be consistent with its truth. The latter view, which falls out of my preferred answer to EQ, allows that objectionable thick terms and concepts can in principle be true more or less of the same things of which they would (at least paradigmatically) be true if they weren't objectionable. What goes wrong with them is that they "misevaluate" because they are closely associated with evaluations that don't "fit" the things of which they are true (Eklund 2011). The former view, which falls out of the semantic view, treats (25)–(26) as more like (29):

(29) Madonna's show isn't lewd, because nothing is lewd.

There are several ways to approach this issue.

One option is to ask whether *lewd*-objectors can consistently think that some claims of the form *x is lewd* are true even if it would typically be misleading for them to *say* so. Imagine in all graphic detail some event that pruders would count as paradigmatically lewd and suppose that a *lewd*-objector is forced to render a verdict on the truth or falsity of the claim that it is lewd. No doubt some *lewd*-objectors have idiolects which render the verdict "false" and count sentences like (29) as straightforwardly true. But some people are happy to say that a *lewd*-objector needn't be making a conceptual mistake if their verdict is that such a sentence can be true.⁴²

⁴¹ The classic discussion of metalinguistic negation is Horn (1989), ch. 6. For criticisms of the three tests for metalinguistic negation proposed by Horn, see Carston (1998) and Geurts (1998). I suggest treating examples like (25)–(26) as instances of metalinguistic negation in Väyrynen (2009). The argument there is insufficient, but I provide a fuller argument in Väyrynen (Unpublished).

⁴² These philosophers include Richard (2008), Eklund (2011), and myself. For instance, Richard writes that "the fact that an act was lustful was not, in and of itself, a reason to condemn it; the fact that an act was chaste was not, in and of itself, a reason to praise . . . it just wouldn't be responsive to [insist that such a person] was laboring under a

It isn't at all obvious that such verdicts must be explained as "inverted commas" or other non-literal uses of *lewd*, even if the cases where *lewd*-objectors are willing to use *lewd* include its non-literal use aimed to mock prudish sensibilities.⁴³ The intuition that (29) seems to be acceptable to *lewd*-objectors could be explained by saying that (29) is naturally interpreted as saying that nothing is lewd in the way prudes have in mind. (This may, again, be the sense that many idiolects attach to *lewd*.)

Another option is to ask whether one can coherently imagine a context where evaluations like (6) and (7) are somehow or other suppressed (perhaps it is known to all conversational participants that most of them reject those evaluations) but someone raises the question whether something is lewd. (Or perhaps the question they raise is whether it is more lewd than something else.) Participants in such a context might well think that asking whether the thing is lewd might not be the best way of putting the question, given what is known about typical uses of *lewd*, but it is much less clear that it would thereby be incoherent to render the verdict that the thing in question is lewd. The point is of course that in such contexts the verdict "true" wouldn't trigger the misleading implication that the thing in question is bad in the way prudes would think it is.

A further option is to consider certain nearby questions. One is what to say about cases where a speaker applies a thick term or concept but leaves the truth of the evaluation that it may be used to convey open. Consider (30):

- (30) Whether or not this is a good thing, John can be truthfully and neutrally described as being considerate.⁴⁴

Uttering (30) might not be acceptable to *considerate*-objectors, but even its acceptability to *considerate*-agnostics is enough to suggest that thick terms and concepts can be applied literally while bracketing the evaluations they are typically used to convey. It is then not clear why speakers should be

'conceptual confusion' (Richard 2008), p. 32. And later he writes: "No (thick) concept is immune from being used by someone who has dropped its evaluative trappings... And since *we* can describe what would be happening in such a situation, *we* can correctly identify a conceptual structure as realizing a certain thick concept *without* making it essential, for a structure's realizing the concept, that it involve any particular evaluative attitude" (Richard 2008), p. 33.

⁴³ Simon Blackburn gives examples where *lewd*-objectors are willing to use *lewd* in ways that are parasitic on the typical, negative uses (1998), p. 105. Elsewhere he implies that such examples don't exhaust the cases (Blackburn 1992). Be that as it may, such examples don't settle the appropriate verdict of truth or falsity when that is what *lewd*-objectors are forced to issue.

⁴⁴ Bergström (2002), p. 5.

committed to regarding the thick terms or concepts they find objectionable as empty.⁴⁵

No doubt these sorts of considerations won't move everyone. One might continue to think that consistency requires *lewd*-objectors to judge all literal utterances of sentences of the form *x is lewd* to be false, at least if forced to render a verdict. More remains to be said on both sides. The best account of (25)–(26) depends, among other things, on fuller assessments of whether they can be instances of metalinguistic negation and whether one can consistently regard thick terms and concepts that one finds objectionable as not empty.

