Philosophers' Imprint

VOLUME 13, NO. 1
JANUARY 2013

How Are Thick Terms Evaluative?

Brent G. Kyle
United States Air Force Academy

© 2013 Brent G. Kyle
This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 3.0 License.
<www.philosophersimprint.org/013001/>

I. Introduction.

Many influential ethicists of the twentieth century turned their attention to thick concepts, like *courageous*, *dishonest*, and *generous* (Murdoch 1971; Foot 1958; Hursthouse 1995). Some ethicists even urged us to stop focusing as much on thin concepts, like *good* and *wrong*, and to expand or shift our attention towards the thick (Anscombe 1958; Williams 1985; Lovibond 1983; Platts 1979). But what is the supposed significance of thick concepts? Very briefly, thick concepts are said to combine descriptive and evaluative elements and have thereby provided focal points for a cluster of related issues, such as whether there is a fact-value gap, whether evaluative language is truth-apt, and whether the evaluative can be reduced to the descriptive.

But critics have responded by downplaying the importance of thick concepts within ethics, and they've done so by arguing that thick concepts are not genuinely evaluative (Brower 1988) or, similarly, that thick terms do not express evaluative meanings (Blackburn 1992). Simon Blackburn, for example, has declared that thick terms "are of no great importance to the theory of ethics", and he has done so by arguing that such terms do not have evaluative meanings (1992, 285).

Contrary to this skepticism, I shall argue that thick terms indeed have evaluative meanings (section II). I call this the Semantic View. Proponents of this view have provided very little argumentative support in its favor. But an argument is certainly needed, since the Semantic View is by no means obvious and has been attacked on many occasions (*e.g.*, Blackburn 1992; Brower 1988; Väyrynen 2009). Although space prohibits discussion of all of these attacks, I shall defend the Semantic View against Pekka Väyrynen's recent challenge arising from objectionable thick concepts (section III).

What exactly is at issue between those who affirm and those who deny the Semantic View? Most ethicists are willing to grant that thick terms are *somehow* associated with evaluations, but they tend to disagree about what exactly this relationship is. The issue for debate is not *whether* but *how* thick terms are associated with evaluations. For example, is a thick term's evaluation pragmatically associated with it?

Or is the evaluation semantically associated with it? The view I shall advance is a specific version of the latter.

Just how is this semantic relationship to be characterized? One way to formulate the Semantic View is to say that sentences containing thick terms have evaluative truth-conditions. Or, similarly, we might say that thick concepts conceptually, analytically, or semantically entail evaluative contents. It makes little difference which of these formulations we choose. So, I shall fix on the following:

> Semantic View: Many thick concepts (if not all) conceptually entail evaluative contents.[1]

1. Two assumptions ought to be addressed briefly. First, the Semantic View assumes there is a suitable way of distinguishing between evaluative and non-evaluative content. But philosophers like Foot, Murdoch, and Williams are commonly believed to have used thick concepts to undermine such a distinction (see Millgram 1995 for doubts about this attribution). Nevertheless, it will make no difference to the substance of my argument whether we accept this distinction. The Semantic View can be reformulated in a way that does not assume a distinction between evaluative and non-evaluative content:

   > Modified Semantic View: Many thick concepts (if not all) conceptually entail the contents expressed by thin terms.

   The thick/thin terminology is typically accepted by those who reject the distinction between evaluative and non-evaluative content (*e.g.*, see Jackson 1998, 135–36, and Williams 1995, 240). In principle, it should be possible to draw a distinction between such terms without committing to a distinction on the level of content. My arguments in this paper will support the Modified Semantic View just as effectively as they support the original formulation. Thus, it will make no difference to the substance of my argument whether we accept a distinction between evaluative and non-evaluative content, since there is a suitable reformulation of the Semantic View that does not rely on any such distinction.

   And second, the Semantic View also assumes a controversial relation — conceptual entailment. In this paper, I do not defend this relation, but I also do not rely on any of the weightier assumptions that are often associated with it — *e.g.*, that it is best understood in terms of synonymy, or that conceptual truths are an epistemologically privileged class (see Williamson 2007, 48–133, for a recent critical survey). Furthermore, it's worth emphasizing that there are also skeptics about nearly all of the linguistic relations appealed to by alternative views. Many have tried to reduce or eliminate relations like conventional implicature (Bach 1999) and presupposition (Atlas and Levinson 1981; Wilson 1975). Nevertheless, for the purposes of this paper, I shall adopt the permissive policy of taking all semantic and pragmatic relations as generally acceptable.

In the next section, I advance my positive argument for this claim. The parenthetical qualification within the Semantic View—*i.e.*, "Many thick concepts (if not all)…"—will be explained in due course (section II.5). In section III, I take up the issue of objectionable thick concepts.

II. An Argument for the Semantic View.

II.1 Good and Bad in a Way

My argument for the Semantic View is an inference to the best explanation of certain linguistic data. After presenting this data (II.1 and II.2), I propose an instance of the Semantic View as my favored explanation (II.3) and then proceed to reject various alternative explanations of that data (II.4).

The linguistic data in question crucially involves two thin evaluative concepts — namely the concepts *good in a way* and *bad in a way*. These concepts are often employed in ordinary conversation. Suppose you're ultimately against the new health care bill, but you nonetheless believe it has certain merits. In this case, you might defend it against an uncharitable attack by saying

> You're right that the bill should be repealed. But I want to point out that it is good in a way — it will lower insurance costs in the long run.

Here the concept *good in a way* is employed in a rather ordinary way; similar things can be said for *bad in a way*. In what follows, I will inquire into exactly how these concepts are related to the thick.

Many thick terms seem to bear a close relationship to either *good in a way* or *bad in a way*.[2] But it's not obvious what exactly this relationship is. At the very least, it seems that the typical utterance of many thick

2. A possible exception, involving multivalent thick terms (*e.g.*, 'eccentric' and 'kinky'), is discussed in section II.5.

terms commits the speaker to a claim involving either good in a way or bad in a way. For example, an utterance of 'Nancy is generous (loyal/kind/courageous)' tends to commit the speaker to the claim that Nancy is good in some way. Similarly, an utterance of 'Nancy is rude (lewd/brutal/unkind)' typically commits the speaker to the claim that Nancy is bad in some way. Of course, in neither case is the speaker obviously committed to the stronger claim that Nancy is good or bad overall, but the speaker is certainly committed to the weaker claim that she is good or bad in some way or other.

