

OBJECTIONABLE THICK CONCEPTS IN DENIALS*

Pekka Väyrynen
University of Leeds

1. Introduction

It is common in moral philosophy, aesthetics, and increasingly in epistemology, to distinguish between “thick” and “thin” concepts. Concepts commonly regarded as thick in moral philosophy include *honest*, *courageous*, *tolerant*, *cruel*, *greedy*, and *deceitful*, those commonly regarded as thin include *good* and *bad*, *right* and *wrong*, and *ought*. We may similarly speak of thick and thin terms, that is, terms that are typically used to express thick and thin concepts.

Where thick and thin terms or concepts are commonly thought to differ — and in what the “thickness” of the former is thought to consist — is that thick terms or concepts have some substantive non-evaluative satisfaction conditions, whereas thin terms or concepts have little or no such content. For instance, even if things which involve causing someone pain just for fun count as both cruel and wrong, the meaning of ‘cruel’ seems to encode this kind of non-evaluative content in a way that the meaning of ‘wrong’ doesn’t. The distinction seems to mark a difference in degree along a spectrum of concepts, rather than a binary distinction in kind (Scheffler 1987: 417). For instance, the non-evaluative demands which terms such as ‘impartial’ and ‘just’ make on the world for their satisfaction seem to be weaker or less specific than those made by ‘cruel’ or ‘courageous’ but stronger or more specific than those made by ‘wrong’ or ‘good.’ Even some paradigmatically thin concepts appear to have some non-evaluative meaning. One example might be *ought*. The principle that ‘ought’ implies ‘can’ is widely endorsed, and it is common to think that ‘A ought to ϕ ’ entails rather than merely implicates ‘A can ϕ ’ and that the entailment is conceptual rather than merely metaphysically necessary.¹

This distinction regarding the non-evaluative contents of the relevant concepts is usually taken to mark a distinction within the class of *evaluative* and *normative* concepts. (The typical examples in moral philosophy are more clearly evaluative than they are normative, insofar as there is a clear distinction. I’ll speak of evaluative concepts, but mainly for the sake of brevity.) Many typical

characterizations of the distinction imply that thick and thin concepts are not merely both evaluative, but also evaluative in the same kind of way. Hence thick concepts are often taken to differ in kind from concepts which are used evaluatively but aren't in themselves evaluative, such as perhaps *pleasant*, *painful*, or *athletic*.²

It is an attractive thought that concepts such as *cruel*, *brutal*, and *generous* are evaluative concepts in some way in which thin concepts such as *good* and *wrong* are evaluative. But this shouldn't be built into the intuitive distinction between thick and thin concepts or otherwise assumed from the start. This isn't just because it isn't clear what makes a term or concept evaluative to begin with. For no doubt there is *some* important sense in which I don't understand what the critics of waterboarding try to communicate when they say "Waterboarding is brutal," if I fail to grasp that calling things 'brutal' is associated with a negative evaluative claim or attitude. But this doesn't show that *brutal* is an evaluative concept. Whether the kind of understanding one would lack in this sort case is conceptual, substantive, or something else, and thus whether thick terms are evaluative in meaning or in any other sense that distinguishes them from terms which are merely used evaluatively, depends on the way in which sentences containing thick terms convey evaluations.

These initial observations specify two tasks for any account of thick concepts. One is to identify the evaluative contents that go with thick concepts. The other is to theoretically locate this evaluative content. For instance, is that content attributable to thick terms or concepts themselves or to certain of their uses, and is it conveyed semantically or pragmatically? Further questions concerning the location of evaluative content will then be whether it makes thick concepts evaluative in some legitimate sense, and if so, whether it makes them evaluative in the same kind of way as thin concepts. When I need a neutral placeholder for the relationship between thick concepts and the evaluative contents they are used to convey, I'll speak of evaluative contents "associated with" thick concepts. (This isn't meant to imply association in a psychological sense, Freudian or otherwise.)

This paper approaches these tasks by considering "objectionable" thick concepts (OTCs, for short). Some concepts or words are regarded as objectionable in a distinctive way. Oscar Wilde was willing to call a certain short story 'disgusting' and 'horrible,' but declined the suggestion that it was blasphemous: "'Blasphemous' is not a word of mine."³ Those who don't accept Catholic sexual morality are unwilling to join the Pope in praising people or their sexual conduct as 'chaste' or condemning them as 'lustful' or 'lewd.' In fact pretty much any thick concept seems to be in principle open to being regarded as objectionable. For instance, certain fundamentalists may regard *forgiving* or *tolerant* as objectionable concepts in a usage which conveys approval. What these examples illustrate is a refusal to use concepts or words for reasons having to do with their evaluative dimension.⁴

Some care is due in getting the phenomenon in focus.⁵ A concept can be objectionable for reasons that don't have to do with its evaluative dimension. For instance, *sinful* might be an objectionable concept for reasons having to do with the theological implications of calling something 'sinful' and *sexual perversion* might be an objectionable concept insofar as calling something 'sexually perverted' involves an idea of unnaturalness that has no application (Slote 1975). A concept may also be objectionable only in some types of applications. For instance, many people who are perfectly happy to call babies and puppies 'cute' find it objectionable to call adult women 'cute' (Blackburn 1998: 103). But this doesn't make *cute* an objectionable concept for reasons to having to do with its evaluative dimension any more than the sense in which it is objectionable to call the painful deaths of babies 'funny' or my coffee grinder 'greedy' makes *funny* or *greedy* objectionable concepts.

The phenomenon of my interest is different. Calling something 'lewd,' for instance, typically conveys some such evaluative content as that it is at least to some extent bad or condemnable for having certain sorts of non-evaluative features having to do with sexual display. Regarding *lewd* as an objectionable concept requires regarding this way of thinking as failing to draw genuine evaluative distinctions — for instance, regarding things that have the relevant non-evaluative features as in no way bad so far the possession of those features goes. A concept can then said to be objectionable if the evaluative content with which it is associated in fact fails in this way to draw genuine evaluative distinctions.

Refusal to use words or concepts one regards as objectionable in this sense seems only to make sense if the evaluative contents associated with them enjoy a certain autonomy with respect to the attitudes and intentions of particular thinkers. For otherwise I should find it acceptable to use concepts I regard as objectionable so long as I don't endorse or intend to convey the evaluative contents I regard as objectionable. But that isn't what we typically find. The evaluative contents associated with thick concepts don't seem to be easily cancelable in this way. Racial slurs provide a robust example of expressions which carry objectionable content autonomously from speaker intentions.⁶ OTCs seem no different in kind (though perhaps they do in degree) in this respect. For instance, insofar as uttering 'Pissing on graves is blasphemous' conveys that pissing on graves is in some way condemnable for its irreverence towards something that is considered sacred, it does so irrespective of whether the speaker intends it to. I'll assume it as a constraint on accounts of thick concepts that they account for the autonomy of evaluative content.

The literature on thick concepts often mentions the phenomenon of objectionable concepts but, I believe, underestimates its significance in determining the location of the evaluative contents associated with thick concepts.⁷ To identify what is at issue when a thick concept is regarded as objectionable is to identify the evaluative content associated with it. To specify the way in which the concept

is associated with this evaluative content is to locate it. And since its theoretical location is a linguistic or conceptual matter, but whether a concept is in fact objectionable is a substantive matter, the location of the evaluative content won't be sensitive to whether the concept is in fact objectionable.⁸ Its location will be more sensitive to facts about regarding a concept as objectionable. Since that phenomenon is widespread, results regarding OTCs may hold with wide generality. There is no reason to expect any theoretically deep divide between the logic and meaning of OTCs and other thick concepts.

The aim of this paper is threefold. The first is to describe some data concerning how OTCs behave in various kinds of denials. The second is to argue that the data strongly suggest that thick terms don't in general have evaluative content as part of their conventional meaning and mention some consequences that would seem to follow if that is correct. The third is briefly to assess the strengths and weaknesses of three other sorts of analysis regarding the location of evaluative content in explaining the data. All this will take place under two caveats. First, there is no reason to assume in advance that thick concepts will be homogeneous with respect to the location of their evaluative content. It might turn out that one account works best for certain thick concepts, a different account for others, and so on. Second, the data regarding OTCs and the issues it raises are highly complex. The paper aims not to settle for good what the data regarding OTCs are or show, but to explore one potentially fruitful approach to evaluative concepts.

2. Objectionable Concepts, Disagreement, and Negation

What sort of disagreement obtains when one speaker regards some concepts of another as objectionable? This question bears on both tasks for accounts of thick concepts. For to identify the kind of evaluation with which those who regard a thick concept as objectionable disagree is to identify the evaluative content associated with the relevant uses of the word or concept. And to identify ways of expressing those disagreements is to identify data which hypotheses about the theoretical location of this content must explain.

If the Pope calls Brad and Janet, an unmarried couple, 'chaste,' we know what he is trying to get across. One thing we know is that if Brad and Janet's conduct is to count as satisfying 'chaste,' it must have some properties which signal some kind of dedication to not being sexually provocative. What count as such properties may be a substantive issue that remains unspecified in the analysis of the concept. One example of what many who go in for this kind of talk would regard as paradigmatically chaste might be abstinence from extramarital sex, but the matter can be controversial among them. Some might think that the relevant kind of dedication requires abstinence even in thought and desire. Even so, there might be non-evaluative constraints on *chaste* which will be common meaning between those who regard *chaste* as objectionable and those who don't.