I conclude that although certain denials of the evaluations conveyed by typical utterances of sentences like (3)/(24) are infelicitous, certain others are licensed as felicitous moves in conversations between prudes and *lewd*-objectors. There are some reasons to think that some of these denials are consistent with the sentences whose utterances convey those evaluations, but such reasons remain inconclusive here. What they do suggest is that it is a serious possibility that the evaluations that thick terms and concepts (or at least those that are open to being regarded as objectionable) may be used to convey are defeasible implications of their utterances that can be removed by context. If that is right, those evaluations don't belong to the meanings of thick terms and concepts.

6. COMPARING EXPLANATIONS

I have argued that the hypothesis that the evaluations that objectionable thick terms and concepts may be used to convey are implications that satisfy Projection provides a good explanation of why even certain complex utterances that embed *lewd* are unacceptable to *lewd*-objectors and why certain denials of those evaluations seem acceptable to them. One response to this is that a different but equally good explanation is available to those who treat the relationship between thick terms and concepts and evaluation as semantic.⁴⁶ This view interprets *lewd*-objectors as thinking that *nothing* is *lewd*, and thus accepting claims along the lines of (29), on the grounds that *nothing* is (roughly speaking) an overt and transgressive display of sexuality and bad in some way for it.

⁴⁵ A related question: are such utterances as *If Madonna's show were accurately described as lewd, I wouldn't count it as thereby bad in any way* really inconsistent or conceptually confused?

⁴⁶ This response is developed in Kyle (2011), pp. 25–34.

Let's first see how this view plays out with (8)–(10). If nothing is lewd, then it would seem odd for someone who thinks so to ask of any particular thing if *it* is lewd, as happens with Madonna's show in (8); the question is already settled in her mind. It would similarly seem odd for her to say that something *might* be lewd, as in (9), or that *if it is lewd*, something else will follow, as in (10).

This explanation of judgments regarding (8)–(10) doesn't follow directly from the semantic view. All that the semantic view contributes to the explanation is that utterances of (8)–(10) don't entail the evaluations in virtue of which *lewd*-objectors regard *lewd* as objectionable. But it follows from independent pragmatic principles to the effect that embedding *p* in the antecedent of a conditional, in a question, and in the scope of a possibility operator carries a "clausal implicature" to the effect that the speaker treats the truth of *p* as a live (but perhaps unsettled) option.⁴⁷ Such implicatures concerning (6) and (7) are unacceptable to *lewd*-objectors. So, although utterances of (8)–(10) don't entail these evaluations, their unacceptability to *lewd*-objectors can be explained on the basis of unwanted implicatures they trigger.

Let's next see how this view plays out with (4). If nothing is lewd, then it follows in particular that Madonna's show isn't lewd. So (4) is straightforwardly true on the semantic view, since it denies both that the show is sexually explicit and the concomitant negative evaluation. But true utterances may be misleading. As noted earlier, (4) fails to specify why the speaker thinks that (3) is false. Utterances of (4) are typically interpreted as denying just that Madonna's show is sexually explicit. We saw some potential problems with this explanation of the unacceptability of (4) to *lewd*-objectors. It requires that (25)–(26) be parallel to (27), but this is questionable with respect to its implications for (18). The explanation also doesn't seem to capture the apparently backgrounded status of general evaluations of lewd things as bad in some way.

One might remedy these problems by appealing to the independent pragmatic principle that the negation of the positive element of an antonym pair (such as *happy/unhappy*) entails just that the positive predicate doesn't hold (which allows a neutral state) but also implicates by "negative strengthening" that the negative predicate holds.⁴⁸ Utterances of (31a) entail (31b) and imply (31c):

⁴⁷ See e.g. Gazdar (1979), pp. 59–62 and Levinson (1983), pp. 436–7; (2000), pp. 108–9) for discussion of clausal implicatures, and Kyle (2011), pp. 33–4 for its application to the present data.

⁴⁸ See e.g. Horn (1989), pp. 331ff. and Levinson (2000), pp. 127ff. for discussion of negative strengthening, and Kyle (2011), pp. 28, 31 for its application to the present data.

- (31) (a) Droopy isn't happy.
 (b) It isn't the case that Droopy is happy.
 (c) Droopy is unhappy.

One might then similarly think that utterances of (32a) not only entail (32b) but also imply (32c):

- (32) (a) Iggy isn't chaste.
 (b) It isn't the case that Iggy is chaste.
 (c) Iggy is unchaste.

Chaste-objectors are just as unprepared to utter *x is unchaste* as to utter *x is chaste*. Both utterances embody values they reject. The same goes for whatever might be an antonym for *lewd*. So the unacceptability of (4) to *lewd*-objectors can again be explained on the basis of an unwanted implicature of its utterances.