To bring this relationship into perspective, let's focus on the thick term 'generous' and consider its relationship to 'good in a way'. As just noted, it seems clear that a typical utterance of

(1) Nancy is generous.

commits the speaker to the truth of

(2) Nancy is good in a way.

But how do we explain this connection? As I'll argue, the connection is best explained by the relevant instance of the Semantic View. That is, we can best explain this by claiming that (2) is a conceptual entailment of (1). But before I state my reasons for accepting this, it's useful to see the faults of another potential way of explaining this connection — namely conversational implicature. As we'll see, the shortcomings of this explanation reveal the advantages of the Semantic View.

II.2. Conversational Implicature

If a sentence S1 conversationally implicates S2, then S1 (or the fact that someone utters S1) normally gives the hearer a defeasible reason to conclude that the speaker also means to convey S2. For example, my utterance of 'Some students attended' normally gives you reason to conclude that I also mean to convey 'Not all students attended.' Typically, however, the connection between S1 and S2 can be reinforced and canceled without oddity. And this is precisely the problem with the present explanation. The connection between (1) and (2) is neither reinforceable nor cancelable.

Let's begin with reinforceability. The basic idea is that conversational implicatures can normally be made explicit without awkward redundancy (Sadock 1978, 295). I can reinforce the above implicature by saying 'Some students attended, but not all', and this does not sound odd. But nothing similar is true for the relationship between (1) and (2). Consider the following attempt:

(3) ? Nancy is generous {and/but} she's good in a way.[3]

This sounds awkward and redundant. So, it seems unlikely that the connection between (1) and (2) is reinforceable.

The second main feature of conversational implicature is cancelability (Grice 1989, 44). In typical contexts, the fact that I utter 'Some students attended' gives you reason to conclude that I also mean to convey 'Not all students attended.' But this implicature can be canceled — in other contexts, you may have no reason to draw this conclusion. Suppose I say, 'Some students attended; in fact, all of them did.' In this case, my addition of 'in fact, all of them did' cancels the implicature associated with the first part. And my utterance seems normal. But again, nothing similar is true regarding the connection between (1) and (2). Consider an attempt at cancelation that parallels the above example:

(4) # Nancy is generous; in fact, she's not good in any way.

This sentence seems highly odd, unlike the previous example of cancelation. So, although (1) implies (2) in some sense, this relationship appears to be neither cancelable nor reinforceable. We therefore cannot explain this connection by appeal to conversational implicature.

3. When '?' appears before a sentence, this signifies that the sentence is at least somewhat odd. And when '#' appears, this signifies that the sentence is highly odd.

II.3. The Semantic View of 'Generous'

Although I will consider other rival explanations in a moment (section II.4), I first want to show that the problems for the conversational implicature explanation are easily handled by a particular instance of the Semantic View. Since *good in a way* is clearly evaluative, the following would be an instance of the Semantic View:

> SV-Generous: The thick concept *generous* conceptually entails *good in a way*.[4]

This view straightforwardly explains the relationship between (1) and (2) by claiming that (1) conceptually entails (2). But how well does SV-Generous explain the sentences that led us to reject the appeal to conversational implicature?

SV-Generous easily explains the fact that (3) seems awkward. According to SV-Generous, the second part of (3) is conceptually entailed by the first part. Thus, the second clause is redundant, and that's why (3) seems odd. Why does (4) seem so highly odd? According to SV-Generous, the second part of (4) contradicts what is conceptually entailed by the first part, and that's why (4) seems so odd. Thus, SV-Generous explains all the data we've seen thus far.

SV-Generous treats 'generous' and 'good in a way' similarly to how we typically treat 'bachelor' and 'unmarried'. And it's worth emphasizing that there is further linguistic data to support this correlation. In particular, the following conjunction seems highly odd:

(5) # Nancy is generous, and she's not good in any way.

And (5) seems nearly as odd as

(6) # Jack is a bachelor, and he's not unmarried.

And much like that of (6), the oddity of (5) cannot be solely attributed to either one of its conjuncts. After all, these seem to be quite normal:

(7) Nancy is generous.

(8) She's not good in any way.

Since (7) and (8) are felicitous in their own right, it follows that neither conjunct within (5) is by itself responsible for the overall inappropriateness of (5). Clearly it's something about their combination that strikes us as highly odd. And SV-Generous explains this by allowing us to claim that the conjunction of (7) and (8) generates a contradiction.[5]

4. It should be noted that SV-Generous does not entail the controversial thesis that the description and evaluation of *generous* can be "disentangled" (see McDowell [1981] for the seminal paper on this). Nothing about SV-Generous entails that we could exhaustively state the descriptive (or non-evaluative) content of *generous*. Neither does it entail that we could exhaust its evaluative content, since *good in a way* need not be seen as its only evaluative content.

   On the other hand, some people may hold a semantic-type view about the evaluation associated with 'generous' but reject my supposition that 'generous' inherits its evaluation from an associated thin concept like *good in a way*. On this view, thick evaluation is sui generis. This alternative view deserves serious consideration in its own right, but it remains unclear how such a view can explain the data discussed in this paper regarding the connection between 'generous' and 'good in a way'. If I am right that SV-Generous provides the best explanation of that data, then it appears that this alternative view does not adequately characterize the way in which 'generous' is associated with evaluation. Moreover, as I note in section II.5, the same goes for other thick terms like 'courageous', 'murder', 'brutal', and 'kind'. However, in that section, I concede that my overall argument may not apply to multivalent thick terms (*e.g.*, 'eccentric', 'kinky', and 'unorthodox'). So, my argument in this paper allows for the possibility that this alternative view could claim some significant territory within the class of multivalent thick terms.

