

What are Thick Concepts?

Author(s): MATTI EKLUND

Source: *Canadian Journal of Philosophy*, Vol. 41, No. 1 (March 2011), pp. 25-49

Published by: Cambridge University Press

Stable URL: <https://www.jstor.org/stable/41302117>

Accessed: 30-03-2026 09:38 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Cambridge University Press is collaborating with JSTOR to digitize, preserve and extend access to *Canadian Journal of Philosophy*

What are Thick Concepts?

MATTI EKLUND
Cornell University
Sage School of Philosophy
Ithaca, NY 14853-3201
USA

Introduction

Many theorists hold that there is, among value concepts, a fundamental distinction between *thin* ones and *thick* ones. Among thin ones are concepts like *good* and *right*. Among concepts that have been regarded as thick are *discretion, caution, enterprise, industry, assiduity, frugality, economy, good sense, prudence, discernment, treachery, promise, brutality, courage, coward, lie, gratitude, lewd, perverted, rude, glorious, graceful, exploited*, and, of course, many others. Roughly speaking, thick concepts are value concepts with significant descriptive content. I will discuss a number of problems having to do with how best to understand the notion of a thick concept. Thick concepts have been widely discussed in the meta-ethical literature. But some important problems concerning what thick concepts are supposed to be have not been squarely addressed even in the most systematic of these discussions. Here I want to highlight these problems.

First I will present some puzzles regarding thick concepts. Then I will consider some proposed accounts of thick concepts and see how well they solve these puzzles. Later I will tentatively defend a particular positive account. But the account relies on the legitimacy of the potentially questionable notion of *epistemic analyticity*; and it also relies on a particular understanding of this notion. Special attention will be paid

to *objectionable* thick concepts: concepts somehow presupposing false evaluative claims. I will argue that objectionable thick concepts present problems for otherwise potentially attractive theories of thick concepts.

While the thin/thick distinction is wholeheartedly embraced in some parts of the literature, other theorists would be more skeptical of the claim that there is a significant distinction there. I will for the most part adopt the working assumption that there is a significant thin/thick distinction. But those inclined to reject that assumption might well take their view to be supported by the problems I here discuss.¹

I Puzzle (I): What Makes a Concept *Thick* Rather than *Thin*?

Someone who holds that there are thick concepts in effect makes two substantive claims. One is that there is some sort of difference between evaluative and descriptive concepts. Already this is a quite substantive claim, and one that would be denied in some quarters. A second is that there is among the evaluative concepts a distinction between thin(ner) and thick(er) ones.

I will for now presuppose that there is a distinction between evaluative and descriptive concepts. Moreover I will, at least at the outset, make the seemingly reasonable assumption is that if C is a positive (negative) value concept then 'x is (pro tanto) good' ('x is (pro tanto) bad') follows analytically from 'x is C.'² Even granted this much, there are problems regarding the thin-thick distinction. To see that there are problems here, note how certain characterizations given in the literature fail.

Here is Allan Gibbard's informal characterization:

(T1) A term stands for a thick concept if it praises or condemns an action as having a certain property.³

1 Throughout the discussion I will simply assume that evaluative and normative predicates genuinely stand for properties. I will simply set aside the simple versions of non-cognitivism that deny this, and I will likewise for simplicity set aside alternative cognitivist suggestions, for example to the effect that 'good' does not function semantically as a predicate but as a predicate modifier.

2 Read this as shorthand for: 'if C is a positive (negative) value concept, then, for any predicate F which expresses C, 'x is F' entails ...'. Throughout this paper I will sometimes talk about concepts where strictly I should speak of *predicates expressing* these concepts.

3 Gibbard (1992), 268f. This is the rough, informal characterization Gibbard starts

One may be uneasy already about the talk of *terms* (as opposed to users thereof) praising or condemning actions. But take such talk on board. There is still a problem. There is as much reason to think that 'good' satisfies this condition as that a term like 'courageous' does. 'Good' stands for the property of being good. And doesn't this term praise actions as having this property, just as 'courageous' praises actions for being courageous? On a non-cognitivist view, such as Gibbard himself defends, it can be denied that 'good' stands for a property. But it is clear that Gibbard does not mean his characterization only to be acceptable to a non-cognitivist.

It may be suggested that Gibbard must mean something like *descriptive* property. We get:

- (T1) A term stands for a thick concept if it praises or condemns an action as having a certain descriptive property.

But what does 'descriptive property' mean? Here is a natural suggestion. Start with a distinction between descriptive and evaluative *expressions* (and since we are talking about properties we can focus on predicates). Then a descriptive property can be said to be one that can be ascribed by a purely descriptive predicate.⁴ But with 'descriptive property' understood this way, (T1) faces problems. For on many metaethical views, even 'thin' moral predicates stand for such properties. A general point may be that even when the distinction between descriptive and evaluative predicates is regarded as unproblematic, it is not quite clear how best to understand the corresponding ontological distinction between descriptive and evaluative *properties*.⁵

I will return below to accounts of thick concepts that are aimed at getting around these problems. But first let me indicate how the problems I have presented for Gibbard's informal characterization arise also for certain other prominent characterizations of thick concepts.

out with. Later in his paper he discusses more sophisticated proposals. I will briefly discuss Gibbard's main positive proposal later.

- 4 'Ascribed,' not 'denoted.' The terminology, and the reason for adopting it, comes from Jackson (1998), 119fn10. If we are currently discussing goodness then 'the property we are discussing' denotes goodness, so if we talked about properties denoted in purely descriptive terms, goodness would count as a descriptive property. Talk about properties being 'ascribed' is meant to get around this.
- 5 Sometimes it may be reasonable to use 'descriptive property' as described. The thesis that evaluative properties are descriptive properties still rules out theories like Moorean platonism and straightforward forms of non-cognitivism. But in conjunction with (T1) the given characterization of descriptive properties is unhappy.

On Jonathan Dancy's (1996) view on thick concepts, a thick concept is, somehow, associated with both a property and an attitude, but these are not in any way separable elements of the concepts since 'the property is best described as that of meriting the attitude, and the attitude is best characterized as the appropriate one given the presence of the property' (1996, 268). Dancy indicates that it is peculiar to thick concepts that they are associated with both a property and an attitude. But as already stressed, on all natural views even thin moral concepts are associated with properties. Moreover, consider the property associated with the concept *good*: the property of being good. It is far from obvious that this property is not best described as that of meriting the relevant attitude. Or take the attitude: it is far from obvious that this attitude is not best described as the attitude it is appropriate to adopt toward good things. Both these speculations can be doubted. A utilitarian, for instance, might think she has a better characterization of the property. But all I need in order to raise the problem is that it is as plausible that Dancy's characterization is satisfied by the concept expressed by 'good' as it is that it is satisfied by thick concepts.