To keep things simple, I'll assume that 'x signals dedication to not being sexually provocative' is in the right neighborhood to be the kind of non-evaluative claim that 'x is chaste' entails as a matter of meaning. But I won't assume that this predicate is anything like extensionally equivalent to 'chaste' or sufficient for it to apply, or even that it captures the full non-evaluative meaning of 'chaste,' whether combined with evaluation or not. There is no reason to suppose that every shade of meaning must be lexicalized in a way which would yield a term with such properties.⁹ We may even lack the resources for dividing between a thick term and some non-evaluative correlate at the level of semantics.

Another thing we know is that the utterance is in some way associated with thinking that Brad and Janet's conduct is good or praiseworthy in a particular way, for having some properties which signal an appropriate kind of dedication to not being sexually provocative. Some of us don't accept this way of thinking about sexual conduct. So it seems that if I regard *chaste* as objectionable but the Pope happily applies it, we have a disagreement that seems substantive. (Neither thinks we can both be right.) But what is it that the Pope thinks but I deny?

Disagreement can usually be expressed through negation. One might utter things like (1) or (2) to express disagreement with anyone who thinks otherwise:

- (1) Alex isn't good.
- (2) Alex isn't kind.

I grant (1) or (2) might sound odd if not uttered in discourse initially but in response to someone's claim that Alex is good, or that he is kind.¹⁰ But the explanation is pragmatic: typically we are expected to supply a reason when we contradict another speaker. The kind of reason one supplies to support a denial helps to clarify the focus of disagreement. One kind of ordinary disagreement can be expressed by amending (1) and (2) along these lines:

- (1') Alex isn't good. He often lies.
- (2') Alex isn't kind. He is often heartlessly hurtful of other people's feelings.

A parallel reply to the Pope's assertion 'Brad and Janet are chaste' would be (3):

- (3) Brad and Janet aren't chaste. They aren't married but often have sex.

A different kind of ordinary disagreement can be expressed by uttering (4):

- (4) Brad and Janet aren't chaste. They abstain from sex but really want it.

(3) expresses the kind of disagreement which someone who knows what Brad and Janet get up to might be having with someone who agrees what it takes to count as satisfying 'chaste' but has different beliefs about Brad and Janet. (4) expresses the kind of disagreement which speakers who have different beliefs about what

counts as satisfying ‘chaste,’ such as a more and a less conservative Catholic, might be having even if they shared the same non-evaluative beliefs about Brad and Janet.

But clearly neither kind of disagreement is one which those who regard *chaste* as an objectionable concept are having over ‘Brad and Janet are chaste.’ They have no stake in what kind of dedication to not being sexually provocative something must signal to count as satisfying ‘chaste.’ Their disagreement is over the general evaluative dimension associated with *chaste*. My hypothesis is that the differences between ways of expressing disagreement with claims made using concepts people regard as objectionable and responses to claims with which they disagree but which are made using one of “their concepts” are relevant to locating the evaluative contents associated with thick concepts.

Three points of methodology are in order. First, while I won’t directly discuss whether it is certain words or what they express that is objectionable, I suspect that what is objectionable with respect to OTCs in the first instance is a certain way of thinking rather than speaking.¹¹ Thus my procedure requires me to assume that words that express such concepts or representations are related to them in some way such that linguistic evidence constrains or gives clues to the theoretical location of evaluative content.¹²

Second, I’ll assume that the linguistic conventions which determine the meanings of words that express OTCs, and the mechanisms which determine what is conveyed by their uses, are not exceptional but probably are the same sorts of conventions and mechanisms as determine these things for the rest of the language.

Third, when I suggest that some sentence is fine, odd, or one that certain speakers would hesitate to use, or would display semantic incompetence, or the like, I only claim that statistical patterns in different speakers’ linguistic judgments tend in the way of these judgments. Since there may be wide variation in these judgments, if only because terms may be vague or unclear, sometimes there may be no linguistic fact of the matter whether, for instance, a particular proposition follows analytically from a sentence that uses a thick term, or is inconsistent with it, and so on.

Let’s now consider how those who regard *chaste* as an objectionable concept could express the kind of disagreement they have with (5).

(5) Abstinence from extramarital sex is chaste.

The lesson of (3) and (4) is that we wouldn’t expect those who regard *chaste* as an objectionable concept typically to be willing to express their disagreement with (5) by uttering either (6) or (7).

(6) Abstinence from extramarital sex isn’t chaste.

(7) Abstinence from extramarital sex isn’t chaste. Chastity requires abstaining even from thoughts and desires concerning extramarital sex.

Uttering (6) alone may sound odd if uttered in response to (5) rather than initially to express disagreement with anyone who says otherwise. But uttering some such expansion as (7) wouldn't express the right sort of disagreement. What is at stake between (5) and (7) is the extent of dedication to not being sexually provocative that something must display to be good or praiseworthy in the relevant particular sort of way or respect. Those who regard *chaste* as objectionable have no stake in this. They think that whether something is chaste or not, it wouldn't be good in any way for being so.¹³ We would typically expect them not to be willing to express their disagreement with (5) by uttering (6) plus a reason for disagreement which patterns with (3) or with (4) and (7).

Those who remain unsure that this would be a bad way for those who regard *chaste* as objectionable to express disagreement with (5) are invited to consider whether they would be willing to express disagreement with (8) by uttering (9):

- (8) The Italian Euro 2008 football team are greasy wops.¹⁴
- (9) The Italian Euro 2008 football team aren't greasy wops.

I would expect that the answer is No (unless one uses a certain special tone of voice).¹⁵ OTCs and racial slurs seem parallel in this regard.

Next consider (10) and (11):

- (10) #Abstinence from extramarital sex isn't chaste. Chastity doesn't reside in having properties that signal dedication to not being sexually provocative.
- (11) Yes, abstinence from extramarital sex is chaste, because it signals dedication to not being sexually provocative.

(10) seems defective. Whatever else 'x is chaste' may conceptually entail, one thing it does entail is something along the lines of 'x has some properties that signal dedication to not being sexually provocative.' Those who regard *chaste* as an objectionable concept would, however, typically not be willing to express their agreement that (5) conceptually entails this condition by uttering (11).

My next group of examples consists of (12)–(14):

- (12) Abstinence from extramarital sex isn't chaste, [but/although] it does signal dedication to not being sexually provocative.
- (13) Abstinence from extramarital sex isn't chaste; that it signals dedication to not being sexually provocative doesn't mean it is good in any way.
- (14) Abstinence from extramarital sex isn't chaste, because it is in no way good for signaling dedication to not being sexually provocative.

(12)–(14) negate (5). But, in contrast to such expansions of (6) as (7), those who regard *chaste* as an objectionable concept should often find it acceptable to express disagreement with (5) by uttering them. Follow-up clauses which merely

affirm a non-evaluative entailment of the assertion need a concessive particle, as in (12). Follow-up clauses which deny the sort of evaluative content which those who regard *chaste* as objectionable take to be associated with it need either a pause, as in (13), or an explanatory connective like ‘because/since/for,’ as in (14). The same goes for the variants of (13) and (14) which specify that abstinence from extramarital sex isn’t praiseworthy in the respect of signaling dedication to not being sexually provocative. The acceptability of each to those who regard *chaste* as objectionable is clearer with pitch-marking where ‘not’ and ‘chaste’ are stressed with a rising intonation.

The follow-up clauses in (13) and (14) seem to need the complexity they display. Abstinence from extramarital sex might be good in all sorts of ways, such as a way of pleasing the priest, avoiding venereal disease, or winning *Burning For It* (imagine a reality T. V. show involving unmarried Catholic couples). But these don’t constitute the evaluative dimension associated with *chaste*. Hence it is irrelevant that a simpler follow-up, as in ‘A isn’t chaste, because it isn’t good in any way,’ might be false. What is being denied by those who regard *chaste* as objectionable is that things are good insofar as they have some properties that satisfy the kind of generic non-evaluative condition that goes with calling things ‘chaste.’ (What specifically these properties are may, again, be left unspecified in the analysis of the concept.)

Finally consider two examples of what I’ll call “concessive denials”:

- (15) Yes, but there is no reason to praise abstinence from extramarital sex just because it signals dedication to not being sexually provocative.
- (16) Abstinence from extramarital sex may well be chaste, but I want to emphasize that it is in no way good for signaling dedication to not being sexually provocative.

The interpretation and acceptability of these examples is the least clear of the bunch.¹⁶ My own inclination is to judge that (15) and (16) can be used by those who find *chaste* objectionable to express their disagreement with (5) when they are read as denying the evaluative content the speaker regards as objectionable and conceding to (5) only some non-evaluative aspect in which the speaker has no stake. An example would be a reading on which the first half of (16), and the ‘yes’ in (15), echo only the claim that abstinence from extramarital sex signals dedication to not being sexually provocative but not any claim which (5) may be used to make about what exactly that takes. Such a reading isn’t *ad hoc*. (7) suggests that disputes among those who don’t regard *chaste* as objectionable about whether chastity requires abstinence merely in act, or also in thought and desire, are substantive rather than semantic. And (12)–(14) suggest that one can direct denial specifically at the associated evaluative content. No doubt this reading isn’t always appropriate. Sometimes it surely is alright to speak concessively for merely pragmatic reasons. Determining whether some instances of (15) and (16) can be read in this way thus requires paying attention, among other things, to the interaction between ways to co-ordinate conversation

(by means of such discourse markers as ‘yes, but...’ and other devices of politeness) and what is said in conversation.