My own view is that this explanation of the data is coherent but inferior to the explanation provided by the hypothesis that the evaluations that thick terms and concepts may be used to convey are themselves implications that satisfy Projection. My case for this claim remains inconclusive in this essay, however, because here I can only briefly describe some claims in support of this conclusion, which I defend at greater length in other work.⁴⁹

One reason to prefer my explanation is that it is unnecessary to assign the evaluations that thick terms and concepts (or at least those that are open to being regarded as objectionable) may be used to convey to their meanings. Those evaluations can instead be explained on the basis of general conversational principles. The aim is to discharge the burden of explaining how those evaluations arise, if not in virtue of the meanings of thick terms and concepts. If those evaluations can be explained on the basis of general conversational principles, then they shouldn't be treated as semantic features. This follows from the widely accepted methodological principle that senses aren't to be multiplied beyond necessity; other things being equal, it is better to postulate pragmatic implications than senses (or conventional implicatures) because the latter are properties of individual expressions that don't typically have a more general explanation.⁵⁰

The conversational explanation of how evaluations arise goes, in a nutshell, as follows. Most of us who use the term *cruel* in the first place believe that cruelty is bad in a certain sort of way, and the same is true of

⁴⁹ See especially chapters 5–7 in Väyrynen (Unpublished).

⁵⁰ This sort of principle can be found e.g. in Stalnaker (1974), p. 50, Grice (1978), pp. 47–8, Levinson (1983), pp. 97–100, 132; (2000), p. 15, Horn (1989), pp. 213–14, 365, 383, and Huang (2007), pp. 7, 37.

many other thick terms, at least for certain historical periods or local cultures. Given that those who use a particular thick term or concept in the first place tend to share certain evaluations not accepted by those who don't go in for using it, it wouldn't be surprising if those evaluations were part of a normal understanding of utterances involving the thick term or concept in question. Attributing those evaluations to the speaker would be reasonable on the basis of general knowledge about people who go in for using the term or concept in the first place.⁵¹ For these reasons I suspect that if we stipulated a language in which the meaning of *lewd* doesn't build in evaluation, but held fixed other facts about the linguistic community (including its non-linguistic practices, the evaluative outlooks of its members, and the like), we would still expect audiences normally to interpret someone who utters (3) as meaning to convey negative evaluation.⁵²

It is important to this account that a conversational explanation of some implication of an utterance needn't be specific to its particular context. Implications can become "generalized" in the sense that they are triggered by default by saying a certain type of thing, without requiring special contextual features to arise.⁵³ Since generalized implications are present in all literal utterances in normal contexts unless special contextual circumstances defeat them, a pragmatic account of the relationship between thick terms and concepts and evaluation can explain why thick terms and concepts are very intimately connected to evaluation and the evaluations to which they are connected may be easily mistaken for aspects of semantic meaning. Differences between information belonging to the meaning of an expression and the generalized implications triggered by its utterances could easily fail to be psychologically significant.

⁵¹ For some circumstantial evidence, consider the succession of semantic change recorded for *lewd* in *The Oxford English Dictionary*: *lay* → *unlearned* → *common, low, vulgar* → *ignorant, foolish, unskillful* → *vile, evil, wicked, base* → *bad, worthless, poor* → *lascivious, unchaste*. What is constant here isn't that *lewd* has evaluative meaning but only that its meaning is such that it can easily come to be used primarily for evaluative purposes, given the sorts of evaluative interests that are common among us. This is encoded in some but not all conventional meanings of *lewd*; but whether it is so encoded doesn't seem like a particularly deep fact (cf. Blackburn 1992, p. 297). Such terms as *painful* and *pleasant* also operate primarily in domains that are commonly invested with evaluative significance and normally carry evaluative connotations, but for all that remain non-evaluative in meaning.

⁵² This test for separating semantic meaning from speaker meaning is due to Kripke (1977).

⁵³ See e.g. Grice (1975; 1978) and Levinson (2000). Discussions tend to focus on generalized conversational implicature, but presuppositions and other backgrounded implications could be generalized in the same sort of way (cf. Burton-Roberts 2009).

This explanation can exploit the widely accepted point in the literature on thick concepts that thick terms and concepts are “parochial” in the sense that their use for evaluative purposes tends to derive its point from its role in some particular evaluative perspective.⁵⁴ (Terms like *lewd* and *chaste* are good illustrations of this idea.) Given this point, interpreting a speaker of (3) as intending to convey negative evaluation should tend to be licensed by such principles as that we should read as much into an utterance, beyond the information it actually contains, as is consistent with what we know about the world, or that although one should make one’s contribution as informative as possible, one should say no more than is necessary for achieving this in the conversational context.⁵⁵ Such shifts in the default interpretation of a word denoting a broader set as denoting only a salient subset or stereotypical member of that set tend to be culturally or socially rather than linguistically motivated, and cases differ with respect to whether the narrowing implication has become fossilized into conventional meaning.⁵⁶ It isn’t easy to see why such conventionalization should be a deep fact if (as seems possible) it happened with thick terms and evaluation.⁵⁷

Another reason to prefer my explanation of the data focuses on the fact that the unacceptability of certain utterances to *lewd*-objectors appears to have a unified basis. My account captures this because it treats the evaluative implications on which the unacceptability of these utterances is based throughout as implications that satisfy Projection. By contrast, the rival account requires two completely distinct pragmatic mechanisms on top of a semantic hypothesis to explain these phenomena, so it is both less simple and less unified.