Bernard Williams (1985, 29) says of thick concepts that they are 'world-guided,' meaning that their application 'is determined by what the world is like.' Thin concepts, by contrast, are supposed to have more of a purely action-guiding role. But is not also the application of the concept *good* determined by what the world is like? One might try to get around this objection by saying that the application of thick concepts is determined by what the world is like *in descriptive respects*. But this faces problems we have already seen: Specifically, why shouldn't also the concept *good* be like this?

A general lesson is that insofar as the 'thick' in thick concepts has to do with a relation to the descriptive, that relation is not best understood as one to descriptive *properties* (however we should conceive of this) but to descriptive *concepts*. For on a reasonable understanding of 'descriptive property,' paradigmatically thin concepts can be related to descriptive properties just as thick concepts are. (Sometimes in the literature, the thin/thick distinction is understood as a distinction between *properties* rather than a distinction between *concepts*. See e.g. Dancy (1993) and McNaughton and Rawling (2000). This seems unwise, in light of the present discussion.)

A first puzzle regarding thick concepts is then: what differentiates them from thin concepts?

II Puzzle (II): Seeming Sufficiency

Thick concepts have been appealed to in arguments that purport to show that sometimes evaluative sentences can be analytically entailed by descriptive sentences, contrary to a common belief (that no 'ought' can ever be analytically entailed by an 'is'). One well-known argument of this kind is Philippa Foot's (1958). In that article, Foot sets out to argue against what she calls the 'breakdown theory' in ethics, according to which disputes about the applicability of an ethical concept are always liable to breakdown, in that there is no objective way of settling ethical arguments.

Foot's example is the concept expressed by 'rude.' She notes that 'it expresses disapproval, is meant to be used when an action is to be discouraged, implies that other things being equal the behaviour to which it is applied will be avoided by the speaker, and so on' (1958, 102). The concept expressed by 'rude' is *evaluative*; specifically, it is negatively evaluative. But the concept also has descriptive conditions associated with it. In Foot's words, 'it can only be used where certain descriptions apply' (ibid). She goes on, 'The right account of the situation in which it is correct to say that a piece of behaviour is rude is, I think, that this kind of behaviour causes offence by indicating lack of respect' (ibid). The concept expressed by 'rude,' although an *evaluative* concept, is hardly strictly a *moral* concept. However, Foot wants to generalize from this example, and the philosophical interest in what Foot says about the concept expressed by 'rude' lies in this generalization to moral concepts.

Since Foot says that 'rude' can *only* be used where certain descriptions apply, it may be thought that her point is only that satisfaction of those descriptions is *necessary* for the applicability of the word 'rude.' But from the discussion that follows it is clear that she regards satisfaction of the descriptive conditions in question as *sufficient* for the applicability of 'rude':

Given that reference to offence is to be included in any account of the concept of rudeness, we may ask what the relation is between the assertion that these conditions of offence are fulfilled — let us call it O — and the statement that a piece of behaviour is rude — let us call it R. Can someone who accepts the proposition O (that this kind of offence is caused) deny the proposition R (that the behaviour is rude)? I should have thought that this was just what he could not do, for if he says that it is not rude, we shall stare, and ask him what sort of behavior is rude, and what is he to say? (1958, 103)

The concept expressed by 'rude' is then, on Foot's view, such that as a conceptual matter, anything that satisfies certain associated descriptive conditions falls under it; and since the concept is evaluative, we

can conclude that if something falls under it, then it is bad, and is to be disapproved of. This is what allows Foot to argue against the breakdown theory in ethics. If Foot's argument works, Foot also provides an initially promising solution to the puzzle of what distinguishes thick concepts from thin ones: if C is a thick concept then there is some substantive descriptive concept D such that 'x is C' follows analytically from 'x is D.' (The 'substantive' is there to rule out concepts like *is not self-identical*. Maybe important problems are glossed over here.) This is an account of thickness that relates thick concepts to descriptive concepts, not descriptive properties.

There is something intuitively attractive about Foot's outlook. But I think it is clear that it must be rejected.

Some thick concepts are, somehow, *objectionable*. Somehow these concepts presuppose or embody values that ought not really to be endorsed. Gibbard (1992) mentions *lewd* as an example: he does not agree on the — prudish — view of sexuality which underlies the employment of this concept. Graham Priest (1997) in effect argues that *sexually perverted* is an objectionable thick concept, whose usage presupposes that sexual behavior which does not fulfill a supposed natural purpose is thereby worthy of condemnation. While the examples can reasonably be doubted, I do not think that the general phenomenon of objectionable thick concepts can be.⁶

The very existence of objectionable thick concepts presents problems both for Foot's argument against the breakdown theory, and the corresponding account of what makes an evaluative concept thick. Take *lewd*. Suppose the associated descriptive concept to be *is an overt display of sexuality*. (No doubt this is much simplified!) Then, by Foot's reasoning, 'x is lewd' follows analytically from 'x is an overt display of sexuality'; and since *lewd* is a negatively evaluative concept, 'x is (pro tanto) bad' follows analytically from 'x is lewd.' Putting this together, we get that 'x is (pro tanto) bad' follows analytically from 'x is an overt display of sexuality.' Someone who happily accepts that descriptive sentences can analytically entail evaluative ones will not be put off by this being the kind of claim that it is. But when we are dealing with an objectionable thick concept, like *lewd*, this is plainly unacceptable. For if we think *lewd* is objectionable, we should also think that 'x is (pro tanto) bad' does not follow — let alone analytically follow — from 'x is an overt display of sexuality.'

6 If we include epithets among thick concepts — as some writers do; see Blackburn (1984 and 1992) and Hurka and Elstein (2009) — then it is even clearer that there are objectionable thick concepts.

Foot's argument goes wrong. However, there is something highly intuitive about what Foot is saying. How can this be accommodated, even while, in light of the problem of objectionable thick concepts, we cannot actually accept either Foot's argument against the breakdown theory or the corresponding account of what distinguishes thick concepts from thin ones? Call this the *puzzle of seeming sufficiency*.

I should immediately further justify calling this a *puzzle*. Not everyone may be on board with the claim that it intuitively seems as though there are analytically sufficient descriptive conditions associated with thick concepts.

Further consideration of what Foot says reinforces such doubts. Does 'causes offence by indicating lack of respect' mean *actually causes offence* or *would cause justifiable offence*? If the former, it arguably fails to be extensionally adequate. If the latter, then the condition is not descriptive. Coming up with a better proposal than that which Foot presents is far from straightforward. However, even if it is hard to come up with something appropriate that is suitably general, one can come up with plausibly analytically sufficient conditions. Just think of the description of some behavior that would be paradigmatically rude. The inference from that description to the conclusion that the behavior is rude may be intuitively analytic, even if general statements designed to capture what rude behaviors have in common are hard to come by.

A more general reason for doubts concerning the supposed puzzle stems from the idea, common in discussions of thick concepts, to say that the descriptive element cannot be separated out. (See e.g. McDowell (1979 and 1981), Williams (1985), and Dancy (1996).)