Replies to (5) which pattern after (13)–(16) can be used to describe the relations of agreement and disagreement which are possible between those who regard *chaste* as objectionable and those who don’t. For they can be used to identify just what it is that the former don’t accept but take to be associated with typical utterances of (5). In each case the follow-up clause specifies which aspect of an utterance of (5) is being denied. For instance, uttering variants of (13) and (14) expresses disagreement with the positive assertions corresponding to their second halves. Such assertions express evaluative contents in the neighborhood of (17) and (18):

- (17) Things are good in a way for having properties that signal dedication to not being sexually provocative.
- (18) Things are praiseworthy to some degree for having properties that signal dedication to not being sexually provocative.

It seems correct that the focus of the disagreement between someone who utters (5) and those who regard *chaste* as objectionable lies in the neighborhood of (17) and (18). I explained above, in connection with (13) and (14), why they seem to need the complexity they display. If the focus of disagreement is something along their lines, then the sorts of disagreements which concern objectionable concepts come out as involving substantive evaluative disagreement and not as (merely) verbal disputes. The kinds of denials which those who regard a concept as objectionable would typically be willing to issue against the assertions made by those who don’t can thus be treated as the sort of genuine engagement with the contents of others’ claims which one would expect to find in genuine disagreement.

A final observation about these examples before moving on to discuss what they show is that the non-evaluative implications of OTCs and the evaluative contents associated with them behave asymmetrically in certain respects. For one asymmetry, compare (19), a variant of (16), and (20):

- (19) Abstinence from extramarital sex may well be chaste, although it is in no way good for that.
- (20) #Abstinence from extramarital sex may well be chaste, although it doesn’t signal dedication to not being sexually provocative.

As noted in connection with (16), there may be disagreement as to whether uttering (19) would be an acceptable way for those who regard *chaste* as objectionable to express their disagreement with (5). But (20) is clearly bad: uttering it would betray a failure to grasp that satisfying ‘chaste’ requires having some properties that signal dedication to not being sexually provocative; here compare (10).

For another asymmetry, consider (12) and (21):

- (12) Abstinence from extramarital sex isn't chaste, although it does signal dedication to not being sexually provocative.
- (21) ?Abstinence from extramarital sex isn't chaste, although things are good in a way for signaling dedication to not being sexually provocative.

The second half of (12) concedes to (5) a non-evaluative implication of *chaste* which uttering (20) would show a failure to grasp. The second half of (21) concedes to (5) the evaluative content which those who regard *chaste* as objectionable don't accept. What aspect of (5) could one then be denying on this concessive reading of (21)? Surely not that which (20) denies. But what then?

3. Evaluative Content and Conventional Meaning

The behavior of OTCs in denials bears on the meaning of thick concepts. The data make it plausible that if we regard *chaste* as an objectionable concept, our objection to (5) isn't (merely) that it is false. If (5) were false, then (6) should be non-problematically true, even if uttering it alone would be pragmatically odd. For the claim that (5) is false can be expressed by applying ordinary truth-conditional negation: taking *p* to *not-p* in this case generates (6). Then there is no reason why we should be unwilling to express our disagreement with (5) by uttering (6) plus some follow-up on the model of (3), (4), and (7) which indicates that our use of negation is truth-conditional. Yet typically we are.

Conditionals exhibit the same phenomenon. If our objection to (5) were that it is false, then conditionals such as 'If abstinence from extramarital sex is chaste, then so is refraining from desiring extramarital sex' should be non-problematically true, due to a false antecedent. But the truth of such conditionals isn't something that those who regard *chaste* as objectionable are typically willing to grant.

These points can be explained as instances of a more general pattern. Truth-conditional negation and conditionalization are presupposition holes: they take scope over the assertion that sentence *S* would have made if uttered, but not other types of information its utterance would have conveyed, such as presuppositions or implicatures. Hesitation to express disagreement with (5) by uttering an expansion of (6) that uses negation truth-conditionally indicates that the evaluative content associated with *chaste* projects past presupposition holes. In other words, assertions of 'A is chaste' and 'A isn't chaste' both convey the evaluative content which is associated with *chaste* but not accepted by those who regard it as objectionable. (Such utterances may also convey that the speaker endorses such an evaluative content or express a corresponding attitude.) What makes them unwilling to express disagreement with (5) by using negation truth-conditionally is that this would typically misrepresent their evaluative outlook.¹⁷

If those who don't accept the evaluative content associated with *chaste* cannot express their disagreement by applying truth-conditional negation to 'A is chaste,' then it doesn't seem to be part of the conventional meaning of (5) that it has some such evaluative truth-condition as (17) or (18). For brevity, I'll put this

as the claim that the evaluative contents associated with thick concepts aren't truth-conditional.¹⁸ Since my argument only assumes that 'chaste' is regarded as objectionable, not that it in fact is objectionable, the argument will apply to any thick terms which pattern like 'chaste,' whether in fact objectionable or not.

The claim that the evaluative contents associated with thick concepts aren't truth-conditional aspects of thick claims is supported by the behavior of OTCs in denials. It fits well with the reading of concessive denials described in section 2. If something like (17) of (18) were part of the truth-conditions of (5) in virtue of its conventional meaning, (15) and (16) should sound contradictory. But they don't, at least not on the reading in question, and many of those who find them somewhat bad won't find them contradictory. So those who accept this truth-conditional analysis must either explain (15) and (16) in some other way or explain them away.¹⁹ Whether such explanations are always the best merits further scrutiny.

If the evaluative contents associated with thick concepts aren't truth-conditional aspects of thick claims, that would also explain why it is hard to read the negation in (13) and (14) as truth-conditional. Those who regard *chaste* as objectionable are typically willing to utter these sentences, but the evaluative content to which they object would escape from the scope of truth-conditional negation, as illustrated by examples like (7). So some instances of (5) seem to be such that when speakers express disagreement with them through some such negation as (13) or (14), the negation focuses on the evaluative content associated with the relevant utterances (5) but this evaluative content isn't part of the truth-conditions of (5).

A better explanation of the data supports the claim that the evaluative contents associated with thick concepts aren't truth-conditional aspects of theirs. The explanation is that (13) and (14), as well as (12), involve a *metalinguistic* use of negation. While the truth-conditional use of negation is a device for objecting to the *proposition* expressed by a sentence S, a metalinguistic use is a device for signaling that there is something wrong with the *utterance* of S. Metalinguistic negation may be directed at any number of features of an utterance, including presuppositions, implicatures, and such formal features as intonation, pronunciation, word-formation, style, and register.²⁰ The range of the phenomenon is evident from (22)–(25):

- (22) The king of France isn't bald, because there is no king of France.
[Existential presupposition]
- (23) He's not rich, he's filthy rich. [Scalar implicature]
- (24) He didn't call the POlice, he called the poLICE. [Pronunciation, register]
- (25) We didn't [have intercourse/make love] — we fucked. [Word choice, style]

The material in square brackets specifies which feature of a positive assertion corresponding to the first half of each example is being commented on by the sentence.

Why think that (13) and (14) are instances of metalinguistic negation? It is very plausible that the second half of each focuses the negation in the first half of each specifically on the evaluative content associated with *chaste*. We saw that it is hard to see how that negation could be a truth-conditional negation of (5). Further support for reading it as metalinguistic comes from cases where negation is indisputably truth-conditional but which aren't relevantly analogous:

- (26) Taylor isn't a vixen, because although Taylor is a fox, he isn't female.
- (27) Abstinence from extramarital sex isn't chaste, because although it does signal dedication to not being sexually provocative, it is in no way good for that.

(26) contradicts 'Taylor is a vixen,' and so the negation is truth-conditional, but (27) needn't be read as contradicting (5). (27) effectively conjoins (12) and (13). If the negation isn't truth-conditional in (12) and (13), then it isn't truth-conditional in (27) either.

If the negation in (13) and (14) is metalinguistic, then there is no reason to think that something like (17) or (18) is part of the truth-conditions of those utterances of (5) to which one could object by uttering (13) or (14). The metalinguistic use of negation is a device for objecting precisely to features of utterances other than their truth-conditional content. Moreover, the patterns in the sentences that are acceptable to speakers who regard *chaste* as objectionable suggest that this evaluative content stays inside the negation in (13) and (14).

The hypothesis that these examples are instances of a metalinguistic use of negation predicts all this, since negation doesn't function in such uses as a presupposition hole. Rather, it is directed precisely at such features of utterances as presuppositions, implicatures, and various formal features. Classifying (13) and (14) as metalinguistic uses of negation is also independently plausible.²¹ They are naturally glossed as comments on an utterance of (5), to the effect that to represent something in terms of what is within the scope of negation is to misevaluate in a systematic way. Utterance commentary is the sort of function that is served by metalinguistic negation.

4. Some Important Consequences

It matters a lot for theorizing about thick concepts in particular and evaluative concepts in general whether the behavior of OTCs in denials shows that the evaluative contents associated with thick concepts aren't truth-conditional in the way just discussed. Writings on the distinction between thick and thin concepts commonly just assume that evaluative content is part of thick concepts.²² Even if this assumption isn't false, it requires defense. If it is false, the consequences cut deep and wide.

The data I have been discussing seem to suggest that the evaluative contents associated with thick concepts are related to thick concepts differently from how the evaluative contents associated with thin concepts are related to thin concepts. In the case of thin concepts it is trivial that if concept *C* is evaluative or normative, then sentences of the form ‘*x* is *C*’ conceptually entail that *x* is good in some way, or that there are reasons to respond to *x* in a certain way, or the like. The behavior of OTCs in denials suggests that many thick terms or concepts aren’t evaluative in this sense as a matter of their conventional meaning.²³ The possibility even remains that thick concepts aren’t really evaluative at all.²⁴ So a good account of thick concepts shouldn’t merely explain the behavior of OTCs in denials in way that accounts for the autonomy that the evaluative contents associated with thick concepts enjoy with respect to the intentions and attitudes of particular speakers. It should also determine, in some principled way, whether thick concepts are evaluative in any sense which distinguishes them from concepts which are merely used evaluatively and whether thin concepts are evaluative in the same kind of way.