One might object that the claim that the rival explanation is less unified than mine is unfair because my account is restricted to thick terms and concepts that are in principle open to being regarded as objectionable and that I am therefore committed to finding the rest a different explanation.

⁵⁴ See e.g. McDowell (1981), Williams (1985), Hurley (1989), Blackburn (1992), Gibbard (1992), Dancy (1995), and Sreenivasan (2001). In the literature this point is usually coupled with the claim that one may not be able to “master” the extensions of thick terms and concepts without making an attempt to grasp the relevant evaluative perspective. My suggestion here is that this claim doesn’t cut deep insofar as it is plausible; see also the end of this section.

⁵⁵ See Levinson (1983), pp. 146–7 and Horn (1989), pp. 194–5, respectively.

⁵⁶ See Horn (1989), pp. 358, 390 and, on semantic change generally, Traugott and Dasher (2002).

⁵⁷ Cf. note 51 above. What about the intuition that thick terms and concepts feel more robustly and intimately evaluative than terms (like *athletic* or *chocolate*) that have widely shared evaluative connotations? Elsewhere I argue that such intuitions can be explained without projecting the differences they draw into semantics (Väyrynen, unpublished).

But that remains an open question. First, it might be that all thick terms and concepts are in principle open to being regarded as objectionable.⁵⁸ Second, if some aren't, this might mark a genuine distinction among thick terms and concepts which a good explanation should register.⁵⁹ The rival explanation discussed above also seems less robust than mine. Holding all other things fixed, the unacceptability of the relevant sentences to *lewd*-objectors would seem to persist even in languages in which clausal implicatures and negative strengthening weren't typical implicatures of the sorts of sentences that actually do trigger them. My account can capture their unacceptability at least insofar as those evaluations are shared in ways that make them generalized in the above sense.

I cannot here say more to support the hypothesis that the evaluations that thick terms and concepts may be used to convey are conversationally triggered implications that satisfy Projection over the standard view that those evaluations belong to the sense or semantic content of thick terms and concepts. Thus, although I think that the latter view offers an overall worse explanation of the data that appear to suggest that the relevant evaluations satisfy Projection and are deniable in certain ways, I should repeat that it has resources for a coherent explanation of those data. I'll now briefly address methodological concerns that arise from the availability of different explanations of the same data.

One concern is that the data themselves may need more scrutiny. Might the different judgments from prudes and *lewd*-objectors reflect the possibility that (some? many? all?) thick terms are systematically polysemous between an evaluative and a non-evaluative sense?⁶⁰ This possibility remains open insofar as one cannot jump from linguistic evidence straight to conclusions about concepts. But the possibility that thick terms are systematically polysemous doesn't undermine my argument here. Judgments from *lewd*-objectors would seem to pattern in the ways discussed above precisely in contexts where *lewd* would be most plausibly interpreted as being used in its putative evaluative sense and not otherwise. Since those judgments

⁵⁸ Even *just*, one of the stronger candidates for an inherently evaluative thick term, may be such. Some utilitarians deny that justice as such is valuable, and in Plato's *Republic* Thrasymachus condemns justice on the basis of his account of justice as what is in the ruler's interest. Since such claims don't seem conceptually confused, it is unclear what would support the semantic posit that just things are good in a certain sort of way. (Presumably being in the ruler's interest isn't a relevant way of being good here; if so, such a posit wouldn't explain our intuitions of univocality.)

⁵⁹ A further complication (which might cut either way) is that it may be difficult or even impossible confidently to classify various concepts as either thick or thin (Scheffler 1987, p. 417).

⁶⁰ This suggestion has been put to me in conversations.

can be explained by treating evaluation as an implication commonly associated with utterances of *lewd*, they provide no reason to posit a conventional sense of *lewd* that builds in evaluation. Positing polysemy would also imply that evaluation isn't an invariable feature of thought and discourse about the domains in which thick terms operate. Nothing about it would seem to privilege the evaluative senses over the non-evaluative ones either in assigning meanings to thick terms or in theorizing about the significance of thick terms and concepts to evaluative thought and judgment.