But first, it is one thing to say that one cannot *separate out* the descriptive component of a thick concept — the component that we get when, so to speak, conceptually subtracting the evaluative aspect — and another, considerably more radical, thing to insist that one cannot find descriptive concepts such that they have *significant analytic connections* to thick concepts. (Even if it is impossible to analyze the concept of knowledge in terms of simpler concepts, it can still be the case that knowledge analytically entails truth.)

Second, most importantly, the point I am making can be made without appeal to the notion of analytic sufficiency. The point is just that even if Foot's argument does not work, a satisfactory account of thick concepts should explain the intuitive attractiveness of the argument. Perhaps no notion of analyticity will occur *in the explanation*. So objections to the idea of analyticity, or to the possibility of analytically separating out a descriptive element, are beside the point.

In relation to the puzzle of seeming sufficiency, the following should also be noted. Thick concepts are, as mentioned, often characterized as merging descriptive and evaluative elements. But if we understand

thick concepts that way, many concepts — for example, concepts of the form *good and D*, for D descriptive, and *good qua F*, for F some sortal — count as thick which never make it onto the lists of thick concepts. These concerning thick concepts such as the claim that the descriptive element cannot be separated out are clearly false when it comes to concepts of the form *good and D*. And while an argument like Foot's has some initial plausibility when run using certain paradigmatic thick concepts, it would be hopeless if run using *good and D*. Compare too the concept *murder*. It somehow merges descriptive and evaluative elements, if it is at all correct to say that any thick concepts do this. But a corresponding argument involving *murder* would not carry the same intuitive force as Foot's original argument. It does not seem more intuitive that one can argue analytically from descriptive premises to the conclusion that there was a murder than that one can argue from descriptive premises analytically to the conclusion that there was something which was a killing and which was wrong that was done.⁷ A lesson to draw is that there are subclasses within the class of thick concepts. A fully general account of thick concepts may be a chimera. I will continue focusing on the subclass to which *lewd* and *rude* belong rather than the subclass to which *murder* belongs.⁸ Consideration of concepts like *murder* helps reinforce the second puzzle, the puzzle of seeming sufficiency. It might have seemed attractive that the difference between *lewd* and *rude* on the one hand and *murder* on the other is that the former are associated with analytically sufficient conditions while the latter is not. The discussion of Foot's account shows that this suggestion will not work. What else, in the vicinity, might do?

III Puzzle (III): Emptiness?

After having given the argument I summarized above, Foot somewhat qualifies her position, as follows:

It is of course possible to admit O without asserting R.... Calling an action 'rude' is a concept which a man might want to reject, rejecting the whole practice or prais-

-
- 7 Note the formulation. I am not saying that one cannot argue from descriptive premises analytically to 'x is murder' or 'x is wrong.' All I am saying it seems equally hard in both cases; nothing is gained, argumentatively, by considering *murder* instead of *wrong*.
- 8 It may be suggested that *murder*, although a value concept with some descriptive content, is not properly described as a *thick* concept. The issue seems merely terminological.

ing and blaming embodied in such terms as 'polite' and 'rude.' Such a man would refuse to discuss matters of etiquette, and arguments with him about what is rude would not so much break down as never begin.

The only recourse of the man who refused to accept the things which counted in favour of a moral proposition as giving him reason to do certain things or take up a particular attitude, would be to....abjure altogether the use of moral terms. (1958, 104)

This is — as Marvin Glass (1973) put it — a kind of lack of nerve on Foot's part (even if, given the above arguments, the lack of nerve is appropriate). As mentioned above, Foot's purpose in bringing up the argument concerning *rude* is to argue against the so-called breakdown theory. But Foot here in effect introduces a 'second-order' breakdown theory: this time the argument will concern whether or not to engage in a particular practice of praising and blaming.

There is a more damaging point to be made. If 'x is D,' for some descriptive concept D (say, *the conditions of offence are fulfilled*), really analytically entails 'x is (pro tanto) bad' via analytically entailing 'x is C' for some thick concept C (perhaps *rude*), my refusal to use the relevant thick concept only amounts to my refusing to give expression to certain truths. The *facts* don't change!

Several other theorists beside Foot who acknowledge the existence of objectionable thick concepts also hold that we simply should not employ them.⁹ Compare Oscar Wilde on blasphemy: 'the word 'blasphemous' is not a word of mine' (Foldy, 1997, 8). But this piece of advice, even if sound, leaves questions about the *extensions* of the concepts unresolved, as just stressed. Specifically, can objectionable thick concepts still truly apply to things, or are they empty?

One cannot happily say that they are empty. If it is said about *lewd* or *sexually perverted*, just think about behaviors that would be regarded as paradigmatically lewd or paradigmatically sexually perverted. (Modesty prevents bringing up really persuasive examples.) So it seems that they are not empty. But there is a problem with this. Suppose I find, say, *lewd* objectionable, while still I must agree that X falls under 'lewd.' I am committed to the truth of 'X is lewd,' to its being the case that X is lewd. The problem is that it may seem that I am then committed to holding that X is bad on the ground of being an overt display of sexuality, just because of what 'lewd' means. But precisely because I find *lewd* objectionable, I will surely want to distance myself from any claim to that effect.

9 See, e.g., Priest (1997) and Slote (1975).

My view is that objectionable thick concepts are not empty. Start by considering a slightly different case: epithets. It is at best unclear whether epithets, including racial epithets, should be classed as thick concepts. (Indeed, I will later present problems for this idea.) But comparing the case of epithets may anyway be instructive. One view on this matter is that when the racist calls a black person 'nigger,' that is to be compared to the racist's calling this person 'black' using a negative or contemptuous tone of voice. What is wrong with the latter utterance does not lie in what it is said, but in what is conveyed through how it is said. On this view on epithets, the non-racist should agree that 'nigger' is actually *true of* black people.¹⁰ One thing that favors this view on epithets is that epithets convey the attitudes with which they are associated even when occurring embedded.

While there intuitively are differences between epithets and thick concepts, I hold that in the relevant dimension there is an analogy. If C is an objectionable thick concept, then its use is objectionable not only when it is claimed of someone or something that it is C: its use is objectionable in the same way when it occurs in embedded sentences.¹¹ If I find 'lewd' objectionable I will find an ordinary utterance of 'Mick isn't lewd' as problematic as I will find an utterance of 'Mick is lewd.' (An exception is provided by metalinguistic negation: I can say 'Mick is not *lewd* — rather, his behavior is *sexually provocative*.' But this corresponds to 'Oliver Hardy isn't *fat* — he's *overweight*.' The negation is metalinguistic. It is the wording that is rejected.) Such embedding behavior is not neatly explained by any hypothesis about the relation between the negative evaluation in 'lewd' and its truth-conditional content that would yield that 'lewd' is empty. As an illustration, consider the simplest hypothesis under which 'lewd' is empty if objectionable: 'x is lewd' means something like 'x is D and blameworthy,' where D is some descriptive concept. 'x is not lewd' then means: 'it is not the case that x is D and blameworthy.' But someone finding 'lewd' objectionable should not find fault with this.