5. Instead of SV-Generous, can we simply hold that *generous* a priori entails *good in a way* (without holding that this entailment is conceptual)? It appears not — this weaker alternative does not adequately explain the awkwardness of the above sentences. For example, if we merely claim that the first part of (5) a priori entails what the second part denies, this would not by itself tell us why (5) is awkward. This is because there need be no awkwardness when one denies what is a priori entailed by the first part of one's utterance. Consider mathematical statements like

   (i) 85 men and 48 women attended, and there were no more than 123 men and women combined.

   The first part of this sentence a priori entails what's denied by the second part. But the sentence itself seems felicitous. The fact that sentences (3)–(5) are infelicitous suggests that they are somehow linguistically impermissible,

At this point, it will be objected that we can envision contexts in which an utterance of (5) would not seem odd. Imagine that the speaker of (5) is someone like Ebenezer Scrooge who disvalues generosity. Or, to make the example more vivid, imagine that Scrooge utters (5) within a community of speakers who also disvalue generosity. Within this context, it seems possible for Scrooge to utter (5) felicitously. The problem is that SV-Generous initially seems to predict that his utterance of (5) would express something contradictory.

This objection can be avoided once we clarify SV-Generous in the right way. In effect, SV-Generous asserts that the concept expressed by 'good in a way' is conceptually entailed by *generous*. But it's misleading to speak about *the* concept expressed by 'good in a way', since this phrase expresses many different concepts in different contexts of utterance. As I will suggest, Scrooge's utterance can be explained if we clarify SV-Generous in a way that is mindful of this context-sensitivity. Let me first explain how 'good in a way' is context-sensitive and then move on to the needed clarification of SV-Generous.

The primary mechanism responsible for the context-sensitivity of 'good in a way' is the constituent expression 'a way'. To see this, consider two contexts where 'a way' occurs:

and SV-Generous holds that this impermissibility is semantic in nature. But the apriority of these sentences does not seem to help out in explaining their infelicity.

One might object by pointing to certain a priori false sentences that are awkward. Consider

(ii) ? Nancy gives much time and money to people in need, and she is not good in any way.

For the reasons stated above, my reply is that the awkwardness of (ii) ought to be explained by other means (aside from merely citing a priori entailment). Ultimately, we may wish to conclude that the complex concept *gives much time and money to people in need* conceptually entails *good in a way*, provided there are no better ways of explaining the awkwardness of (ii). I am not opposed to this potential conclusion. Those who reject the possibility of conceptual entailments from non-evaluative to evaluative sentences might be troubled by this conclusion. But I am in no way committed to rejecting this possibility.

Gas Station Context

Traveler: "Is there a way to get from Ithaca to New Haven?"

Gas Station Clerk: "Yes, there is a way — just head east on Highway 17."

Bus Station Context

Traveler: "Is there a way to get from Ithaca to New Haven?"

Bus Station Clerk: "Unfortunately, there's not any way. You'll have to take the bus into NYC and then take the train to New Haven."

The bus clerk says "there's not any way", while the gas clerk says "there is a way". But it seems plausible that both of their utterances could be true.[6] Relative to different contexts, their utterances express different propositions. The bus clerk is asserting that there's no way by bus, whereas the gas clerk is merely asserting that there is a way by car.

In general, when we utter that there is (or is not) a way to do such and such, we are very seldom asserting that there is (or is not) a logically possible way. Typically, the scope of the utterance is restricted by a contextual parameter — in this case, a quantifier domain. Here's one way of representing how the above utterances are associated with domains:

> 'There is a way' is true relative to a context C if and only if there is some way w within the domain provided by C.[7]

6. Under some pragmatic views of context-sensitivity, we must say that the bus clerk's utterance expresses a false proposition, but he communicates something true nonetheless. In what follows, I assume a semantic theory of context-sensitivity, according to which they can both say something true relative to their own contexts. (See Stanley and Szabó [2000] for a discussion of these two views.) My assumption, however, will make little substantive difference to what I go on to say. There are ways of rephrasing my claims so as to reflect a pragmatic view instead of a semantic one.

7. This model is probably too simplistic to handle more complicated quantifier expressions. But it will suffice for current purposes. See Stanley and Szabó

Under this model, the bus clerk's utterance can be true relative to his context, even if the gas clerk is right about there being a way of getting to New Haven by car. Traveling by car is a way that is not contained within the domain of the bus station context.

Something similar can be said for 'good in a way'. When we utter that there is (or is not) a way in which A is good, we are very seldom asserting that there is (or is not) a logically possible way. Here again, the utterance is restricted by being contextually associated with a particular domain:

> 'A is good in a way' is true relative to a context C if and only if there is some way w within the domain provided by C such that A is good in w.

Under this model, the sentence 'She is good in a way' might be true relative to some contexts and false relative to others, depending on which ways of being good are contained within the domain of each context.

These considerations motivate a particular way of clarifying SV-Generous, one that is mindful of the context-sensitivity of 'good in a way'. Instead of referring to concepts, this clarified version refers to sentences whose truth-values can be relativized to contexts:

> SV-Generous Clarified: The sentence 'A is generous' conceptually entails that 'A is good in a way' is true relative to select contexts.

To put this in a slightly different way, we can say that, whenever a person is generous, it follows conceptually that there is a select class of contextual domains each containing a way in which that person is good.[8]

With this clarified version in mind, let us now return to the main objection. Does this view predict that Scrooge's utterance of (5) must be contradictory? It certainly does not predict this, provided we specify the relevant contexts in the right way. SV-Generous Clarified predicts only that (5) expresses a contradiction *within select contexts*. But, as long as Scrooge's context is not among those selected, the view in question does not predict that he would be expressing something false or contradictory in uttering (5). In other words, this view need not hold that the first part of (5) conceptually entails what the second part denies *within Scrooge's context*.

But what are the select contexts mentioned in SV-Generous Clarified? This can be answered if we use our linguistic data as a guide. If sentences like (3)–(5) sound odd within a given context C, then we should find a principled way of including C among the relevant contexts. Similarly, if those sentences are felicitous within some context C*, then we should find a principled way of excluding C* from the relevant contexts. It is far beyond the scope of this paper to fully specify the relevant class of contexts,[9] but I think it is clear that we could have a principled way of specifying this class. For example, let's suppose that (3)–(5) are awkward within all contexts except those in which a person like Scrooge is the speaker. In this case, we should hold that the relevant class includes all and only those contexts in which the speaker is not presumed to disvalue generosity. No doubt, this is oversimplified, but the basic strategy should be clear: first, we find a distinctive feature of all the contexts in which (3)–(5) are

(2000, 248–58) for a critique of this simple model and a development of a different semantic model.