Another concern is that judgments from speakers may be insufficient to decide between different hypotheses concerning the relationship between thick terms and concepts and evaluation. This is a complication with which all parties may have to live. Since semantically significant distinctions may not be psychologically significant, it may not be possible to "read off" of speaker judgments whether some feature of an expression belongs to its conventional profile (and, if so, whether it is a truth-conditionally relevant constraint or some other type of conventional constraint) or is merely associated with its use. Do utterances of sentences of the form *x is lewd* sound off in the mouths of *lewd*-objectors because they are linguistically non-defective but (regarded by such speakers as) systematically false? Or do such utterances sound off because they fail some kind of constraint on non-defective use which is consistent with their truth?⁶¹ Or do they sound off simply because such utterances carry negative connotations as generalized implications thanks to certain sociological and cultural facts? These options are relevant to whether regarding a thick term or concept as objectionable commits one to regarding it as empty, since their availability means that the set of things of which a term is true may diverge from the set of things to which its application may be licensed by various kinds of conventions, rules, or customs. The empirically correct choice between them is underdetermined by the mere judgment that sentences of the form *x is lewd* sound off when uttered by non-prudes and may be difficult to decide.

The fact that the distinction between the meaning of an expression and other information associated with it may not be psychologically significant has further significant implications. One popular argument for the view that thick terms and concepts are inherently evaluative in meaning is that evaluation seems to "drive the extension" of thick terms and

⁶¹ Such constraints may be subjective, requiring a speaker who applies *lewd* to a sexual display to regard it negatively, or objective, requiring that an overt sexual display must in fact be bad in some way for *lewd* to be appropriately used (Predelli 2010).

concepts.⁶² For instance, conversational participants can use *lewd* univocally while disagreeing over its extension. This is possible if (but not only if?) they agree that lewd things are sexually explicit to a degree (whatever it may be) that is bad-making. In that case, their judgments about just *which* sexual displays are lewd are sensitive to evaluation. But again it may be unnecessary to posit that the evaluative implications on which they agree belong to the meaning of *lewd*. If semantically significant distinctions may not be psychologically significant, then judgments about what falls into the term's extension may easily be driven in part by factors that aren't part of its meaning, especially if they are implications (such as generalized implications) that are easily mistaken for aspects of meaning. The ability to track such judgments may therefore similarly involve grasping both semantic and non-semantic information. (How strongly extension judgments are driven by semantic information is particularly controversial in the case of objectionable thick terms and concepts, since it is controversial whether these must be regarded as empty.)

This isn't to deny that the application and withholding of thick terms and concepts may be sensitive to the evaluative perspectives from which their use for evaluative purposes derives its point. (The way in which thick terms and concepts are parochial is perfectly compatible with the account I am offering.) The point is rather that the role that evaluation may play in determining whether a user of a thick term or concept applies or withholds it may be explicable without locating evaluation in its meaning. For instance, although pruders may disagree about the extension of *lewd*, conversations among them concerning its extension typically involve evaluative interests. They invest the question whether something is correctly called *lewd* with evaluative significance. It would then not be surprising if even the potential extension of *lewd* relative to such contexts were restricted to those sexual displays which it is coherent, from the perspective of prudish sensibilities, to regard as bad. In that case patterns of application and withholding would be sensitive to evaluation. But such conversational facts can be explained simply by treating evaluations like (6) and (7) as background assumptions of the relevant conversational contexts. They don't require locating those evaluations in the meaning of *lewd*. They can therefore be explained by the hypothesis that the evaluations that thick terms and concepts may be used to convey are not semantic entailments but (defeasible) implications that satisfy Projection. In closing I'll briefly discuss some broader implications of this hypothesis for the philosophical significance of thick terms and concepts.

⁶² See especially Dancy (1995) and Elstein and Hurka (2009).

7. SOME BROADER IMPLICATIONS

Insofar as the evaluations that thick terms and concepts may be used to convey are generalized but defeasible conversational implications of utterances involving such terms and concepts, those evaluations won't be in any interesting sense essential to thick terms and concepts. This conclusion has largely deflationary consequences for claims to the effect that thick terms and concepts have deep and distinctive significance to evaluative thought and judgment.

Thick terms and concepts are sometimes invoked to challenge the so-called "fact-value distinction" (or "gap"). The idea is that they combine non-evaluative description and evaluation in some way that collapses or bridges the distinction.

It is hard to see this as a cogent move insofar as thick terms and concepts aren't inherently evaluative in meaning. The problem arises irrespective of whether the fact-value distinction is understood as the claim that the factual and the evaluative are mutually exclusive, the claim that evaluative and (other?) factual claims have different "directions of fit" (or differ in some other such constitutive respects), or the claim that one cannot derive an *ought* from an *is* or otherwise reasonably infer evaluative conclusions from purely non-evaluative premises.⁶³ Insofar as evaluation isn't inherent to thick terms and concepts, they cannot serve as counterexamples to any of these claims. For instance, insofar as premises involving thick terms and concepts fall to the *is* side of the *is/ought* distinction, their content will make no distinctive contribution to supporting inferences from non-evaluative premises to evaluative conclusions without auxiliary evaluative assumptions.