Compare again the Oscar Wilde-style policy on the use of objectionable value words. The policy is natural. But if what was objectionable about the value words was a matter of truth-conditional content, then one could well use the words even if one finds them objectionable: it is just that one would not assert of anything that one of these words applies to it.

10 For interesting criticism of this view on epithets, see Richard (2008).

11 See Väyrynen (2009) for discussion.

If objectionable thick concepts are not empty, they present counterexamples to the above characterization of what it is for something to be a positive (negative) value concept. For something C can then be a positive (negative) value concept even if 'x is (pro tanto) good' ('x is pro tanto bad') does not follow analytically from 'x is C.'

Objectionable thick concepts present problems for a number of otherwise rather different accounts of thick concepts. Take a simple conjunctive account on which a thick concept is associated with a descriptive and a normative condition and is true of an object exactly when both conditions are satisfied. E.g., a negative thick concept would have an analysis of the form 'is D and (pro tanto) bad.' On this account it is unclear how a thick concept could be objectionable, for it is unclear how thick concepts would presuppose evaluative claims. More plausibly it can be suggested that the analysis of a positive (negative) thick concept has the form 'is good (bad) in virtue of being D.' Given this sort of account it can be said that a thick concept somehow presupposes that goodness and being D are connected, and a thick concept is objectionable when this is not true. But this sort of account entails that all objectionable thick concepts are empty.

Turning to some specific proposals found in the literature, consider Christine Tappolet's (2004) account of thick concepts as determinates of thin concepts. On Tappolet's view, positive thick concepts are determinates of the determinable (*pro tanto*) good; negative thick concepts are determinates of the determinable (*pro tanto*) bad.¹²

According to Tappolet's account, if C is a negative thick concept, then, necessarily if something is C then it is (pro tanto) bad. But this runs into trouble with objectionable thick concepts. Stick with *lewd* as our example of an objectionable thick concept. Some things that would be regarded as paradigmatically lewd are indeed lewd. But if 'lewd' is objectionable, not all of these things need be bad, even pro tanto. We then have a counterexample to Tappolet's account. Tappolet is committed to taking objectionable thick concepts to be empty.¹³

Earlier I presented problems for some accounts based on the idea that what distinguishes thick concepts from thin ones is that the thick

12 This account in effect rejects the assumption that thick concepts are more descriptive than thin ones. A determinable/determinate distinction does not immediately go hand in hand with a non-descriptive/descriptive distinction. A worry about the account that I will set aside is that the determinable/determinate distinction is better conceived of as a distinction between properties than between concepts.

13 Although Tappolet considers locutions of the form 'too F' where the term 'F' expresses a positive thick concept (2004, 209f), she does not consider the worry concerning objectionable thick concepts.

concepts have a descriptive element that thin concepts lack. Another type of informal characterization of what distinguishes thick concepts is that they are more *specific* than the more general thin concepts. Objectionable but non-empty thick concepts present a *prima facie* problem for such an informal characterization, for given such concepts it is not generally the case that for C a positive (negative) objectionable thick concept, the extension of C is a subset of the concept of being *pro tanto* good (bad).

Consider next the account Gibbard ends up defending in his (1992), stated for the special case of 'lewd.' He first says that to 'use the term with linguistic propriety,' a speaker must 'accept the general importance of limiting sexual display' (1992, 280). More specifically, such a speaker 'must think that there are certain limits on sexual display, such that passing beyond these limits warrants feelings of L-censoriousness toward the person who does so' (*ibid.*), where 'L-censoriousness' is 'the feeling of outraged shock and censure that goes with finding something lewd' (1992, 279). This talk of using 'lewd' with linguistic propriety is motivated by concerns having to do with the fact that speakers might find certain thick concepts objectionable.

Given this as background, Gibbard then states his account:

'Act X is lewd' means this: L-censoriousness toward the agent is warranted, for passing beyond those limits on sexual display such that (i) in general, passing beyond those limits warrants feelings of L-censoriousness toward the person doing so, and (ii) this holds on no further grounds or on grounds that apply specifically to sexual displays as sexual displays. (1992, 280f)

(Gibbard is of course a non-cognitivist, but I do not think his non-cognitivism matters to the evaluation of his proposal regarding thick concepts.)

One may regard the above as a statement of the truth conditions of sentences of the form 'act X is lewd.' If so, it immediately runs into problems with objectionable thick concepts — oddly enough, given that Gibbard himself calls attention to the problem. If the truth is that L-censoriousness is never warranted because of an agent's going beyond certain limits of sexual display, then it follows that nothing is lewd. Gibbard is committed to the emptiness of objectionable thick concepts.

The emphasis on which speakers may use the term with linguistic propriety clouds the issue. The idea is somehow if I do not satisfy the condition in question, I cannot use 'lewd' appropriately. In some sense this is surely right: my use of it would serve to convey acceptance of a certain view on sexual morality that I do not accept, so my use of it would be *misleading*. But does Gibbard mean to make a claim that goes beyond this, to the effect that my utterances would not even be *truth-evaluable*? He is not explicit. (Maybe the reason he does not get into

the issue is that given the non-cognitivism he subscribed to, questions about truth-evaluability are moot.) He only stresses that someone finding the term objectionable ought not to use it. But this just sidesteps the question of the truth of sentences containing 'lewd.'

I call this puzzle *the puzzle of emptiness*. Already I have argued that objectionable thick concepts can be non-empty, so the puzzle I have in mind is not whether objectionable thick concepts are empty or not. Rather, what I call the puzzle of emptiness is a collection of issues related to the non-emptiness of thick concepts. First, it is incumbent on a satisfactory account of thick concepts that it allow that thick concepts can be objectionable yet non-empty. I have already given examples of accounts which do not meet this condition. Second, even if objectionable concepts can be non-empty it is, as earlier noted, awkward to say something of the form '___ is C' for someone who finds C objectionable. A satisfactory account of thick concepts should explain this awkwardness.