8. This view might be trivially true if we select only contexts that have unrestricted domains—*i.e.*, domains that contain all logically possible ways of being good. That is, if we allow all logically possible ways to count as ways of being good (*e.g.*, being good for use in a philosophical discussion on goodness), then it might be conceptually true that everything is good in some way or other. See Thomson (2008, 10) for an argument for this. However, this difficulty can be avoided if we hold that the select contexts include contexts with restricted domains (as well as unrestricted ones).

9. This is partly because the contexts are numerous but also because there is bound to be disagreement and uncertainty about whether (3)–(5) are felicitous within certain contexts.

awkward, and then we specify that the relevant class includes all and only the contexts that have this feature.

Most likely, the select contexts will exclude all contexts in which the relevant conversational participants are presumed to disvalue generosity. But is it ad hoc to exclude these contexts? Is this exclusion motivated solely by the need to explain Scrooge's utterances? No. There are independent reasons for holding that the relevant contexts will have domains that don't contain the ways of being good that conversational participants are presumed to disvalue. To see this, consider sentence (8):

(8) She's not good in any way.

The truth-value of an utterance of (8) can depend on what conversational participants are presumed to disvalue. For example, suppose we are in a context in which it's known that all conversational participants disvalue being a good thief (*e.g.*, because we all have contempt for thieves in general). In this context, it seems that I could truthfully utter (8) even though the person I'm referring to is a good thief. This is because the domain of our context does not contain being a good thief among the ways of being good. Being a good thief is outside our domain, precisely because we're presumed to disvalue this way of being good. Thus, it is independently plausible that the relevant contexts have domains that do not contain the ways that conversational participants are presumed to disvalue. Moreover, since it is taken as common knowledge that Scrooge disvalues generosity, it is no surprise that certain of his contextual domains do not contain the ways of being good that are associated with generosity.

If the above is correct, then the felicity of Scrooge's utterance of (5) poses no problem for the clarified version of SV-Generous. When understood properly, this view predicts that (5) expresses something false and contradictory in most contexts. But it allows for the possibility that (5) might be true relative to certain contexts in which conversational participants are presumed to disvalue generosity. Thus, SV-Generous is able to explain the linguistic data that we've encountered.

II.4. Two Rival Explanations

So far, we've seen that SV-Generous can explain the relationship between (1) and (2), and can also predict the oddity of (3)–(5). Conversational implicature does not adequately account for this data. In order to solidify my case for SV-Generous, I now argue that two of the more likely alternative hypotheses fail to explain particular parts of this data. The first alternative appeals to conventional implicature; the second appeals to presupposition.

II.4.a. Conventional Implicature

The relation of conventional implicature differs from conceptual entailment in that the former is detachable. More precisely, a sentence S1 carries I as a conventional implicature only if there could be another sentence S2 that is truth-conditionally equivalent to S1 but does not carry I as an implicature.[10] Nothing similar is true for conceptual entailment. Conventional implicatures are triggered by particular lexical items or linguistic constructions. For example, it's plausible that 'Smith hasn't arrived yet' conventionally implicates that Smith is expected to arrive and that this implicature is triggered specifically by the word 'yet'. Other conventional implicatures are said to be triggered by words like 'but', 'so', 'even', 'still', 'damn', and 'therefore', as well as appositives and parenthetical constructions.

As an alternative to SV-Generous, we may then wish to explain the connection between (1) and (2) by claiming that (2) is a conventional implicature triggered by the word 'generous' as it occurs in (1). Let's call this explanation CI-Generous. This explanation may also be able to account for the oddity of (5), since conventional implicatures cannot

10. R.M. Hare seems to believe that the evaluations associated with thick terms are detachable (Hare 1963, 188–89). And Blackburn's example of 'fat' and 'fat↓' appears to be a case where the former is the detached counterpart of the latter (1992, 290; 1998, 95).

be canceled without oddity. For example, it would seem odd in many contexts to say 'Smith hasn't arrived yet, and no one is expecting her'.

The main problem with CI-Generous is that it does not lead us to expect the awkwardness of (3):

(3) ? Nancy is generous, {and/but} she's good in a way.

It is distinctive of conceptual entailments that they are not reinforceable — they cannot normally be made explicit without awkward redundancy.[11] So, SV-Generous correctly predicts the oddity of (3). But linguists are highly reluctant to claim that conventional implicatures are not reinforceable.[12] And this is for good reason. Notice that many paradigmatic examples of conventional implicatures can be reinforced:

(9) a. Smith has not arrived **yet**, but he is expected.

b. **Even** Bill passed the test, and he was among the least likely.

c. Sophie is a baby, **but** she's quiet, and most babies are not quiet.

d. It's my turn to mow the **damn** lawn, and I hate mowing the lawn.

In each example, the boldfaced words trigger the implicatures stated explicitly in the rightmost clause.[13] But each of (9a–d) sounds normal. Thus, it appears that many conventional implicatures are reinforceable.[14] And this means that CI-Generous does not by itself predict the infelicity of (3).[15]

11. An exception to this is discussed below with regard to sentence (10).

12. Potts (2007, 668) claims that the implicatures triggered by some particular constructions are not reinforceable, but he shies away from the general claim. See footnote 15 for discussion on these constructions.

13. It is possible to quibble over whether the implicatures are stated exactly as they should be. But I doubt very much that the appropriate changes will make a difference to what I say here and below. For example, (9b) could be rephrased as 'Even Bill passed the test, and that was surprising/unexpected' without creating infelicity.