The behavior of OTCs in denials also bears significantly on at least two issues that have been much discussed in the literature on thick concepts. The first is that their behavior suggests that we should be able to imagine someone acquiring and using concepts like *chaste*, *lewd*, or *lustful*, but later coming to think that in applying those concepts they were systematically miscalculating. Thus we should be able to imagine such a person adopting a way of thinking on which the fact that an act manifests lust isn’t even normally a reason to condemn it, the fact that it involves passing certain limits on sexual display doesn’t even normally make it in any way bad, and so on. It doesn’t seem that changing one’s evaluative dispositions in this way must involve losing the concept, as opposed to continuing to possess it but not using it. But if so, then the evaluative contents associated with thick concepts wouldn’t seem to be essential to them. And if so, there is no obstacle in principle for someone who regards some thick concept as objectionable nonetheless to become more or less competent in predicting how those who don’t regard the concept as objectionable would apply it. These claims are controversial in the literature.²⁵ But the behavior of OTCs in denials seems to weigh in their favor.

The second issue is whether the evaluative and non-evaluative features of thick concepts form some kind of amalgamated whole or can be divided or “disentangled” into distinct components.²⁶ This is a murky issue, in part because it is unclear just what relation this disentanglement is supposed to be. (For instance, is it supposed to hold analytically or otherwise *a priori*?) But if the evaluative contents associated with thick concepts aren’t part of the truth-conditions of claims using such concepts in virtue of their conventional meaning, then there would seem to be a sense in which disentanglement is in principle possible. For sentences of the form ‘*x* is *T*’ (where ‘*T*’ is a thick term) would have some truth-conditional content and some non-truth-conditional content, and their evaluative content would be confined to the latter sort.²⁷ Thus it seems

that those who deny that thick concepts can be disentangled, at least in any sense to which linguistic evidence is relevant, should resist the kinds of claims that the behavior of OTCs in denials seems to support. For the data concerning that behavior would seem to provide an argument against their position unless they can explain the data in some other way or else explain it away.

The behavior of OTCs in denials also bears on whether objectionable concepts are empty. (This claim often comes up in conversations.) If nothing like (17) or (18) is entailed by the conventional meaning of ‘A is chaste,’ then ‘A is chaste’ doesn’t mean something like *A has properties which signal dedication to not being sexually provocative, and is good in some way for having them* or *A has properties which signal dedication to not being sexually provocative, and is to some degree praiseworthy insofar as those properties go*. The hypothesis that the conventional meaning of thick concepts encodes evaluative content in this kind of way predicts precisely that OTCs are empty. But regarding ‘chaste’ as objectionable only requires thinking that nothing satisfies this kind of condition, not that the condition is part of what ‘chaste’ means. Nor is such an analysis required for saying that the mistake in thinking of something as chaste involves thinking that things are good or praiseworthy in some way for signaling dedication to not being sexually provocative.

So how plausible is the claim that OTCs are empty? Concessive denials suggest, on the reading described in section 2, that those who regard *chaste* as objectionable can attribute something true to some utterances of ‘*x* is chaste,’ namely whatever it is that uttering (15) and (16) would concede to (5), provided that these speakers clarify that they aren’t endorsing the evaluative content associated with *chaste*. This isn’t what we would expect if OTCs were empty. Nor would we expect to find such speakers unwilling to say things like ‘No one is chaste.’ It also isn’t clear why, on this kind of analysis, the non-evaluative and evaluative contents associated with OTCs should exhibit the kinds of asymmetry as we find in (19)/(20) and (12)/(21).

A more straightforward problem is that the hypothesis that OTCs are empty doesn’t by itself provide a satisfactory account of what the problem is with using objectionable concepts. It seems plain that what is wrong with wielding such concepts as *chaste* or *lustful* isn’t merely the sort of fault that is involved in wielding such empty concepts as *phlogiston* or *Bigfoot*. So the former sort of fault isn’t fully explained by the hypothesis that OTCs are empty.

The most plausible way of developing the idea that OTCs are empty points away from locating their evaluative contents in the truth-conditions of assertions that use them. If ‘chaste’ has the above kind of evaluative meaning, then either (5) is false or it lacks truth-value. If (5) were false, then instances of (6) which use negation truth-conditionally should be non-problematically true, contrary to what we have seen. So the better option would be to say that (5) is neither true nor false, and similarly for the relevant instances of (6). However, lack of truth-value is most typically attributed to expressions with false presuppositions. So the claim that (5) is neither true nor false seems to tend away from the idea that

the evaluative content associated with a thick concept T is a truth-conditional aspect of sentences of the form 'x is T' towards the idea that they presuppose it (see section 5).

It is worth stressing how difficult locating the evaluative contents associated with thick concepts has become. For what we have seen is that natural assumptions concerning evaluative concepts in general and OTCs in particular seem jointly inconsistent. Suppose *lewd* is an objectionable thick concept. If *lewd* isn't empty, then there are some truths of the form 'x is lewd.' But if *lewd* is an objectionable concept, then not everything that is lewd is thereby bad in any way. Thus it seems that *lewd* cannot be an evaluative concept in the sense that 'x is lewd' conceptually entails negative evaluative content. In what sense are thick concepts then supposed to be evaluative, and evaluative in the same way as thin concepts?

At least three types of analysis of thick concepts remain on the table if the evaluative contents associated with them aren't truth-conditional in virtue of their conventional meaning. I only have space for a brief assessment of the respective strengths and weaknesses of each view against the above constraints. This will select none as clearly better than others. Hence the discussion to follow is far from a last word on the theoretical location of the evaluative contents associated with thick concepts. Much more will remain to be said.²⁸

5. Thick Concepts and Presupposition

There are two obvious proposals to consider concerning the theoretical location of the evaluative contents associated with thick concepts if they really are related to thick concepts differently from how the evaluative contents of thin concepts are related to thin concepts and if they can be denied by a metalinguistic use of negation. The proposals are that the evaluative contents of thick concepts are presupposed and that they are implicated.²⁹ In this section I'll discuss the presupposition analysis.³⁰

The idea that the evaluative content associated with a thick concept T is presupposed by sentences or utterances of the form 'x is T' is well-suited to capture various behaviors of OTCs. The reactions of those who find *chaste* objectionable allow that if (5) presupposes some such objectionable content as (17) or (18), then, under normal conditions, one can reasonably infer it from either an assertion of 'A is chaste' or an assertion of 'A isn't chaste' (read as truth-conditional negation). Similarly, a natural way for those who find *chaste* objectionable to indicate that they have no stake in disputes about what sort of sexual conduct counts as satisfying 'chaste' is to say that whether a particular bit of conduct is chaste or not, it wouldn't be in any way good for that. Instead of accepting or rejecting any particular claim of the form 'A is chaste,' they demur. Such reactions are to be expected if a claim carries a presupposition one rejects. Witness (28):

- (28) Abstinence from extramarital sex is neither chaste nor unchaste, because things are in no way good for signaling dedication to not being sexually provocative.

(28) sounds fine to many ears. A presuppositional analysis predicts that.

A presuppositional analysis can also explain why (12) should sound consistent:

- (12) Abstinence from extramarital sex isn't chaste, although it does signal dedication to not being sexually provocative.

If negation took broad scope over both the evaluative content associated with (5) and whatever non-evaluative content it has, (12) should sound inconsistent. One constraint in deciding what presuppositions to assign to utterances is to avoid inconsistency if possible. Hence a presuppositional analysis predicts a reading on which negation takes scope only over the associated evaluative content. This fits with a metalinguistic use of negation to plug presuppositions.

Finally, a presuppositional analysis implies an asymmetry between the evaluative content associated with 'chaste' and whatever non-evaluative content is at issue in utterances like (5), since presumably the latter is non-presuppositional. This contrast could perhaps be used to explain the asymmetries illustrated by (19)/(20) and (12)/(21).

A presuppositional analysis has two possible forms, depending on whether evaluative presuppositions are regarded as pragmatic or semantic. The difference is between anything that is intuitively felt to be a requirement on context (or, "taken for granted" by a sentence or utterance), whether triggered conventionally or conversationally, and a requirement on context which is conventional.³¹ A standard way to characterize this difference is that a proposition *p* is a semantic presupposition of sentence *S* if *p* must be true in order for *S* to be true or false, and that *p* is a pragmatic presupposition of *S* in context *c* if utterance of *S* is appropriate in *c* only if *p* is mutually assumed by the speaker and hearers. So pragmatic presuppositions are restrictions on the common ground: the assumed truth of *p* is a precondition for felicitous utterance of *S* in *c*. If *p* is a semantic presupposition of *S*, then *S* has no truth value if *p* is false, whereas if *p* is a pragmatic presupposition of an utterance of *S*, then the utterance can say something true even if *p* is false. Desire to allow that presupposition failure needn't imply a lack of truth-value has made a pragmatic approach prevalent.