IV Deflationary Accounts of Thick Concepts

One response to the problem of what is special about evaluative thick concepts is radical: to deny that the concepts regarded as thick properly speaking are evaluative. There are two different ways this can be done. There is first a radical suggestion. Some concepts are used evaluatively without themselves being evaluative. Maybe *athletic* is a case in point. It is often used positively evaluatively; but it does not seem really to be a value concept. (Specific examples can always be criticized. Perhaps *athletic* is in fact an evaluative concept. But even if *athletic* is not a good example, the general point should be clear.) It may be held that the concepts commonly regarded as thick are like the concept *athletic*. One strand in Blackburn (1992) points in this direction. Blackburn notes how we take utterances like 'last year's carnival wasn't lewd enough' in our stride (1992, 296).¹⁴ This observation can be taken to show that the evaluativeness of ordinary uses of 'lewd' is shifty in a way that does not sit well with 'lewd' being an evaluative concept. Second, there is a more moderate suggestion. It can be held that although the terms

14 Compare too the discussion in Blackburn (1998, 103f). Blackburn's example sentence is about what is not 'lewd enough,' and one may want to seek to explain the case by appeal to what 'enough' contributes. But where Blackburn has an envisaged speaker complain that last year's carnival was not 'lewd enough,' he could equally well have someone complaining that there was an essential element missing from last year's carnival — it 'simply was not lewd.'

ordinarily regarded as expressing thick concepts are evaluative as part of their meaning, their evaluativeness is a matter of what in discussions of Frege is sometimes called *tone*, rather than *sense*; a related suggestion would be that when C is a positive (negative) thick concept, it is conventionally or conversationally *implicated*, rather than said, that x is good (bad) when 'x is C' is assertively uttered. I will start by discussing the more radical suggestion.

First, it is well known that there are similar examples involving terms expressing thin concepts. ('Evil, be thou my good.') Anyone justifying the radical deflationary line with respect to thick concepts by appeal to examples like Blackburn's carnival-example must either conclude that not even the concepts commonly regarded as thin evaluative concepts are in fact evaluative, or else find a disanalogy between the examples.

Second, why, exactly, should we say that 'athletic' does not express an evaluative concept even if it is often used evaluatively? (Assuming that 'athletic' indeed is a good example to use.) This is a difficult question, but the following thought experiment provides a hint. Take first a paradigmatically evaluative concept such as *good*. Compare two hypothetical linguistic communities, G1 and G2. Both speak languages exactly like English, except for the following possible differences (and differences immediately related to them): In G1, they systematically apply 'good' to pretty much what we apply 'good' to, but they do not use it to praise — they use it to condemn, or they use it neutrally. In G2, they use 'good' to praise in the manner we do, but they systematically apply it to different things. Now ask: which community — if either — uses 'good' with the same meaning that we use it with? The answer seems clearly to be G2. Now turn to 'athletic' — our example of a concept that, although often used to evaluate, is not evaluative — and consider two hypothetical linguistic communities A1 and A2. Again both speak languages exactly like English except for the possible differences to be noted (and differences immediately related to them). In A1, they apply 'athletic' to pretty much what we apply 'athletic' to, but they do not use it to praise — but either to condemn, or perhaps simply neutrally. In A2, they use 'athletic' to praise in much the way we do, but systematically apply 'athletic' to different kinds of things than the ones we apply 'athletic' to. The community that uses 'athletic' with the same meaning as that we use 'athletic' with seems to be A1 rather than A2.¹⁵

15 Of course the thought experiment is very briefly described, and more can be said. One qualification is immediately in order: by 'systematically,' I do not mean simply *regularly*; rather, I mean that the applications are ones that are stable, in that these are the applications they tend to converge on as more information about the different cases is acquired.

Any argument based on thought experiments like these is bound to be speculative, especially when the relevant scenarios are as briefly described as here. But consider thick concepts. If a thick concept like that expressed by 'lewd' is both descriptive and evaluative, as we have been supposing, then we should expect that both a community which uses 'lewd' evaluatively differently from how we use 'lewd' (compare G1/'good' and A1/'athletic') and a community which uses 'lewd' descriptively differently from how we use 'lewd' (compare G2/'good' and A2/'athletic') use 'lewd' with a meaning different from that we use it with. But if a radical deflationary account is true, we should expect 'lewd' rather to be like 'athletic.' It appears the former hypothesis is more plausible.

Turn now to the moderate deflationary suggestion. Compare again epithets. As mentioned, one common view is that epithets are necessarily coextensive with their neutral counterparts and only differ from them semantically in ways that do not affect truth-value. There are different more specific views on epithets, but one natural view is that epithets differ from their neutral counterparts through what they *conventionally implicate*. Maybe we should conceive of the evaluative element in thick terms in the same way.¹⁶

This hypothesis would deal nicely with the question of the non-emptiness of thick terms. If thick terms have the same truth-conditional contents as their neutral counterparts, they are true of exactly what their neutral counterparts are true of; but what is conventionally implicated by the use of an objectionable thick term is false. The hypothesis also deals nicely with the problem of seeming sufficiency. Given the hypothesis there will be purely descriptive concepts such that if they are true, the thick term truly applies. It is only that the applicability of these descriptive concepts does not suffice for the truth of what is conventionally implicated. Turning to the first of our puzzles, concerning what characterizes thick terms, a proponent of the view under consideration can say that they merge a descriptive sense with an evaluative implicature.

But while the account deals nicely with the second and third puzzles, the account of the first puzzle is unsatisfactory. First, epithets come out as thick terms on this view: but one may think that there are clear differences between epithets and thick concepts, and an account which treats these classes of terms the same way is unsatisfactory.¹⁷ Second, if this

16 As mentioned earlier, these claims about epithets can be denied; see e.g. Richard (2008).

17 As mentioned in fn. 6, some theorists take epithets to be examples of thick terms.

is not immediately persuasive — some authors are after all happy to run together pejoratives and thick terms — consider paradigmatic *thin* terms. Either the evaluativeness of thin terms is a matter of their senses and not of what their use implicates, or also the evaluativeness of thin terms is a matter of (conversational or conventional) implicature, as authors like Stephen Barker (2000), David Copp (2001) and Stephen Finlay (2005) have suggested. In the former case, there is a dramatic difference between the evaluativeness of thin terms and the evaluativeness of thick terms: when it comes to thin moral terms the evaluation is a matter of their sense but when it comes to thick ones the evaluation is a matter of their tone. But as Samuel Scheffler noted in his review of Williams (1985), we do not have here a simple distinction between two kinds of terms but rather a spectrum of cases: for instance, is ‘just’ a thin term or a thick one? (1987, 417) Suppose then that also the evaluativeness of thin terms is a matter of implicature. Then if thin terms have senses, these senses are purely descriptive. But if so, thin terms come out as thick on the present account: they merge descriptive senses with evaluative implicatures. Someone defending this account of thick terms can get around the problem by denying that thin terms genuinely have senses; their use only serves to convey something about the speaker’s attitude. This would be a form of non-cognitivism. Discussing the merits of non-cognitivism would take us too far afield. Suffice it for here to say this. Thick terms have traditionally been held to pose problems for non-cognitivism. It would then be ironic if only non-cognitivists have the tools to say what is special about thick concepts. (This is not in itself an objection to the Barker/Copp/Finlay claim that evaluativeness of thin concepts is a matter of implicature. The point is just that if Barker, Copp and Finlay are right, then it has not been accurately explained wherein the thickness of a thick concept consists.)