14. Proponents of CI-Generous might resist my objection by trying to explain away the felicity of (9a–d). In particular, they might claim that each boldfaced word is ambiguous in that it can carry different implicatures in different contexts. For instance, the first part of (9d) can, in some cases, implicate that the speaker hates *mowing the lawn*, but in other cases, it can implicate that the speaker hates *the lawn itself*. Thus, the proponent of CI-Generous may wish to explain the felicity of (9d) by insisting that its second clause is not completely redundant, because it clarifies the ambiguity of the first clause and thereby adds something new. However, this type of explanation fails to explain the felicity of (9d), because it incorrectly predicts that the following should also be felicitous:

? I went to the **bank**, and I went to a place that deals with money.

The boldfaced word in this sentence is ambiguous between riverbanks and financial institutions, but the second clause clarifies that the first clause is about financial institutions. The second clause therefore adds something new in the same way that was attributed to the second clause of (9d). But the example provided here still seems awkward. The felicity of (9a–d) therefore cannot be explained away as suggested.

15. It might be contended that, although some conventional implicatures are reinforceable, there are well-defined types of conventional implicatures that are not — and perhaps the infelicity of (3) can be modeled after those particular types. For example, it has been argued that parenthetical constructions, like the following, trigger non-reinforceable conventional implicatures — 'Lance Armstrong, **the cyclist**, battled cancer.' According to Christopher Potts (2007, 668), this sentence carries the implicature that Lance Armstrong is a cyclist. And this implicature is obviously not reinforceable, as shown in (i):

(i) ? Lance Armstrong, the cyclist, battled cancer. And he is a cyclist.

However, it is clear that the infelicity of (3) cannot be modeled after that of (i). As Potts (2007, 671) points out, sentence (i) retains its redundancy when that sentence is transposed. By transposing (i), we get

(ii) ? Lance Armstrong is a cyclist — he, the cyclist, battled cancer.

And (ii) seems just as redundant as (i). But, unlike that of (i), the infelicity of (3) does not survive transposition. By transposing (3), we get a perfectly normal sentence that displays no redundancy:

(iii) Nancy is good in a way — she's generous.

It might be thought that SV-Generous falls prey to the same objection that I've raised against CI-Generous. In particular, one might point out that there are examples of conceptual entailment that sound felicitous when reinforced. Consider an example adapted from Horn (1991):

(10) Bush won by a small margin, but win he did.

In this example, the first clause conceptually entails the second, but the whole sentence seems perfectly felicitous. Does this mean that SV-Generous falls prey to the same objection raised against CI-Generous — *i. e.*, that it does not lead us to expect the infelicity of (3)?

No. The difference is that there are established ways of accounting for the fact that sentences like (10) are reinforceable, without predicting that (3) should be reinforceable. But no similar account seems in the offing for (9a–d). For instance, Laurence Horn has argued that sentences like (10) are instances of "rhetorical opposition". In asserting the first clause, the speaker makes a concession (*e.g.*, she concedes Bush's relative lack of popularity). But in the second clause, the speaker affirms something "on the opposite side of the argumentative or emotive ledger" from what was conceded in the first clause (Horn 1991, 334). This concession/affirmation structure is signaled by the fact that a 'but' connective (or similar device[16]) is necessary for its felicity — replacing 'but' with 'and' makes (10) infelicitous. This concession/affirmation structure, according to Horn, explains why sentences like (10) are felicitous.[17] But it clearly does not predict that (3) should be similarly felicitous, since (3) lacks this concession/affirmation structure and involves no rhetorical opposition.

The second clause of (iii) is naturally heard as expanding upon the first clause. Thus, when (3) is transposed, its redundancy disappears. Therefore, the redundancy of (3) cannot be modeled after that of parenthetical constructions like (i).

16. Similar devices include 'nonetheless', 'just the same', 'be that as it may', or 'despite this'.

17. For a similar explanation, see Ward (1988, 191).

Does rhetorical opposition explain why the above examples of conventional implicature in (9a–d) are felicitous? At best, it would account only for the felicity of (9a), leaving (9b–d) untouched. The concession/affirmation structure required for rhetorical opposition is not present in (9b–d). This can be seen from the fact that there is no need for a 'but' connective between the first and second clauses of (9b–d).

In short, we've seen that certain paradigmatic examples of conventional implicatures are reinforceable. I do not claim that *all* are reinforceable — only that *some* are. And this is enough to establish that an appeal to conventional implicature does not by itself predict the infelicity of (3). Thus, unlike SV-Generous, it seems that CI-Generous cannot adequately account for all the data we've encountered.

II.4.b. Presupposition

Can we appeal to presupposition to explain this data? Presupposition can be understood in two general ways.[18] First, it may be understood in terms of a speaker taking a proposition for granted (i. e., assuming its truth) in making an utterance (Stalnaker 1970). For example, in uttering

(11) Smith regrets that he drank Pabst.

the speaker clearly takes the following for granted:

P: that Smith drank Pabst.

A second way of understanding presupposition is in terms of backgrounding a proposition. Consider the following:

(12) Smith, who drank Pabst, is feeling ill.

(12) clearly implies P in some sense, but it does not take P for granted — after all, (12) might convey P as completely new information. Still, there is a legitimate sense in which (12) seems to background P. The basic idea is that a sentence, such as (12), can convey a number

18. See Chierchia and McConnell-Ginet (1990, 281–83) for a general discussion.

of different propositions, some of which are the main message or point of the utterance (*i.e.*, the foreground entailments), while others are backgrounded (Levinson 1983). The main message of (12) is that Smith is feeling ill, whereas P is merely conveyed as background.

For present purposes, we can simply take the disjunction of these two views as a necessary condition for presupposition. More precisely, a speaker presupposes a proposition P in uttering a sentence only if she either backgrounds P or takes P for granted in making that utterance.

How does this help us explain the data we've encountered? An interpreter's acceptance of an utterance's presupposition can act as a precondition for the felicity of that utterance. If you believe that Smith never drank Pabst, you might find it odd for someone to utter either (11) or (12). But it is worth noting that presuppositions may occur in some contexts and disappear in others. For example, in (11), the factive verb 'regrets' triggers P as a presupposition, but that presupposition disappears in other contexts where that verb is used—*e.g.*, 'Smith does not regret drinking Pabst, because he never did!' But in most contexts, that verb will trigger the relevant presupposition. In sum, the presupposition-based explanation of our data would state that the thick term 'generous' triggers a presupposition in at least some contexts and that this allows us to explain the connection between (1) and (2) as well as the infelicity of (3)–(5).