The claim that thick terms and concepts challenge the fact-value distinction also requires a particular stand on an issue that occupies much of the literature on thick concepts. That issue is whether thick terms and concepts represent irreducible fusions of evaluation and description or whether their evaluative and non-evaluative aspects are distinct components that can be separated or "disentangled" from one another.⁶⁴ (Thick

⁶³ These different fact-value distinctions crop up in Hume (1739/1740, 3.1.1), Putnam (1981; 2002), Hurley (1989), Smart (1999), and Bergström (2002), among other places.

⁶⁴ For separabilism, see Hare (1952), Blackburn (1992; 1998a), Miller (2003, pp. 244–54), and Elstein and Hurka (2009). For inseparabilism, see McDowell (1981), Williams (1985), Hurley (1989), Dancy (1995), Kirchin (2010), and Roberts (2011).

terms and concepts can be invoked to challenge the fact–value distinction only if their evaluative and non-evaluative aspects are inseparable.) The account of the relationship between thick terms and concepts and evaluation proposed in this essay cuts through much of this debate. Insofar as the evaluations that thick terms and concepts may be used to convey are implications of their utterances, they will be separable from the descriptions entailed by claims involving thick terms and concepts in whatever way the semantic entailments of sentences are separable from the implications of their utterances. Similarly, insofar as those evaluations are, more specifically, defeasible conversational implications, they will again be separable from the contents of the utterances that trigger them in whatever way pragmatic implications are separable from semantic content. These relationships may not be straightforward, since the semantics/pragmatics interface may be far from clean-cut. There may also be no adequate “thickness-free” paraphrases of thick terms and concepts to offer a psychologically significant separation of thick terms and concepts from the evaluations they may be used to convey. But separability there will be all the same.

There are other broader consequences that may be thought to be of significance. One is that insofar as judgments involving thick terms and concepts aren't in themselves evaluative judgments, then those judgments are irrelevant to the debate between cognitivist and non-cognitivist accounts of evaluative judgment in the context of which thick terms and concepts are often invoked.⁶⁵ Another is that insofar as thick terms and concepts aren't inherently evaluative in meaning, the intuition that they (or the properties they ascribe) are more specific than (the properties ascribed by) thin terms and concepts will seem as ill-defined as the claim that *acidic* is more (or less) specific than *blue*. The meanings of thick terms and concepts will provide no grounds for treating the properties they ascribe as specific ways of being good (or bad) in any sense in which being good in a specific way determines being good. Thick terms and concepts will more generally be unlikely to have any distinctive significance to issues about whether evaluative facts aren't brute but obtain in virtue of other facts.

Thick terms and concepts have also been invoked to argue that moral theory in the traditional style of Kant, Mill, Sidgwick, and Ross oversimplifies our evaluative thought. That enterprise is taken to presume that one or another moral concept or consideration is basic and others are to be explained in terms of it. This is supposed to require neglecting

⁶⁵ See e.g. Foot (1958), Hare (1952, ch. 7; 1981, pp. 17–18, 73–5), McDowell (1981), Williams (1985), Hurley (1989), Blackburn (1992), Dancy (1995), Miller (2003, pp. 244–54), and Kirchin (2010).

thick terms and concepts, and thereby oversimplify our evaluative thought, on the grounds that no particular thick term or concept will be basic in this sense and that they cannot in general be explained in terms of thin terms and concepts.⁶⁶ I won't here discuss whether moral theory in the traditional style is guilty as charged or whether that would reflect unfavorably on the enterprise.⁶⁷ The crucial question for my purposes is whether thin terms and concepts are conceptually or explanatorily prior to thick terms and concepts.⁶⁸

Insofar as thick terms and concepts aren't themselves inherently evaluative in meaning or content, they clearly cannot be invoked to argue that thin terms and concepts aren't basic in evaluative thought and judgment. For how could they in that event play a more fundamental role than thin terms and concepts in evaluative thought, judgment, or theory? Note that it wouldn't follow that thin terms and concepts are prior to thick terms and concepts. That claim is most naturally understood as assuming that thick and thin terms and concepts alike fall into the class of evaluative terms and concepts. It seems more likely that thick and thin terms and concepts would exhibit no systematic interesting relations of conceptual priority or interdependence.

All these consequences are compatible with a wide range of other interesting and important relationships between the thick and the thin. We would be just as able as before to say that things are good in certain ways because they are generous, fair, or satisfy a vital need, and bad in certain ways because they are cruel, selfish, or invade privacy, just as we can say that things are good in a certain way in virtue of being pleasant and bad in virtue of being painful. These relations typically hold, if they do, as a substantive matter that is the province of first-order theory of value. (Virtue ethicists tend not to defend their view by mere appeal to the meanings of virtue and vice terms and concepts.) More generally, various non-evaluative terms and concepts, *pain* and *pleasure* among them, have been recruited to play the sort of roles in normative ethics that thick terms and concepts have also been recruited to play. Thick terms and concepts can play those roles even if they aren't inherently evaluative in meaning. So, for better or worse, the claim that thick terms and concepts aren't inherently evaluative seems to carry no corrosive or revisionary consequences for moral theory.