The thought that thick concepts semantically presuppose the evaluative contents associated with them has difficulty with concessive denials. For if (5) had no truth-value, as it wouldn't if it had a false semantic presupposition, then one shouldn't be conceding to (5) what (15) and (16) seem to concede. So a semantic presupposition view must explain why the reading of concessive denials described in section 2 is unavailable. If the relevant presuppositions are conversationally triggered, however, then an utterance of (5) can be conceded to have said something true even if it has a false presupposition.³²

Whether the evaluative contents associated with thick concepts should be regarded as semantically or pragmatically presupposed depends on how such presuppositions are triggered and, hence, on what explains their presence. The systematic patterns we have found in the behavior of OTCs in denials don't show that their evaluative presuppositions would be semantic. Many linguists hold that some presuppositions which work systematically (for instance, the factive presuppositions of verbs like 'know' and 'regret') are conversationally rather than lexically triggered, or even that all are.³³ Hence even if thick terms formed a distinct category of presupposition triggers, on par with definite descriptions, factive verbs, implicative verbs, change of state verbs, expressions of repetition and temporal relations, clefts, stressed constituents, and questions, this alone wouldn't show that the presuppositions they trigger are part of the conventional meaning of words or grammatical constructions.³⁴

The claim that evaluative presuppositions are semantic thus implies that they can only be triggered lexically by words used to express thick concepts. But it is unclear what grammatical property of these words would directly create the relevant presuppositions. Ascribing them some additional property just to explain the presuppositions would seem to be *ad hoc*. Such theoretical commitments might nonetheless be tolerable if there were strong independent reasons to think that the evaluative contents associated with thick terms are presuppositional but not explained by conversational considerations. But if evaluative presuppositions were to arise out of general conversational considerations rather than features of particular contexts, the pragmatic hypothesis might have resources to explain why these evaluative presuppositions work systematically the way they do and why they enjoy a certain autonomy with respect to the intentions and attitudes of particular speakers. So in this respect it might be no worse off than its semantic cousin. But explanations of these putative evaluative presuppositions which appeal only to general conversational principles have yet to be given.

The force of these considerations is twofold. First, if the evaluative contents associated with thick concepts are presuppositional, then an appropriate explanation of how these putative evaluative presuppositions are triggered will look different depending on whether the presuppositions are supposed to be semantic or pragmatic. Second, a satisfactory explanation of neither sort has as yet been given. Whether such explanations are available remains to be determined. But even if they are available, any presuppositional analysis will have other outstanding issues.

One general issue concerns the backgrounded character of presuppositions. For an utterance to be appropriate in a given context, its presuppositions (whether semantic or pragmatic) should already be part of the common ground. This constraint on context manifestly fails in the case of disagreements between those who regard a given thick concept as objectionable and those who don't. Often when this appropriateness constraint fails, it may be easy for the hearers to accommodate the presupposition.³⁵ But it is equally manifestly not the case

that the hearers are typically willing to accommodate evaluative presuppositions they regard as objectionable, nor can the speaker reasonably expect them to.

This raises two problems. First, a presuppositional analysis predicts that users of OTCs will knowingly and systematically make conversationally inappropriate utterances in such contexts. Sometimes this needn't prevent them from achieving their conversational purpose, such as when their purpose is rhetoric (for instance, to exhort or persuade others whom they know to disagree to adopt their view) or symbolic (for instance, simply to put their opinion on public record, without any expectation of accommodation). But supposing that this is what is always going on when one is addressing those who find one's concepts objectionable implies that speakers' intentions will vary solely depending on whether they believe that their hearers regard their concepts as objectionable: genuine engagement only with the orthodox, only rhetorical or symbolic expression with the unorthodox. What evidence is there for that? Second, given a presuppositional analysis, a competent speaker would know that if no one objects to the use of an objectionable concept, the speaker and hearer are entitled to assume that all presuppose a certain sort of evaluation regarding its targets. For the analysis would mark these concepts as having the purpose of slipping a certain sort of evaluative presupposition into the common ground. But insofar as most contemporary speakers don't want certain words to have such properties, one wonders how those words could have them.³⁶ So it is precisely OTCs that will be hard cases for a presuppositional analysis of thick concepts.

Another general issue concerns the relationship between thick and thin concepts. A presuppositional analysis portrays thick concepts as evaluative in the sense that they trigger evaluative presuppositions. If thin concepts, too, were evaluative only in this sense, what would their non-presuppositional content be? A better option for a presupposition theorist might be to question the idea that thick and thin concepts are evaluative in the same kind of way. One way to do this would be to point out that it wouldn't be a compelling objection to a theory of pejorative expressions if it implied that pejoratives and thin concepts are evaluative in different ways. Why couldn't one then simply regard it as a substantive result that thick concepts turn out to be more like pejoratives than like thin concepts, or that they turn out to be different from both, as the case may be?³⁷ This should make us wonder why it seems attractive that thick and thin concepts are evaluative in the same kind of way. Such a similarity isn't required, for instance, in order for thick concepts to be able to play many, if not all, of the important roles they have variously been thought to play in normative ethics, since various non-evaluative concepts, such as hedonic ones, are sometimes recruited to play those roles as well. This an important issue that deserves a much fuller discussion than is possible here, so I'll leave it as a loose end. Essentially this reply to the parallel problem is also available for an implicature analysis of the evaluative contents associated with thick concepts. Hence I'll now turn to that suggestion.

6. Thick Concepts and Implicature

The hypothesis that the evaluative contents associated with thick concepts are carried by implicatures comes in two different forms, depending on whether the implicatures would be conventional or conversational.³⁸ The conventional implicatures of an expression are part of its conventional meaning, although not its truth-conditions, whereas conversational implicatures arise from the making of utterances and are supposed to be calculable from what was said plus a Gricean co-operative principle and maxims of conversation (Grice 1975: 26–30).

Neither conventional nor conversational implicatures are backgrounded in the way presuppositions are supposed to be, so an implicature analysis avoids that problem with a presuppositional analysis. The hypothesis that thick concepts conventionally implicate evaluative contents explains their autonomy with respect to the intentions and attitudes of particular speakers by treating it as conventionally encoded. But conversational implicatures may enjoy such autonomy as well if they are appropriately generalized.³⁹ Securing such a status for the evaluative contents associated with thick concepts would require giving Gricean derivations of propositions like (17) or (18) from the asserted content of (5) without appeal to special context.

One problem for the conversational implicature analysis is that conversational implicatures are supposed to be reinforceable: they allow conjoining a statement that generates an implicature with an overt statement that says the same as the implicature without getting any feeling of redundancy.⁴⁰ But consider:

- (29) Abstinence from extramarital sex is chaste, and it is good in a particular way for signaling dedication to not being sexually provocative.
- (30) Some athletes smoke, but not all do.

The second half of (29) sounds redundant, but reinforcing a generalized conversational implicature in (30) doesn't.

Things become less clear when we turn to the diagnostics that conventional implicatures are supposed to be detachable but non-cancelable whereas conversational implicatures are supposed to be cancelable but non-detachable (Grice 1975: 25–26).

Implicature content is cancelable (in Grice's sense) just in case the same sentence can be uttered, without linguistic impropriety, without expressing that content. Insofar as the evaluative contents associated with thick concepts enjoy a certain autonomy with respect to the attitudes and intentions of particular speakers, many typical devices for canceling implicatures will often fail to cancel that content. It is unclear that the first half of (16) has the same asserted content as (5), even if it concedes something to (5), and in many contexts (31) sounds questionable:

- (31) Abstinence from extramarital sex is chaste, but I don't mean to imply that it is in any way good for signaling dedication to not being sexually provocative.

The claim that the evaluative contents associated with thick concepts are cancelable in the way conversational implicatures are supposed to be at least requires further defense.⁴¹

Implicature content is detachable just in case the same asserted content can be expressed in the same context without the implicature, typically by substituting a different but coextensive term for the word carrying the implicature. (Conversational implicatures are supposed to be non-detachable, since they are calculated on the basis of the whole content of what was said in the given context, not its linguistic form.) There is no reason to expect that different co-extensive terms will in general be available to substitute for thick terms. But one might still think, counterfactually, that if (5) had a truth-conditional equivalent that didn't use the word 'chaste,' that sentence could avoid being associated with (17) or (18).

One way to test the detachability of such evaluative content as (17) and (18) is offered by the idea that (15) and (16) avoid such association because they use 'chaste' only in inverted commas. One function of inverted commas is to deny asserted content. The new "friends" of a lottery millionaire, although he might call them such, aren't his friends.⁴² But they might also be used to overcome expressive limitations of language. So perhaps inverted commas could be used in this way to detach the evaluative contents associated with thick concepts.

Imagine a people called the Bulli, who use the word 'wumpua' to praise and rejoice in acts which we call 'brutal' on pretty much those grounds on which we call things 'brutal' (Gibbard 2003: 165).⁴³ Suppose the Bulli and we witness a police officer subdue someone about to commit a violent crime and kick him hard when he is down. It seems we can understand each other without having any neutral word for the extension of 'brutal' and 'wumpua' and that we can express our disagreement with the Bullish claim 'That arrest was wumpua!' by saying things like these:

- (32) The arrest wasn't wumpua; it was brutal.
 (33) The arrest wasn't wumpua; acts aren't praiseworthy for being unnecessarily violent.
 (34) Yes, the arrest was "wumpua," but it isn't praiseworthy for having been unnecessarily violent.

(32) and (33) seem to involve a metalinguistic use of negation.⁴⁴ But then it should be plausible to read also (34) as denying something other than the truth-conditional content of 'The arrest was wumpua!' Yet (34) seems to feature a perfectly acceptable use of inverted commas.⁴⁵ So it isn't clear that inverted commas cannot be used to detach the evaluative contents associated with thick

concepts. What exactly this would show in the bigger scheme of things is less clear: at least any lexically triggered presuppositions are also usually regarded as detachable.