I shall argue that the infelicity of (5) cannot be explained in this way. But, in order to assess this potential explanation, we must be clear on what presuppositions might conceivably be triggered by the word 'generous'. In particular, what is the putative presupposition of (1)?

For starters, we cannot plausibly claim that (1) directly presupposes (2). This is because the connection between (1) and (2) lacks one of the distinctive features of presupposition—"constancy under negation" (Huang 2007, 67). If the utterance of an affirmative sentence S presupposes a proposition P, then we typically expect that a simple negation of S will also presuppose P. For example, the negations of (11) and (12)—appearing respectively as (13) and (14) below—both seem to presuppose that Smith drank Pabst:

(13) Smith does not regret drinking Pabst.

(14) Smith, who drank Pabst, is not feeling ill.

But nothing similar is true for (1). It is clear that the negation of (1) does not presuppose (2). More precisely, the following statement,

(15) Nancy is not generous

clearly does not take for granted or background the proposition that Nancy is good in a way. And this seems true for any context of utterance.

Alternatively, it might be claimed that the relevant presupposition of (1), as well as (15), is one of the following:

> Q: that people who are liberal in giving and sharing are good in a particular way.
>
> R: that if Nancy is liberal in giving and sharing, then she's good in a particular way.[19]

And since the combination of (1) with either Q or R would plausibly entail (2), this might help us explain the connection between (1) and (2).

I doubt very much that there are any contexts in which Q or R are triggered as presuppositions by the word 'generous'.[20] But let's suppose, for the sake of argument, that this does occur in at least some contexts. For example, let's suppose that when I utter negations like 'Nancy is not generous—she's greedy and selfish', I have thereby presupposed one of Q or R. Even so, this assumption does not help us explain the oddity

19. The suggestion that (1) may presuppose either Q or R is a possibility that I have loosely modeled after some claims that Väyrynen (2012) makes with regard to 'lewd'. He does not explicitly accept a presupposition account, but this view is implicit in his discussion.

20. The attempted justification for this claim would most likely appeal to the behavior of objectionable thick concepts in negations, modals, disjunctions, and conditionals (see Väyrynen 2012). But, as I argue in section III, we already have widely accepted pragmatic mechanisms for explaining this behavior (*e.g.*, negative strengthening and clausal implicature) without postulating that Q or R are presuppositions.

of (5). The reason is that there are many normal contexts in which neither Q nor R are presupposed by uses of 'generous', but in which (5) would still sound odd. To see this, let's focus on one particular context in which a speaker utters the negation of (1):

> *Bob*: "Nancy is highly controlled in her giving and sharing, and that's what makes her fiscally smart. She is not generous. But she's not selfish either. I admire her approach to finances."

It seems clear that Bob's utterance of the negation of (1) does not imply, let alone presuppose, either Q or R. In particular, there is no sense in which Bob has taken Q or R for granted, nor is there any sense in which he has conveyed Q or R as background. Thus, it seems unlikely that any of these propositions are presupposed by Bob's use of 'generous' in this particular context. The trouble, however, is that the context just described is one in which (5) would sound odd. For example, imagine that Bob's interlocutor, Sue, responds to him as follows:

> *Sue*: "I disagree with you, Bob. Nancy is generous, and she's not good in any way."

Here, Sue's utterance of (5) is clearly infelicitous. And this means that the oddity of (5) outstretches contexts in which propositions like Q or R are presupposed (if there are any). After all, the context just described is one where those presuppositions are absent while the oddity of (5) remains. Thus, it's hard to see how an appeal to presupposition could explain the infelicity of sentences like (5).

II.5. Generalizing the Argument

It looks quite plausible that SV-Generous provides the best explanation of the data we've seen. I've not surveyed all possible explanations, but I've shown that the main contenders are unable to explain the infelicity of one or another of (3)–(5). Thus, I think there is good reason to believe that SV-Generous is true. However, SV-Generous is a very particular claim about a single thick concept (*i. e.*, *generous*), while the Semantic View is a general claim that's supposed to be true of many thick concepts. How do we get from SV-Generous to the more general Semantic View?

For the sake of simplicity, the above argument has focused only on data involving the thick term 'generous'. But there are parallel data involving other thick terms. For example, in the above argument, it was noted that (3) seems awkwardly redundant. But the same goes for other positive thick terms, like 'kind' and 'courageous':

(16) ? Jones is kind, {and/but} he's good in a way.

(17) ? Smith is courageous, {and/but} he's good in a way.

Moreover, negative thick terms, like 'murder' and 'brutal', exhibit similar behavior with respect to a different evaluative concept — *bad in a way*:

(18) ? That was a murder, {and/but} it was bad in a way.

(19) ? That is brutal, {and/but} it's bad in a way.

(16)–(19) seem at least somewhat odd. And this is just what we would expect if the Semantic View were true of 'murder', 'brutal', 'kind', and 'courageous'. On this view, (16)–(19) seem odd because they are redundant: their first conjuncts conceptually entail what their second conjuncts assert.

Another important datum used in the above argument is the fact that (5) is highly odd, even though each of its conjuncts is perfectly normal. But similar claims seem true regarding 'kind', 'courageous', 'murder', and 'brutal':

(20) # Jones is kind, and he's not good in any way.

(21) # Smith is courageous, and she's not good in any way.

(22) # That was a murder, and it was not bad in any way.

(23) # That is brutal, and it's not bad in any way.

(20)–(23) seem highly awkward, even though each of their conjuncts is perfectly normal. And this too is just what we would expect if the Semantic View were true of these thick terms. According to this view, (20)–(23) sound highly odd because they are contradictory: their second conjuncts contradict what their first conjuncts conceptually entail.

This brings up two important points. The first is that it looks like the argument advanced regarding 'generous' will apply with equal efficacy to these other thick terms — 'murder', 'brutal', 'kind', and 'courageous'. The rival explanations discussed earlier would be inferior to the relevant instance of the Semantic View for the same general reasons that were brought up with regard to 'generous'. So, we have strong reason to believe the Semantic View is also true of these thick terms.