V The Conception View

I now turn a different type of positive account, what I will call *the conception view*. The conception view faces immediate objections. But I think it points the way to what a better account might be like. In the next section I will present my positive account, one designed to get around the problems faced by the conception view.

Let the *conception* associated with a concept be a collection of widely held and firmly entrenched beliefs involving the concept. Of course some beliefs that are parts of conceptions may be false. Both descriptive and causal theories of reference allow conceptions a role in reference-determination. Contemporary descriptivists like David Lewis and Frank Jackson take terms to be implicitly defined by associated

conceptions.¹⁸ Causal theorists take the relation between conceptions and reference to be looser, but still allow conceptions a role, for example by saying that terms refer to what causally regulates the conceptions.

Given this notion of a conception, we can attempt to say the following regarding thick concepts: What makes thick concepts thick is that both beliefs linking them to substantive descriptive concepts and beliefs linking them to evaluative concepts are part of the conception. That is the answer to the first puzzle. As for the second puzzle, seeming sufficiency, it can be suggested that what we mistake for an analytic entailment from a descriptive claim to an evaluative claim is just a widely held and firmly entrenched belief to the effect that the descriptive claim entails the evaluative claim. When it comes to the third puzzle one of two things can be said. On a causal theory the reference of a thick concept is what is causally related in the appropriate way to the conception — say, causes, in the right way, the beliefs that are part of the conception. On a neo-descriptivist theory it is what comes closest to satisfying the conception (in the best case it satisfies the conception, but even something falling short of fully satisfying the conception might be the referent of the concept — it might deserve the label even if it is an imperfect deserver of the label).¹⁹ Either way, we can explain how objectionable thick concepts can be non-empty. Now, one of the problems in saying that objectionable thick concepts are non-empty was that saying ‘*x* is *C*,’ where *C* is a negative thick concept, seems to commit one to the truth of ‘*x* is bad.’ But the view we are considering can perhaps explain this. It can be said that if it is part of our conception of *F*s that all *F*s are *G*s, then it is only to be expected that saying of something that it is *F* can serve to convey that this something is *G*.

I think the conception view is on the right track. It provides a relatively neat solution to the third puzzle. And what it says about the second puzzle is fairly attractive, though one may think that the difference between being *analytic* and being *part of the conception* is dramatic enough that an account of how we could possibly systematically conflate the two would be desirable. But what it says about the first puzzle is obviously far from satisfactory. The given characterization of thick concepts is satisfied by many concepts that are not thick. For example, thin concepts can satisfy the characterization. For surely there can be widely held and firmly entrenched beliefs of the form *if x is D then x is good* or *if x is good then x is D*, for some substantive descriptive concept *D*, and then the concept *good* would come out thick.

18 See Jackson (1998), and Lewis, e.g. (1997).

19 See e.g. Lewis (1997) for discussion of the last of these options.

The conception view cannot be exactly right. But it indicates what a better positive view might be like. Suppose we can find a condition on the relation between concepts and claims that is stronger than *the claim is part of the conception associated with the concept* (for our discussion of the conception view suggests that something stronger is needed) yet weaker than *the claim is an analytic truth about the concept* (for basing a characterization of thick concepts on claims that satisfy this condition threatens to make objectionable thick concepts come out empty). To put things intuitively: what we are looking for is a sense in which it can be constitutive of — or ‘part of’ — a concept that a certain claim should come out true, even while the claim in question can fail to be analytically true, or even true at all. If we can find such a condition, then we can attempt to characterize thick concepts as those concepts that are related by claims that satisfy this condition both to evaluative concepts and to substantive descriptive concepts. A positive view like the conception view but centered on this relation promises to build on what seems right about the conception view. In the next section, I will describe a candidate for being this relation, familiar from elsewhere in the literature.

VI Epistemic Analyticity

As authors like Paul Boghossian (see especially his 1996 and 2003) and Jamie Tappenden (1993a and 1993b) have noted, there are several distinct ideas underlying the traditional notion of analyticity. There is the *metaphysical* idea of truth in virtue of meaning. And then there is the *epistemological* idea of — roughly — something’s being a sentence such that it is part of competence with the expressions involved to be disposed to accept it. We can thus distinguish two notions of analyticity, one metaphysical and one epistemological. Focus on the epistemological idea. It is possible, at least in principle, that a sentence should be analytic in the epistemological sense without being true at all.

If it is indeed possible that sentences can be epistemically analytic without being true, this can provide an important part of the solution to many philosophical puzzles. Consider for example the liar and sorites paradoxes. It is widely agreed that the liar paradox shows that not all instances of the schema ‘*s* is true iff *p*’ (the *T-schema*), where for ‘*p*’ a sentence is substituted and for ‘*s*’ a name of this sentence is substituted, are true. But it seems somehow meaning-constitutive for ‘true’ that this schema should be valid. We can resolve the conflict by saying that the *T-schema*, or the instances thereof, are epistemically analytic without being true. The sorites paradox appears to show that a principle like ‘if *x* and *y* are indistinguishable with respect to color to casual observation

then if one looks red then so does the other' cannot be true. Yet it seems to follow from the fact that 'looks red' is an observational color predicate that this principle should be true. We can resolve the conflict by saying that the relevant principle, and analogous principles (so-called *tolerance principles*) for other vague expressions are epistemically analytic without being true.²⁰

An immediate objection to these suggestions is that it seems clearly possible to deny instances of the T-schema, or tolerance principles, without displaying incompetence. For example, experts on the liar and sorites paradoxes routinely deny these things! This observation might be generalized. One might think that for any purported untrue epistemic analyticity, if someone can recognize that it is untrue she can thereby come to reject it without manifesting incompetence. The most important response to this objection is the following. What underlies the idea that there is a non-trivial notion of epistemic analyticity is that there are sentences such that competence with the expressions involved entails standing in some distinguished cognitive relation R to these sentences. It is common to identify R as *believing* or *accepting*, or *being disposed to believe* or *accept*. When introducing epistemic analyticity I have followed suit. But already independent considerations show that such an identification is too simple-minded. If anything has the status of an epistemic analyticity, the sentences expressing basic logical laws surely do. But for pretty much any basic logical law, one can find logicians and philosophers denying it.²¹ (Call this the phenomenon of *radicalism*.) Hence, the identification of R is somewhat trickier than might have been expected. And given that the simple-minded suggestions regarding R do not work, the path to untrue epistemic analyticities lies open, where a sentence is epistemically analytic just in case full competence with the constituent expressions entails standing in relation R to this sentence. The phenomenon of radicalism indicates that R does not entail belief or acceptance.