A conventional implicature analysis has trouble with concessive denials. Even if (15) and (16) don't sound as natural as the denial of a generalized scalar implicature in (35), they sound better than the denial of conventional implicature in (36):

(35) A: John had six pints. B: Yes, but in fact he had eight.⁴⁶

(36) #A: Basil Fawlty is English, therefore he is brave. B: Yes, but his being English doesn't imply his being brave.

The only way I can rescue (36) is by extensive pragmatic reanalysis that takes B to concede the truth-conditional content worded in a way that detaches the conventional implicature, such as 'Basil Fawlty is an Englishman and he is brave.' But hearing (15) and (16) as acceptable seems to require much less radical repair.

The force of these considerations is that the evaluative contents associated with OTCs display some features that are traditionally attributed to conversational implicatures, but also some that are traditionally attributed to conventional implicatures, and that they aren't clearly either with respect to yet other features. Hence neither sort of implicature analysis is a good fit with the behavior of OTCs in denials.⁴⁷ But even if this problem can be solved by taking issue with the data, any implicature analysis will have other outstanding issues.

One issue is that many denials of the evaluative contents associated with thick concepts are non-contrastive contexts of negation, whereas implicature denials are typically contrastive, as illustrated by (37) and (38).⁴⁸

(37) John didn't have six pints but seven.

(38) I'm not happy, I'm ecstatic.

But consider a contrastive denial like (39):

(39) #Abstinence from extramarital sex isn't chaste, but rather things are in no way good for signalling dedication to not being sexually provocative.

(39) sounds markedly worse than such non-contrastive denials as (13) and (14).

Another issue concerns the relationship between thick and thin concepts. As noted at the end of section 5, here implicature theorists have much the same resources available to them as presupposition theorists.

7. Thick Concepts and Contextualism

Insofar as presupposition and implicature suggestions don't satisfy, one might reconsider the idea that evaluative content is a truth-conditional aspect of claims made using thick terms. A more nuanced analysis than the one considered

in sections 3–4 might say that the conventional meaning of a thick concept *T* is such that whether the proposition expressed by a sentence of the form ‘*x* is *T*’ has evaluative content is determined by context of utterance.⁴⁹ Factors which determine a context of utterance include who is speaking, when, where, to whom, and with what intentions; what has been said before; and what presuppositions are being made.⁵⁰

The form of contextualism which the behavior of OTCs in denials would seem to require says that contextual factors determine the extension of terms such as ‘chaste’ as follows: on some occasions of use something satisfies ‘chaste’ only if it has properties which signal dedication to not being sexually provocative, whereas on other occasions of use it satisfies ‘chaste’ only if it also is good or praiseworthy in some way for having those properties.⁵¹ So, generalizing, whether a sentence of the form ‘*x* is *T*’ is associated with a certain evaluative content *E* is determined by contextual factors, but when the two are thus associated, *E* is part of the truth-conditions of the sentence.⁵² The hope for contextualism is to explain why ‘chaste’ has evaluative content in certain utterances but not others by pointing to relevant differences in factors that determine context of utterance. This strategy requires that context must be able to select evaluative standards or ideals as part of what determines the extension of thick terms. For what other kind of contextual factors could determine evaluative truth-conditions?

The most straightforward contextualist explanation of why an ordinary utterance of (5) is associated with evaluative content is that such content is part of the truth-conditions of (5) and that this is because the speaker’s evaluative standards imply it. This explanation has two flaws. First, it predicts that (5) is true, albeit relative to standards regarded as objectionable. If (5) were true, the negation in the first half of (13) and (14) should sound at least odd — but typically it doesn’t. Contextualists might treat the denial of evaluative content in the second half of (13) and (14) as directed not at the proposition expressed but at the standards which help determine what proposition is expressed. But it isn’t clear that denials directed at factors which determine truth-conditions count as instances of metalinguistic negation. Second, the explanation requires that the relevant standards depend on the speaker’s attitudes or intentions in such a way that it is possible to shift that feature of context by making one’s intention to do so sufficiently explicit. But many utterances of (5) are associated with evaluative content like (17) or (18) even if the speaker doesn’t endorse or intend to convey it. So the autonomy which the evaluative contents associated with thick concepts enjoy with respect to the intentions and attitudes of particular speakers constrains which contextual factors the relevant kind of contextualism can allow to determine truth-conditions.

Another contextualist explanation is that ordinary utterances of (5) have no evaluative truth-conditions: they are true, but relative to context which has no relevant evaluative factors in play. The first problem above applies also to this explanation, but the potential reply isn’t available to it. Contextualism is also redundant in it. When (5) expresses a non-evaluative proposition, explaining

why the utterance still conveys evaluative content in those many cases where it plainly seems to do so would have to invoke independent pragmatic mechanisms. Finally, the explanation must apply only in a limited range. If the proposition expressed by (5) isn't partly evaluative even when uttered by the Pope to a parish priest while discussing what Brad and Janet get up to, when is something like (17) or (18) ever going to be part of the truth-conditions of (5) and not merely pragmatically conveyed?

Thus it seems that any satisfactory contextualist explanation must appeal to contextual factors which are either objective or highly selective with respect to the attitudes of particular speakers.⁵³ The relevant standards might be, for instance, the correct standards, whatever they may be.⁵⁴ So something like (17) or (18) is part of the truth-conditions of (5) only if the correct standards entail it, not simply because the speaker's standards do. Disagreement about (5) is then explained as disagreement in substantive beliefs concerning what the correct standards are.

This contextualist view is also problematic. Either (5) is true, or it is false, relative to the correct standards. If it is true, then the problems with the first contextualist explanation arise again. But if (5) is false (or taken to be so by someone who regards *chaste* as objectionable), then again denials of (5) that use negation truth-conditionally should sound fine and concessive denials should sound odd. (16), for instance, should sound odd rather than fine or even unclear, because if (5) is false, why should one concede anything to it in the way that 'yes, but . . .' constructions seem often to do? The same dilemma arises for other contextualist variants which either further idealize the speakers' or hearers' evaluative standards or fasten on to various objective contextual factors.

Concessive denials like (15) and (16) raise another problem for contextualism. Either (5) as uttered in a context has some evaluative content as part of its truth-conditions or it doesn't. If it does and concessive denials echo this, then contextualism predicts that concessive denials are contradictory: they both concede and deny the evaluative content. And if it does but concessive denials don't echo this, then the reply talks past (5) instead of engaging with its content in a way we expect to see in cases of genuine disagreement. On the other hand, if (5) as uttered in a context doesn't have evaluative truth-conditions, then it again becomes unclear when contextualism is ever going to imply that an utterance of (5) expresses an evaluative proposition.

I can see only two replies to this problem available to contextualism. The first is to claim that uttering concessive denials like (15) or (16) can retrospectively shift the context of the previous utterance of (5) from one where evaluative standards enter into determining the extension of 'chaste' to one where they don't. Such a backwards influence seems mysterious and *ad hoc*. The second is again to claim either that concessive denials are too unclear to count as part of the data or that, if they do so count, it is alright to speak concessively for pragmatic reasons.

A further problem for contextualism concerns how to predict the readings of (12)–(14) in which negation takes scope over the evaluative content associated

with ‘chaste’ but not its non-evaluative implications. According to contextualism, utterance of (5) sometimes expresses some such evaluative proposition as (40):

- (40) Abstinence from extramarital sex has properties which signal dedication to not being sexually provocative, and things are good in a certain way for having those properties.

When (5) expresses a proposition like (40), negation is naturally read as taking scope over the whole proposition. For the entailments of a sentence are generally not preserved under negation. But if so, then (12), which concedes to (5) a certain non-evaluative entailment of (5), should sound at least odd, and the second half of (13) and (14) should sound redundant. Neither prediction is borne out by the data to which I have been pointing.

Another further problem is that contextualism seems to contribute little if nothing to explaining the asymmetries illustrated in (19)/(20) and (12)/(21) between the evaluative contents associated with thick concepts and their non-evaluative implications. When (5) expresses a proposition like (40), these contents aren’t asymmetrical in the relevant way. And when (5) expresses a non-evaluative proposition, any evaluative content will be conveyed by some independent pragmatic mechanisms, and so its behavior won’t be explained by contextualism.

Finally, the present form of contextualism requires some account of the relationship between thick and thin concepts. For although it isn’t implausible that various thin evaluative terms are context-sensitive in various ways, none seem to be context-sensitive with respect to whether they have semantic evaluative content at all.

Much more could and has been said about how different kinds of contextualism account for various kinds of disagreement. (The same applies to semantic relativism, which I haven’t tried to consider here.) But even on the present showing I can conclude that while the idea that the evaluative contents associated with thick concepts are contextually determined might explain some of the behavior of OTCs in denials, many serious problems remain with such contextualism — perhaps more so than with presupposition and implicature views.