The second point is that the five thick terms discussed herein are a representative sample of the thick terms that are either purely positive or purely negative (*i.e.*, not multivalent). We therefore have good reason to think that many other thick terms are likely to exhibit similar behavior. In particular, it's likely that many other positive thick terms (*e.g.*, 'loyal', 'chaste', 'wise', 'honest', etc.) are also semantically associated with *good in a way* and that many negative thick terms (*e.g.*, 'lewd', 'rude', 'liar', 'dishonest', 'unkind', etc.) are semantically associated with *bad in a way*. The number of thick terms to which the Semantic View applies seems to expand out quite significantly. I conclude, then, that we have good reason to accept the Semantic View — many thick concepts, if not all, conceptually entail evaluative concepts.

The Semantic View, as I've stated it, extends to many thick terms, although I've left open the possibility that it might not extend to all. This possibility is left open primarily because the style of argument I have advanced in support of the Semantic View may not apply with respect to at least two kinds of terms.

First, consider slurs (*e.g.*, 'kraut', 'gringo', 'tart', etc.). There is no consensus about whether slurs should count as thick. Some ethicists, such as R.M. Hare (1963, 25) and Mark Richard (2008, 14), have no qualms about including them among the thick. Others, like Jonathan Dancy (1995, 264) and Allan Gibbard (2003, 300ff), are hesitant. Rather than taking a stand on this issue, let me instead assert the following conditional: If slurs are thick terms, then we've seen no automatic reason to extend the Semantic View to cover them. In particular, the argument I have advanced in favor of the Semantic View relies on data like (18)–(19) and (22)–(23). But it's doubtful that similar data will be available for all slurs. For example, the sentence 'Dr. K is a gringo, and he's not bad in any way' strikes me as infelicitous, but its infelicity may be wholly attributable to the first conjunct, 'Dr. K is a gringo.' Thus, whether or not we have reason to extend the Semantic View to all thick terms may depend on whether we include slurs among the thick. I do not deny that the Semantic View is true of slurs — I merely claim that my argument does not clearly establish this.

And second, consider multivalent thick terms (*e.g.*, 'unorthodox', 'eccentric', 'quirky', 'kinky', and 'grotesque').[21] Roughly speaking, these are thick terms that are commonly used to evaluate both positively and negatively, perhaps in different contexts.[22] Do these terms present a problem for the argument I have advanced? The first thing to note is that the phenomenon of multivalence is perfectly compatible with the view that all thick concepts conceptually entail either good in a way or bad in a way. In general, a concept can be used to evaluate, say, negatively, even if it conceptually entails good in a way — e.g., 'Jones is a good thief and therefore not to be trusted.' We could even suppose, for the sake of argument, that both the positive and negative evaluations of a multivalent thick term are conceptually encoded (i.e., part of its meaning). Even this supposition is perfectly compatible with the view that all thick concepts conceptually entail either good in a way or bad in a way, because it is possible that each multivalent thick

21. Thanks to an anonymous referee for reminding me of multivalence and pointing out how my argument may not extend to such terms.

22. This rough account of multivalence might be so broad as to include too many thick terms as multivalent, but nothing of substance in this paper will rely on the account. For a sustained discussion on multivalence, see Väyrynen (2011). As he notes, the positive and negative use of these terms is meant to be restricted to literal usage (2011, 4).

concept conceptually entails both good in a way and bad in a way.[23] This kind of multivalence is not unprecedented: if liar is a negatively evaluative concept, then good liar would likely fit this mold of having both positive and negative evaluations conceptually encoded.

It is one thing to show that the Semantic View is consistent with multivalence, but it's another to show that my argument for the Semantic View can be applied to multivalent thick terms. And unfortunately, the argument may not apply, because the analogous linguistic data may not be available with regard to some multivalent thick terms. For instance, it seems felicitous to utter 'Max is unorthodox, and he's not bad in any way', and the same is true if we replace 'bad' with 'good'. If this is right, then the type of argument I have advanced with regard to 'generous' could not be applied to 'unorthodox'. This situation may not exist for all multivalent thick terms, but it's likely to be more prevalent among them because our usage of such terms may not typically imply any one of our two main evaluative concepts — good in way and bad in a way. Thus, it's possible that the type of argument I have advanced would not apply to some multivalent thick terms, because the type of linguistic data I have relied on is not available. However, just as with slurs, I do not deny that the Semantic View is true of multivalent thick terms — I only claim that my argument does not clearly demonstrate this for some such terms.

In short, slurs and multivalent thick terms are two potential exceptions to the Semantic View. There may be others. The notion of a thick term is not well-defined (as illustrated by the disagreement about whether slurs are thick). And so it isn't surprising that we find importantly different classes of terms grouped together as thick. Nonetheless, I have argued that the Semantic View is plausible with

23. One attempt at explaining the phenomenon of multivalence is to hold that this view is not merely possible but true. This would not explain all types of multivalence, however, since it still needs to be explained how the use of some multivalent thick terms can lose a valence in a given context. For example, some uses of the term 'kinky' are nothing but negative and carry no positive valence whatsoever. What then happens to the positive valence in this context? This is different from the behavior of 'good liar', which appears to carry both evaluations in all contexts.

regard to a sizeable group of thick terms, including many of our virtue and vice terms (e.g., 'courageous') as well as other important moral expressions (e.g., 'murder'). Now that I have presented my case for the Semantic View, we can consider how well this view stands up against opposition. In the next section, I defend the Semantic View against the most recent critique by Pekka Väyrynen stemming from objectionable thick concepts.

III. Objectionable Thick Concepts.

Väyrynen (2009) has argued that the use of objectionable thick terms within certain contexts sheds light on the question of how thick terms are associated with evaluations. And he thinks this data strongly suggests that these evaluations are not semantically associated with thick terms. But what are objectionable thick concepts? And how might they be a problem for the Semantic View?