These consequences of the view that the evaluations that thick terms and concepts may be used to convey are a certain kind of defeasible conversational implications of their utterances are stated with qualifications

⁶⁶ See especially Williams (1985) on this charge, as well as Hurley (1989).

⁶⁷ For an insightful discussion of these issues, see Scheffler (1987).

⁶⁸ Susan Hurley calls this view "centralism" (Hurley 1989, p. 11).

("insofar as," "in general"). If not all thick terms and concepts are in principle open to being regarded as objectionable, then my argument leaves open the possibility that some thick terms and concepts are inherently evaluative in meaning and, indeed, that the evaluative character of some of them might be irreducible to thin evaluations, bridge the distinction between fact and value, and so on. The broader consequences of the account of the relationship between thick terms and concepts and evaluation proposed in this essay might therefore turn out to be restricted in scope. I am happy to leave such issues open to future work. Saying that thick terms and concepts have deep and distinctive significance to evaluative thought and judgment will typically (if not with respect to all of the above issues) require more than just that some limited range of thick terms and concepts bear a semantic relationship to evaluation. This stronger claim is what I have sought to challenge in this essay.

REFERENCES

- Anderson, L. and Lepore, E. Forthcoming. "Slurring Words," *Noûs*.
- Barker, S. 2003. "Truth and Conventional Implicature," *Mind* 112, 1–34.
- Beaver, D.L. and Geurts, B. 2011. "Presupposition," in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Summer 2011 Edition)*. <<http://plato.stanford.edu/archives/sum2011/entries/presupposition/>> (last accessed November 15, 2011).
- Bergström, L. 2002. "Putnam on the Fact–Value Dichotomy," *Croatian Journal of Philosophy* 2, 1–13.
- Blackburn, S. 1992. "Through Thick and Thin," *Proceedings of the Aristotelian Society*, supplementary volume 66, 284–99.
- 1998. *Ruling Passions*. Oxford: Oxford University Press.
- Burton-Roberts, N. 2009. *The Limits to Debate: A Revised Theory of Semantic Presupposition*. Cambridge: Cambridge University Press.
- Carston, R. 1998. "Negation, 'Presupposition' and the Semantics/Pragmatics Distinction," *Journal of Linguistics* 34, 309–50.
- Chierchia, G. and McConnell-Ginet, S. 2000. *Meaning and Grammar: An Introduction to Semantics*, second edition. Cambridge, MA: MIT Press.
- Dancy, J. 1995. "In Defence of Thick Concepts," *Midwest Studies in Philosophy* 20, 263–79.
- Eklund, M. 2011. "What Are Thick Concepts?," *Canadian Journal of Philosophy* 41, 25–50.
- Elstein, D.Y. and Hurka, T. 2009. "From Thick to Thin: Two Moral Reduction Plans," *Canadian Journal of Philosophy* 39, 515–36.
- Fintel, von K. 2004. "Would You Believe It? The King of France is Back! (Presuppositions and Truth-Value Intuitions)," in A. Bezuidenhout and M. Reimer (eds), *Descriptions and Beyond*. Oxford: Oxford University Press, 269–96.