8. Conclusion

This paper raises more questions than it answers. But it shows that theorizing about thick concepts cannot ignore the way in which thick concepts which are regarded as objectionable in a certain sense behave in various kinds of denials. Their behavior suggests that the evaluative contents associated with thick concepts aren’t generally conceptually entailed by them. It seems to suggest a sense in which the evaluative and non-evaluative features of thick concepts can be disentangled, without those evaluative features being essential to them. Their

behavior is amenable to rigorous treatment with tools from the philosophy of language, thereby providing a methodologically fresh opening into these much disputed issues in the literature on thick concepts. And yet, explaining their behavior also turns out to be difficult if those evaluative contents are thought of as presupposed or implicated by expressions involving thick concepts or as part of the contextually determined truth-conditions of such expressions. It is unclear what other options there are for theoretically locating the evaluative contents associated with thick concepts. Thus, even if the concepts which are usually classified as thick may not be homogeneous with respect to which analysis fits them best, there is a wide and important range of thick concepts which presently lack a satisfactory account. How they are accounted for will bear consequences for such important issues as whether thick and thin concepts are evaluative in the same kind of way and what it is for a concept to be evaluative to begin with. Thus it seems safe to say that the phenomenon of objectionable thick concepts provides fertile ground for further work on evaluative concepts.

Notes

*Thanks to audiences at University of Leeds, University of Nottingham, and the Thick Concepts conference at University of Kent, Canterbury, for stimulating discussions, and to Simon Kirchin for the invitation to speak at the conference. Special thanks are due to Carrie Jenkins, Andrew McGonigal, Robert May, and especially Daniel Elstein and Adam Sennet for helpful conversations, and to Matti Eklund, Brent Kyle, and Adam Sennet for very useful written comments. I acknowledge support from the European Community's Seventh Framework Programme (FP7/2007–2013) under grant agreement n° 231016.

1. An interesting further question is whether *can* individuates *ought* in the way that the substantive non-evaluative content of *cruel* or *generous* seems partly to individuate them. This might matter to whether some concepts commonly regarded as thin are to some degree thick.
2. See Eklund (ms).
3. The occasion was the 1895 Wilde vs. Queensberry libel trial. A transcript is available at <http://www.law.umkc.edu/faculty/projects/ftrials/wilde/WildeLibelTranowcross.html> (accessed on July 31, 2009).
4. I cannot here pursue the interesting question whether thin concepts are open to being regarded as objectionable in this sense, and if not, why. There are different senses in which thin concepts may be regarded as objectionable. For instance, the claim that the concepts *duty* and *obligation* require a law-giver where there is none is an objection to normative distinctions which presuppose a certain kind of source, not to any particular grounds for drawing normative distinctions whatever their source. A similar point would apply to a more general sort of error theory of morality.
5. Thanks to Matti Eklund and Brent Kyle for providing care packages here.
6. See especially Hom (2008).
7. Blackburn (1992, 1998) and Gibbard (1992) are honorable exceptions.

8. It would be sheer magic if the location of evaluative content were determined in one way for all thick concepts which in fact are objectionable and in another way for those which aren't.
9. For another discussion with a similar tenor, see Gibbard (1992: 274–8).
10. This difference matters to my discussion because it bears on who takes the lead in determining how a word is being used in a particular discourse, and that in turn can matter to what kind of responses are conversationally appropriate. Thanks to Adam Sennet here.
11. Of course it also matters what concepts are. For a quick run-through of some different notions of concepts, see Williamson (2007: 13–17, 29–30).
12. Much more needs to be said about this assumption. It should accommodate the possibility of using thick and thin terms non-evaluatively merely to classify things and the observation, due to Simon Kirchin, that words which are typically used to express thick concepts can be used to make thin evaluations, as with such interjections as 'Neat!' or 'Wicked!' (I am inclined to regard these as mere slang or idiolect meanings; compare 'Rad!') A bigger question lurking in the background which I won't be able to address here concerns whether the mechanisms of language are such that words could semantically express genuinely thick concepts. This depends on such issues as whether the main semantic property is a connection to truth, dividing up possible worlds, or what have you.
13. Similar observations can be found in Gibbard (1992: 281–2).
14. Attributed to politician Alun Cairns on <http://news.bbc.co.uk/1/hi/wales/7454545.stm>.
15. The tone of voice one uses when uttering 'greasy wops' in (9) can effect a way to express disagreement with (8) by uttering (9), namely when it effects a metalinguistic use of negation. See section 3.
16. Some speakers report that they wouldn't be willing to say that a claim like (5) is true, but misleadingly put (e.g. Gibbard 2003: 168). Others don't find this objectionable. How the former reaction bears on (15) and (16) depends on whether these sentences can at best amount to saying 'True, but I wouldn't put it that way'; the latter bears on how bad it would be if they did.
17. This is perhaps no surprise, given the autonomy which the evaluative contents associated with thick concepts seem to enjoy with respect to the intentions and attitudes of particular speakers.
18. This is a simplification. The evaluative content associated with thick concepts might be a truth-conditional aspect of theirs, but not solely because of their conventional meaning, if a certain form of contextualism is true (see section 7). Note also that the claim isn't that the evaluative contents associated with thick concepts aren't truth-apt. For all I say, they might or might not be.
19. For instance, as noted earlier, one hypothesis is that it is alright to say 'yes' for pragmatic reasons, such as that 'yes, but . . .' can be a polite way of coordinating conversations that involve disagreement.
20. The *locus classicus* on the metalinguistic use of negation is Horn (1989: ch. 6).
21. In reality, the issue is more complicated. For instance, no standard test for metalinguistic negation is uniformly reliable (see e.g. Carston 1996 and Geurts 1998). But this means that even if (13) and (14) pass only some of these tests, this

doesn't undermine the metalinguistic explanation. Another source of support is that (13) and (14) look like "garden-path utterances": when an audience reaches the second half, they often realize that in order for the sentence to correctly understood, they need to go back and reanalyze the negation in the first half. Such garden-pathing and need for pragmatic reanalysis are typical characteristics of metalinguistic negation.

22. See e.g. Foot (1958), McDowell (1981), Williams (1985), and Dancy (1995).
23. I generalize because the same behavior will be exhibited by many concepts which are regarded as objectionable by some but in fact aren't objectionable.
24. It seems fair to attribute this kind of view to Hare (1952) and Brower (1988).
25. Those in favor include Gibbard (1992, 2003: 165–8), Blackburn (1992, 1998: ch. 4), Sreenivasan (2001), and Richard (2008: 30–3). Those against include McDowell (1981) and Williams (1985: chs. 7–8).
26. Those in favor include Hare (1952: ch. 7), Blackburn (1992, 1998: ch. 4), and Hurka and Elstein (forthcoming). Those against include McDowell (1981), Williams (1985: chs. 7–8), and Dancy (1995).
27. Alternatively, if some relevant form of contextualism is true (see section 7), they would have conventionally determined truth-conditional content and contextually determined truth-conditional content, and evaluative content would be of the latter sort.
28. One proposal which I won't discuss is worth mentioning briefly, because it is radical and I don't know what to say about it. This is the view that it is indeterminate whether the evaluative content associated with a thick concept *T* is semantically or pragmatically associated with utterances or sentences of the form '*x* is *T*.' (Thanks to Carrie Jenkins for suggesting this view to me.) Such indeterminism may be well-suited to explain the puzzling patterns in the examples of section 2. It is puzzling how one can express disagreement with a claim both by using negation and by using such concessive constructions as 'yes, but . . .' If it is indeterminate whether the evaluative contents associated with thick concepts are conventionally or pragmatically conveyed, it would be no surprise if we were inclined to say sometimes that denying those evaluative contents doesn't involve contradicting another speaker and sometimes that it doesn't. Indeterminism can also explain how thick and thin concepts might be evaluative in the same kind of way. The thought would be that whichever evaluative features are semantic and pragmatic, these features are either determinately semantic or determinately pragmatic in the case of thin concepts, but only indeterminately so in the case of thick concepts. But I find it hard to assess indeterminism with any great confidence because it isn't clear to me how to assess theories of this general type.
29. There may be others. Discussions of these proposals should bear in mind that it may often be hard to discern whether a particular content is asserted, presupposed, or implicated. It should also be noted that different theorists may variously classify one and the same thing as an implicature or presupposition.
30. For one presuppositional analysis, see Gibbard (1992). The more complex of the two patterns of "reductive" analysis of thick concepts defended by Hurka and Elstein (forthcoming) embeds an inference that can (but needn't) be construed as a presupposition.

31. See e.g. Stalnaker (1974), Levinson (1983), Chierchia and McConnell-Ginet (1990), and Kadmon (2001).
32. Compare this exchange: A: “Sam doesn’t realize that he is late.” B: “Yes, but he isn’t late.” What ‘yes, but . . .’ indicates is that B denies that Sam is late but agrees that Sam doesn’t think he is late. Insofar as B’s reply sounds odd, this may simply be induced by presupposition failure.
33. Some: Chierchia and McConnell Ginet (1990). All: Levinson (1983). What is more, even some pragmatic presuppositions seem to be lexically triggered. See especially Kadmon (2001: ch. 11).
34. It should be clear that most words that express thick concepts fall into none of these other categories of words or constructions which are generally recognized as presupposition triggers.
35. On presupposition accommodation, see Lewis (1979).
36. Richard (2008: 21) makes basically this point about racial and ethnic slurs.
37. Thanks to Matti Eklund for raising this question. As he notes, someone who takes this line could explain the appearance of a spectrum from thick to thin by saying that it is indeterminate whether the borderline cases are thick or thin concepts. Eklund (ms) raises the general issue in the context of a conventional implicature account of the evaluative contents of thick concepts.
38. The hypothesis is different from, but compatible with, the idea that evaluative terms both express properties and implicate that the speaker holds certain conative attitudes towards the bearers of those properties, whether conventionally (Copp 2001) or conversationally (Finlay 2005).
39. On generalized conversational implicature, see Grice (1975: 37) and Levinson (2000). For a recent sophisticated account of conventional implicature, see Potts (2005). Williamson (2009) defends a conventional implicature analysis of pejorative expressions, but as he notes (2009: 151 n. 14), the relationship between pejoratives and evaluative terms is a large question that I leave open as well. Simon Blackburn may be best read as suggesting a generalized conversational implicature analysis of thick concepts. He holds that the evaluative contents associated with thick concepts aren’t a matter of conventional meaning and can usually be communicated by intonation and tone, but that such effects can nonetheless be predictable (Blackburn 1992: 297–8). It also seems plausible to read R. M. Hare as an implicature theorist. Hare holds that thick concepts have a “primary meaning” that is non-evaluative and a “secondary meaning” that is prescriptive or evaluative. But what he says about this secondary meaning makes it look like either a generalized conversational implicature or a conventional one (see Hare 1952: 118–22).
40. On reinforceability, see Sadock (1978).
41. The sense in which conversational implicatures are supposed to be cancelable should be distinguished from a broader sort of deniability. For instance, the factivity of ‘before’ is deniable in a certain sense. ‘Sally died before she finished her thesis’ can be true and felicitous although it falsely presupposes that Sally finished her thesis (Kadmon 2001: 210; Potts 2005: 22–4). This raises the question whether it is possible to utter things like ‘This year’s carnival wasn’t lewd,’ and not merely as a rhetorical trope, and yet deny that the carnival would have been the worse for it otherwise (cf. Blackburn 1992). One explanation of how the