What else might R be? Any discussion here will perforce be brief. But one suggestion is that R is something like *accepting as the default position that [...]*. It is striking that counterexamples to taking R to be something like (dispositions to) belief or acceptance always involves someone having a *special reason* to deny the supposed analyticity. Timothy Williamson has prominently called attention to Vann McGee's rejection of modus ponens when criticizing the notion of epistemic analyticity, and he has also called attention to how speakers can for theoretical reasons

20 See Eklund (2002) and also (2005).

21 See Williamson (2003), (2006) and (2007).

reject instances of the schema ‘every F is an F.’ But strikingly, the examples involve thinkers armed with specific reasons to reject the epistemic analyticities. If one imagines someone rejecting modus ponens or instances of ‘every F is an F’ without special reason to do so, it becomes considerably more doubtful that this thinker — social externalism aside — uses these terms with their standard meanings.²²

Of course, other reactions to the phenomenon of radicalism are possible. One might draw the conclusion that any philosophy of language that demands that there are sentences such that competence entails that the speaker stands in distinguished cognitive relation R to them must be abandoned.²³ That is to say, one can simply reject the notion of epistemic analyticity. Or one might think that even a ‘radical’ philosopher or logician has the disposition to accept the relevant sentence. It is just that she has a mistaken *theory* about her own competence. (Compare undergraduates with beliefs to the effect that ‘or’ as they use it is exclusive. It can still plausibly be argued that they still implicitly accept, or are disposed to accept, the usual rules for inclusive ‘or.’) If this is the right thing to say about a ‘radical’ philosopher or logician, I do not see that it is any less plausible to say this concerning someone who denies a tolerance principle for a vague predicate or an instance of the T-schema.

Suppose that there can be epistemic analyticities known to be untrue. This would provide the materials for a nice resolution to the problems of thick concepts that we have been discussing. Foot’s (1958) discussion in effect suggested an account of thick concepts according to which C is a thick concept if and only if

- (i) there is some substantive descriptive concept D such that ‘x is D’ analytically entails ‘x is C’;
- (ii) if C is positive (negative), ‘x is C’ analytically entails ‘x is (pro tanto) good’ (‘x is (pro tanto) bad’).

Such an account does not work, as shown for example by the discussion of objectionable thick concepts. But consider this account of thick concepts modified, so that ‘analytically’ in (i) and (ii) is understood only to mean *epistemically* analytically:

- (i’) there is some substantive descriptive concept D such that ‘x is D’ epistemically analytically entails ‘x is C’;

22 For discussion see Williamson (2007), and Eklund (2007) and (forthcoming)

23 In effect, this is how Williamson reasons in e.g. his (2007).

- (ii) if C is positive (negative), 'x is C' epistemically analytically entails 'x is (pro tanto) good' ('x is (pro tanto) bad').

Intuitively: we are only saying that it is *somehow part of the meaning* of C that 'x is C' is entailed by 'x is D,' and that it is *somehow part of the meaning* of C that 'x is C' entails 'x is (pro tanto) good' ('x is (pro tanto) bad'); we are not saying that 'x is C' *actually is entailed by* 'x is D' and *actually entails* 'x is (pro tanto) good' ('x is (pro tanto) bad').

Recall now our three puzzles. The first concerned how to account for the thickness, or descriptiveness, of thick concepts. The second was how to accommodate the intuitiveness of the type of account suggested by Foot's discussion. The third concerned the non-emptiness of objectionable thick concepts.

It should be clear how the first puzzle is dealt with: (i') and (ii') are supposed to characterize thick concepts. The answer to the second puzzle, given the present suggestion, is straightforward, for what is defended is a straightforward variant of the account suggested by Foot's discussion. With respect to the third puzzle, the friend of epistemic analyticity can simply take over what the friend of the conception view says. First, concerning how objectionable thick concepts need not, given the present suggestion, be empty; this is straightforward. The defender of the idea that there can be untrue epistemic analyticities can just take over neo-descriptivism's account of reference-determination. Naturally, problematic questions can be raised with respect to this account. For example, the notion of closeness that is appealed to when it is said that the referent is the closest satisfier can be problematized.²⁴ But the appeal to epistemic analyticity does not give rise to any *new* problems. Second, when it comes to the seeming oddity of saying of an objectionable thick concept that it is not empty, one may suggest that if it is epistemically analytic that all Fs are Gs, then if someone says that something is F she thereby suggests that it is G.²⁵

24 For further discussion of the appeal to the notion of closeness, see Eklund (2005, 50ff).

25 There has recently been some discussion of the example 'Boche,' originally discussed in Dummett (1973). Some authors (see e.g. Brandom (2000) and Boghossian (2003)) have, following Dummett, taken 'Boche' to express a defective concept, somehow governed by the rules

From 'x is German' to infer: 'x is Boche.'
From 'x is Boche' to infer 'x is cruel,'

which are defective qua concept-constituting rules since it is not true (let alone analytically true) that all Germans are cruel. Critics like Williamson (see 2003 and

Neither the present account nor the conception view immediately positively entails that objectionable thick concepts are non-empty. It could be that there is no closest and good enough satisfier, so an objectionable thick concept has an empty extension. What these accounts do, though, is to explain how objectionable thick concepts could be non-empty: they are non-empty if there are non-empty extensions that are the closest (and good enough) satisfiers.²⁶

Of course, those who find a particular positive (negative) thick concept *C* objectionable will not find themselves disposed to reason from 'x is *C*' to 'x is (pro tanto) good' ('x is (pro tanto) bad'). But one point of the excursion into general problems that can be raised for epistemic analyticity views is that similar problems arise for the idea of epistemic analyticity more generally, and that the condition for being epistemically analytic might anyway have to be seen as something more complex.

So the account of thick concepts based on (i') and (ii') deals nicely with the puzzles brought up. *If* adequate sense can be made of epistemic analyticity and it can be made plausible that there are rich epistemically analytic connections between concepts, it must be said that the view is attractive. Of course, absent a more compelling characterization of epistemic analyticity, many will be skeptical of appeal to the notion. However, it should be stressed that if one cannot deal with the three puzzles without appeal to the notion, that itself is indirect evidence that adequate sense should in principle be able to be made of it.²⁷

2009) have said that the evaluativeness of 'Boche' is a matter of what its use conventionally implicates, and not of its *sense*.

It seems to me that much of this debate is marred by the fact that 'Boche' is not a terribly good example for the Dummettian side to use. Objectionable thick concepts would make for better examples. Indeed, the positive account of thick concepts that I will go on to propose is a close relative of what Dummett proposes in the case of 'Boche.'

- 26 One possibility that suggests itself on either account is that it is *indeterminate* whether the objectionable thick concepts are empty. An assignment of semantic values on which they are does equally well in terms of the desiderata on a theory of concepts as a theory on which they are not.

I have earlier argued that objectionable thick concepts are non-empty. But already the indeterminacy hypothesis is enough to pose problems for those accounts that entail that objectionable thick concepts are empty.