- Foot, P. 1958. "Moral Arguments," *Mind* 67, 502–13.
- Gazdar, G. 1979. *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.
- Gendler, T.S. 2000. "The Puzzle of Imaginative Resistance," *Journal of Philosophy* 97, 55–81.
- Geurts, B. 1998. "The Mechanisms of Denial," *Language* 74, 274–307.
- Gibbard, A. 1992. "Thick Concepts and Warrant for Feelings," *Proceedings of the Aristotelian Society*, supplementary volume 66, 267–83.
- Grice, H.P. 1957. "Meaning," repr. in his *Studies in the Way of Words*, 213–23.
- 1968. "Utterer's Meaning, Sentence-Meaning and Word-Meaning," repr. in his *Studies in the Way of Words*, 117–37.
- 1975. "Logic and Conversation," repr. in his *Studies in the Way of Words*, 22–40.
- 1978. "Further Notes on Logic and Conversation," repr. in his *Studies in the Way of Words*, 41–57.
- 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hare, R.M. 1952. *The Language of Morals*. Oxford: Clarendon Press.
- 1981. *Moral Thinking*. Oxford: Clarendon Press.
- Horn, L.R. 1989. *A Natural History of Negation*. Chicago, IL: University of Chicago Press.
- Huang, Y. 2007. *Pragmatics*. Oxford: Oxford University Press.
- Hume, D. 1739/1740. *A Treatise of Human Nature*. (Several editions.)
- Hurley, S.L. 1989. *Natural Reasons*. Oxford: Oxford University Press.
- Kadmon, N. 2001. *Formal Pragmatics*. Oxford: Blackwell.
- Karttunen, L. and Peters, S. 1979. "Conventional Implicatures in Montague Grammar," in C.-K. Oh and D. Dineen (eds), *Syntax and Semantics 11: Presupposition*. New York: Academic Press, 1–56.
- King, J. and Stanley, J. 2005. "Semantics, Pragmatics, and the Role of Semantic Content," in Szabó (ed.), *Semantics versus Pragmatics*, 111–64.
- Kirchin, S. 2010. "The Shapelessness Hypothesis," *Philosophers' Imprint* 10(4), 1–28. <www.philosophersimprint.org/010004/> (last accessed November 15, 2011).
- Kripke, S. 1977. "Speaker Reference and Semantic Reference," *Midwest Studies in Philosophy* 2, 255–76.
- Kyle, B. 2011. *Knowledge as a Thick Concept: New Light on the Gettier and Value Problems*. Ph.D. Dissertation (Cornell University).
- Levinson, S.C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- 2000. *Presumptive Meanings*. Cambridge, MA: MIT Press.
- McDowell, J. 1981. "Non-Cognitivism and Rule-Following," reprinted in his *Mind, Value, and Reality*. Cambridge, MA: Harvard University Press (1998), 198–219.
- Miller, A. 2003. *Introduction to Contemporary Metaethics*. Cambridge: Polity Press.
- Plato. *Republic*. (Several editions.)
- Platts, M. 1979. *Ways of Meaning*. London: Routledge, Kegan & Paul.
- Potts, C. 2005. *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.

- Predelli, S. 2010. "From the Expressive to the Derogatory: On the Semantic Role for Non-Truth-Conditional Meaning," in S. Sawyer (ed.), *New Waves in Philosophy of Language*. London: Palgrave Macmillan, 164–85.
- Putnam, H. 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- 2002. *The Collapse of the Fact/Value Dichotomy and Other Essays*. Cambridge, MA: Harvard University Press.
- Richard, M. 2008. *When Truth Gives Out*. Oxford: Oxford University Press.
- Roberts, D. 2011. "Shapelessness and the Thick," *Ethics* 121, 489–520.
- Scheffler, S. 1987. "Morality Through Thick and Thin: A Critical Notice of *Ethics and the Limits of Philosophy*," *Philosophical Review* 96, 411–34.
- Simons, M. 2006. "Foundational Issues in Presupposition," *Philosophy Compass* 1, 357–72.
- , Tonhauser, J., Beaver, D., and Roberts, C. 2010. "What Projects and Why," in N. Li and D. Lutz (eds), *Proceedings of SALT (Semantics and Linguistic Theory) 20*. Ithaca, NY: CLC Publications, 309–27.
- Smart, J.J.C. 1999. "Ruth Anna Putnam and the Fact-Value Distinction," *Philosophy* 74, 431–7.
- Soames, S. 1989. "Presupposition," in D. Gabbay and F. Guenther (eds), *Handbook of Philosophical Logic*, Volume 4. Dordrecht: Reidel, 553–616.
- 2008. "Drawing the Line Between Meaning and Implicature—and Relating Both to Assertion," *Noûs* 42, 440–65.
- 2010. *Philosophy of Language*. Princeton: Princeton University Press.
- Sreenivasan, G. 2001. "Understanding Alien Morals," *Philosophy and Phenomenological Research* 62, 1–32.
- Stalnaker, R.C. 1974. "Pragmatic Presuppositions," repr. in his *Context and Content*. Oxford: Oxford University Press (2002), 47–62.
- Sundell, T. 2011. "Disagreements about Taste," *Philosophical Studies* 155, 267–88.
- Szabó, Z.G. 2005. (ed.), *Semantics versus Pragmatics*. Oxford: Oxford University Press.
- Traugott, E.C. and Dasher, R.B. 2002. *Regularity in Semantic Change*. Cambridge: Cambridge University Press.
- Väyrynen, P. 2009. "Objectionable Thick Concepts in Denials," *Philosophical Perspectives* 23, 439–69.
- Unpublished. *The Lewd, the Rude, and the Nasty: A Study of Thick Concepts*.
- Weatherston, B. 2004. "Morality, Fiction, and Possibility," *Philosophers' Imprints* 4(3), 1–27. <www.philosophersimprint.org/004003/> (last accessed November 15, 2011).
- Williams, B. 1985. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.
- Yablo, S. 2006. "Non-Catastrophic Presupposition Failure," in J.J. Thomson and A. Byrne (eds), *Content and Modality*. Oxford: Oxford University Press, 164–90.