evaluative contents associated with thick concepts might be deniable in this way would be the hypothesis that such contents can remain merely potential rather than actual contributions to context of utterance. It would be satisfying to see rigorous attempts to develop and assess such suggestions.

42. Finlay (2005: 14) makes this point against Copp's (2001: 35) appeal to inverted commas to defend the detachability of evaluative attitudinal content from moral terms. Insofar as inverted commas can function to overcome expressive limitations, I doubt the significance of Finlay's quibble that inverted commas fail to produce a *different* co-extensive term.
43. Gibbard uses this example for a different purpose and takes 'wumpua' and 'brutal' to be almost but not quite co-extensive. But the scenario has a coherent variant in which they are co-extensive.
44. A further worry about the conversational implicature view is that the interaction of OTCs with different uses of negation seems systematic and predictable but the behavior of conversational implicatures under negation is unpredictable, since they are calculated on the basis of the whole content of what was said in the given context.
45. Thus it seems unclear what exactly the right inverted commas reading of (15) and (16) would be and whether it would conflict with the reading described in section 2. This bears on doubts about whether (15) and (16) constitute the kind of evidence which I suggest they provide.
46. The scalar implicature that is denied here is *John had exactly six pints*.
47. These analyses have some plausibility regarding Blackburn's example of using 'fat↓' ('fat' pronounced with a sneer) to convey that those with excessive body fat are contemptible (Blackburn 1992: 289–91). Sneer may well have a quasi-conventional or generalized function of conveying contempt. While the status of any particular tone is an empirical matter, it seems that if hearers didn't think that the "fattist" idiolect uses sneer to convey contempt, using 'fat↓' to convey it would require more from context than using expressions such as 'gross' for this purpose. Incidentally, this may tell against Blackburn's claim that fattists could perfectly well jettison 'gross' for 'fat↓' (1992: 290).
48. See Geurts (1998: 279–81). Contrastive negation emphasizes a part of the negated sentence and contrasts it with a corresponding part in an affirmative statement; one form in English is *Not X but Y*. Geurts also claims that presuppositional denials are typically non-contrastive. This might support a presuppositional analysis. But the import of such contrastive examples as 'Abstinence from extramarital sex isn't chaste but prudish' and 'The early Almodóvar movies aren't lewd, they are exhilarating' is unclear.
49. Literature on the semantics of modal expressions, including deontic modals like 'ought,' understands contextualism more narrowly, as the view that modals are quantifiers over possibilities which include parameters for contextually determined values (such as what possibilities are in the domain and which of those possibilities are best, in the case of 'ought'). But few if any terms that express thick concepts are modal expressions.
50. This notion of context of utterance follows Lewis (1979). Other instances of this kind of contextualist hypothesis are that context determines to whom the word 'I' refers (this is determined by who is speaking) and that context determines

- how costly something must be to fall within the extension of 'expensive' (this is determined by factors such as whether the speaker or his audience have just been talking about paper clips, vintage guitars, or health care reform proposals).
51. Contextualism should perhaps also allow the possibility that on some occasions something satisfies 'chaste' only if it is bad or condemnable in some way for having those properties.
 52. At least some thick terms are likely to be context-sensitive also in another way. Many of them are gradable adjectives (for instance, 'chaste,' 'greedy,' 'kind,' and 'courageous') and many gradables are context-sensitive (for instance, 'tall,' 'loud,' 'athletic,' and 'expensive'). But the respect in which gradable expressions typically are context-sensitive only concerns the degree to which something must exhibit certain qualities to count as satisfying the given term in context. (I discuss this in the context of thick epistemic concepts in Väyrynen 2008.) It doesn't in any obvious way concern whether this quality is evaluative in some contexts but not so in others.
 53. 'Highly selective' because (5) can easily come out as true in many contexts even if the relevant standards are those given by what really matters to the speaker or by his fundamental ideals.
 54. The autonomy of evaluative content can allow speaker intentions to determine this.

References

- Blackburn, Simon (1992). "Through Thick and Thin." *Supplement to the Proceedings of the Aristotelian Society* 66: 284–99.
- Blackburn, Simon (1998). *Ruling Passions*. Oxford: Oxford University Press.
- Brower, Bruce (1988). "Virtue Concepts and Ethical Realism." *Journal of Philosophy* 85: 675–93.
- Carston, Robyn (1996). "Metalinguistic Negation and Echoic Use." *Journal of Pragmatics* 25: 309–30.
- Chierchia, Gennaro, and McConnell-Ginet, Sally (1990). *Meaning and Grammar. An Introduction to Semantics*. Cambridge, MA: MIT Press.
- Copp, David (2001). "Realist-Expressivism: A Neglected Option for Moral Realism." *Social Philosophy and Policy* 18: 1–43.
- Dancy, Jonathan (1995). "In Defence of Thick Concepts." *Midwest Studies in Philosophy* 20: 263–79.
- Eklund, Matti (ms). "Some Problems Regarding Thick Concepts." Unpublished.
- Finlay, Stephen (2005). "Value and Implicature." *Philosophers' Imprint* 5(4): 1–20.
- Foot, Philippa (1958). "Moral Arguments." *Mind* 67: 502–13.
- Geurts, Bart (1998). "The Mechanisms of Denial." *Language* 74: 274–307.
- Gibbard, Allan (1992). "Thick Concepts and Warrant for Feelings." *Supplement to the Proceedings of the Aristotelian Society* 66: 267–83.
- Gibbard, Allan (2003). *Thinking How To Live*. Cambridge, MA: Harvard University Press.
- Grice, H. P. (1975). "Logic and Conversation." Reprinted in his *Studies in the Way of Words*, 22–40. Cambridge, MA: Harvard University Press.
- Hare, R. M. (1952). *The Language of Morals*. Oxford: Clarendon Press.
- Hom, Christopher (2008). "The Semantics of Racial Epithets." *Journal of Philosophy* 105: 416–40.
- Horn, Laurence R. (1989). *A Natural History of Negation*. Chicago, IL: University of Chicago Press.

- Hurka, Thomas, and Elstein, Daniel Y. (forthcoming). "From Thick to Thin: Two Moral Reduction Plans." *Canadian Journal of Philosophy*.
- Kadmon, Nirit (2001). *Formal Pragmatics: Semantics, Pragmatics, Presupposition, and Focus*. Oxford: Blackwell.
- Levinson, Stephen C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Levinson, Stephen C. (2000). *Presumptive Meanings: A Theory of Generalized Conversational Implicature*. Cambridge, MA: MIT Press.
- Lewis, David (1979). "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339–59.
- McDowell, John (1981). "Non-Cognitivism and Rule-Following." In *Wittgenstein: To Follow A Rule*, ed. S. Holtzman and C. Leich, 141–62. London: Routledge and Kegan Paul.
- Potts, Christopher (2005). *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.
- Richard, Mark (2008). *When Truth Gives Out*. Oxford: Oxford University Press.
- Sadock, Jerry (1978). "On Testing for Conversational Implicature." In *Syntax and Semantics 9: Pragmatics*, ed. P. Cole, 281–97. New York: Academic Press.
- Scheffler, Samuel (1987). "Morality Through Thick and Thin: a Critical Notice of *Ethics and the Limits of Philosophy*." *Philosophical Review* 96: 411–34.
- Slote, Michael A. (1975). "Inapplicable Concepts." *Philosophical Studies* 28: 265–71.
- Sreenivasan, Gopal (2001). "Understanding Alien Morals." *Philosophy and Phenomenological Research* 62: 1–32.
- Stalnaker, Robert C. (1974). "Pragmatic Presuppositions." In *Semantics and Philosophy*, ed. M. Munitz and P. Unger, 197–214. New York: New York University Press.
- Väyrynen, Pekka (2008). "Slim Epistemology with a Thick Skin." *Philosophical Papers* 37: 389–412.
- Williams, Bernard (1985). *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.
- Williamson, Timothy (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.
- Williamson, Timothy (2009). "Reference, Inference, and the Semantics of Pejoratives." In *The Philosophy of David Kaplan*, ed. J. Almog and P. Leonardi, 137–58. Oxford: Oxford University Press.