- 27 The discussion here is obviously of relevance to the debate over *Hume's Law*, the thesis that an evaluative conclusion cannot be validly inferred from non-evaluative premises. The 'law' amounts to different things depending on how we understand the inference. If Hume's Law is understood to say that an evaluative conclusion cannot be *necessitated by* non-evaluative premises, then it is to say the least far from obvious that it is true. The same goes if it is understood to say that an evaluative conclusion cannot follow *a priori* from non-evaluative premises. Typically it

There is a more general point in the vicinity. Suppose one thinks of content in a purely *referentialist* way — the content of an expression simply is what it refers to. What should one then say distinguishes evaluative terms? There are essentially two options. One is to say that evaluative terms are distinguished by what they refer to; the other is to say that evaluative terms are distinguished by something which is not strictly part of their content, for example by what their use conventionally implicates. Both these options are problematic. Considerations from earlier on present problems for the first option: an evaluative predicate can in principle be coextensive with a descriptive one. And the idea of appealing to conventional implicature was in effect criticized in section four. In response to these problems, a referentialist can seek to deny that any terms or concepts are in and of themselves evaluative: we tend to use some terms and concepts for evaluative purposes, but that is all. But a different reaction is to hold that in order to capture the fact that some terms and concepts are evaluative, one must add richer conceptual structure: reference alone is not sufficient. And to add conceptual structure is, among other things, to add relations between concepts, such as analytic connections.

Received: April 2009

Revised: December 2009

Revised: March 2010

References

Barker, S.J. 2000. 'Is Value Content a Component of Conventional Implicature?' *Analysis* 60: 268-79.

Blackburn, S. 1984. *Spreading the Word*. Oxford: Clarendon Press.

_____. 1992. 'Morality and Thick Concepts II — Through Thick and Thin.' *Proceedings of the Aristotelian Society* Suppl. Vol. 66: 267-83.

is rather understood to mean that an evaluative conclusion cannot follow *analytically* from non-evaluative premises. Understood thus, Hume's Law sounds more likely true. Indeed, when early on I gave a tentative characterization of what distinguishes evaluative expressions, I in effect relied on Hume's Law conceived of this way.

Distinguishing between different notions of analyticity gives rise to different forms of Hume's Law. Formulated using a notion of epistemic analyticity, Hume's Law is false given the present proposal regarding thick concepts. It is consistent with this that Hume's Law is correct if stated using a notion of metaphysical analyticity.

- _____. 1998. *Ruling Passions*. Oxford: Clarendon Press.
- Boghossian, P. 1996. 'Analyticity Reconsidered.' *Noûs* 30: 360-91.
- _____. 2003. 'Blind Reasoning.' *Proceedings of the Aristotelian Society*, Suppl. Vol. 78: 225-48.
- Brandom, R. 2000. *Articulating Reasons*. Cambridge, MA: Harvard University Press.
- Copp, D. 2001. 'Realist-Expressivism: A Neglected Option for Moral Realism.' *Social Philosophy and Policy* 18: 1-43.
- Dancy, J. 1993. *Moral Reasons*. Oxford: Blackwell.
- _____. 1996. 'In Defense of Thick Concepts.' In *Moral Concepts*, Midwest Studies in Philosophy 20, P.A. French, T.E. Uehling Jr., and H.K. Wettstein, eds. Notre Dame, Indiana: University of Notre Dame Press.
- Dummett, M. 1981. *Frege: Philosophy of Language*. 2nd ed. Cambridge, MA: Harvard University Press.
- Eklund, M. 2002. 'Inconsistent Languages.' *Philosophy and Phenomenological Research* 64: 251-75.
- _____. 2005. 'What Vagueness Consists In.' *Philosophical Studies* 125: 27-60.
- _____. 2007. 'Meaning-Constitutivity.' *Inquiry* 50: 559-74.
- _____. 2010. 'Review of Timothy Williamson, *The Philosophy of Philosophy*.' *Australasian Journal of Philosophy* 88: 752-4.
- Finlay, S. 2005. 'Value and Implicature.' *Philosophers' Imprint* 5: 1-20.
- Foldy, M. 1997. *The Trials of Oscar Wilde*. New Haven and London: Yale University Press.
- Foot, P. 1958. 'Moral Arguments.' *Mind* 67: 502-13. Reprinted in *Virtues and Vices*, Oxford: Blackwell 1978.
- Glass, M. 1973. 'Philippa Foot's Naturalism: A New Version of the Breakdown Theory in Ethics.' *Mind* 82: 417-20.
- Hurka, T. and D. Elstein. 2009. 'From Thick to Thin: Two Moral Reduction Plans.' *Canadian Journal of Philosophy* 39: 515-35.
- Jackson, F. 1998. *From Metaphysics to Ethics*. Oxford: Clarendon Press.
- Lewis, D. 1997. 'Naming the Colours.' *Australasian Journal of Philosophy* 75: 325-42.
- McDowell, J. 1979. 'Virtue and Reason.' *Monist* 62: 331-50.
- McDowell, J. 1981. 'Non-Cognitivism and Rule-Following.' In *Wittgenstein: To Follow a Rule*, S. Holtzman and C. Leich, eds. London: Routledge & Kegan Paul.
- McNaughton, D. and P. Rawling. 2000. 'Unprincipled Ethics.' In *Moral Particularism*, B. Hooker and M. Little, eds. Oxford: Oxford University Press.
- Priest, G. 1997. 'Sexual Perversion.' *Australasian Journal of Philosophy* 75: 360-72.
- Richard, M. 2008. *When Truth Gives Out*. Oxford: Oxford University Press.
- Scheffler, S. 1987. 'Morality Through Thick and Thin: A Critical Notice of *Ethics and the Limits of Philosophy*.' *Philosophical Review* 96: 411-34.

- Slote, M. 1975. 'Inapplicable Concepts and Sexual Perversion.' In *Philosophy & Sex*, R. Baker and F. Elliston, eds. Buffalo, NY: Prometheus Books.
- Tappenden, J. 1993a. 'The Liar and Sorites Paradoxes: Toward a Unified Treatment.' *Journal of Philosophy* **90**: 550-77.
- _____. 1993b. 'Analytic Truth — It's Worse (or Perhaps Better) than You Thought.' *Philosophical Topics* **21**: 233-261.
- Tappolet, C. 2004. 'Through Thick and Thin: Good and its Determinates.' *Dialectica* **58**: 207-20.
- Väyrynen, P. 2009. 'Objectionable Thick Concepts in Denials.' *Philosophical Perspectives* **23**: 439-69.
- Williams, B. 1985. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.
- Williamson, T. 2003. 'Understanding and Inference.' *Proceedings of the Aristotelian Society*, Suppl. vol. **77**: 249-93.
- _____. 2006. 'Conceptual Truth.' *Proceedings of the Aristotelian Society*, Suppl. vol. **80**: 1-41.
- _____. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.
- _____. 2009. 'Reference, Inference and the Semantics of Pejoratives.' In *Essays for David Kaplan*, J. Almog and P. Leonardi, eds. Oxford: Oxford University Press.