Very roughly, objectionable thick concepts are concepts that embody values that ought to be rejected. The concept *lewd*, for example, seems to embody the view that overt sexual behavior is somehow bad. But many people believe this evaluative perspective ought to be rejected. So, it appears that *lewd* is a candidate for being an objectionable thick concept. In effect, the question of whether a given thick concept actually is objectionable depends on the potentially controversial question of whether its associated values ought to be rejected. So, there's much room for debate about which thick concepts are actually objectionable. Concepts like *lewd*, *chaste*, *blasphemous*, and *sexually perverse* are commonly seen as paradigmatic examples, although some might dispute whether these are really objectionable. Nonetheless, speakers who in fact reject the values embodied by a given thick concept (whether or not they should) tend to exhibit some interesting linguistic behavior. For the sake of simplicity, I'll refer to these speakers as objectors.

Objectors are often reluctant to use the thick terms they regard as objectionable.[24] Of course, we may expect that someone who rejects the values embodied by chaste (i. e., a chastity-objector) would be highly reluctant to utter an affirmative sentence of the following form:

(24) A is chaste.

This type of sentence clearly endorses the kind of values rejected by the chastity-objector. What is surprising, however, is that chastity-objectors are also reluctant to utter the negation of (24) — namely

(25) A is not chaste.

Their unwillingness to assert (25) is initially puzzling, since it's plausible that chastity-objectors should take (25) to be true. It looks as if these negations also endorse the kind of values rejected by the chastity-objector.

Väyrynen thinks the reluctance of chastity-objectors to assert (25) presents a problem for the Semantic View. In particular, he advances the following claim, which I shall call VC:

> *VC:* If those who reject the evaluative content associated with *chaste* are not willing to apply truth-conditional negation to (24), then, ceteris paribus, that evaluative content is not part of the truth conditions of (24).[25]

24. It is worth pointing out that some people will be reluctant to employ certain thick terms even if they wholly accept the associated values. For example, many people who believe that premarital sex is wrong would nonetheless want to avoid using the word 'fornicate'. Something similar can be said for a host of thick terms like 'sacrilegious', 'holy', 'pure', 'defile', 'sinful', 'infidel', 'profane', 'heretical', and 'heathen'. But the unwillingness of objectors to use these words is of questionable significance, since many people who wholly accept the relevant values are also reluctant to use them. The worry is that certain paradigmatically objectionable thick terms, like 'chaste', 'blasphemous', and 'perverse', may belong in this category. If they do, then their significance to this discussion would be questionable. For the sake of argument, however, I will assume that an objector's reluctance is different in kind from that exhibited by wholehearted believers who wish to avoid 'fornicate', 'defile', etc.

25. I add the ceteris-paribus clause so as to weaken VC in a way that approximates Väyrynen's own statement. He says that, if the antecedent of VC is true, then the evaluative content "doesn't seem to be" part of the truth-conditions of 'A is chaste' (2009, 448).

The basic thought behind Väyrynen's claim is this: If what the chastity-objector rejects is located within the truth-conditions of (24), then truth-conditional negation of (24) should be a perfectly acceptable way of expressing her disagreement (Väyrynen 2009, 448). This basic thought is initially attractive, and it has been accepted by others as well. Matti Eklund (2011, 34) makes a similar point about objectionable thick terms:

> …[I]f what was objectionable about the value words was a matter of truth-conditional content, then one could well use the words even if one finds them objectionable: it is just that one would not assert of anything that one of these words applies to it.

Eklund's claim here is slightly more general than VC, but the same idea lies behind them both. The Semantic View seems to predict that objectors should be willing to use objectionable thick terms in non-affirmative sentences, like (25).

If VC were correct, it would spell serious trouble for the Semantic View. As noted, chastity-objectors are not typically willing to apply truth-conditional negation to (24). In other words, they are not willing to assert (25). But this datum, in combination with VC, entails that (ceteris paribus) the evaluative content is not part of the truth-conditions of (24).[26] Therefore, it doesn't seem that the Semantic View will be true of 'chaste'. Moreover, Väyrynen notes that, since any thick concept can, in principle, be regarded as objectionable, his argument will apply to other thick concepts as well (2009, 449). So, Väyrynen's argument threatens to pose a more general problem for the Semantic View.

26. Technically speaking, this argument does not show that there is *no* evaluative content within the truth conditions of 'chaste', but only that the one rejected by the speaker is not part of those conditions. But for all that matters, there might be multiple evaluative contents associated with 'chaste'. I shall ignore this possibility in what follows.

But I think VC is clearly mistaken. It is important to notice that any person who rejects the evaluative content associated with chaste will also reject at least one other evaluative content — namely the evaluative content associated with unchaste.[27] But Väyrynen overlooks the possibility that a speaker's reluctance to assert (25) might have something to do with her reluctance to be saddled with a claim about A's being unchaste. In particular, notice that

(25) A is not chaste.

clearly seems to imply

(26) A is unchaste.

This kind of implication is what linguists call "negative strengthening" (Levinson 2000, 127). Oftentimes, when 'not' is combined with certain words, such as 'happy', 'believe', 'like', 'good', and 'bad', the speaker is committed to something stronger than what she literally said. For example, if I utter, 'Smith is not happy', this utterance will likely convey the stronger claim that Smith is unhappy. Linguists typically say that claims like (25) and (26) are not truth-conditionally equivalent, but that (25) implies (26) by virtue of conversational implicature (Horn 1989, 331ff; Levinson 2000, 127ff). Let's assume this is true for present purposes.[28] The relationship between (25) and (26) appears to exhibit

27. Since *chaste* and *unchaste* embody the same general evaluative perspective, any chastity-objector who does not also reject *unchaste* would seem to be holding an unintelligible position. Her tendencies about asserting claims like (24) and (25) should therefore not be taken seriously as data for this discussion.

28. A more controversial explanation would rely on the claim that 'not' is ambiguous. Following John Lyons, we might say that, on some occurrences, 'not' "converts a proposition into its contradictory…" (1977, 772). On this reading, (25) is equivalent to the following:

(25*) It's not the case that Smith is chaste.

However, on other occurrences, 'not' converts the proposition into its "contrary". Since the contrary of 'chaste' is 'unchaste', this explains the connection between (25) and (26). On this "contrary" reading of 'not', (25) is actually equivalent to (26). If this sort of view is correct, then we can explain the speaker's reluctance to utter (25) by citing the fact that this utterance is ambiguous

negative strengthening. That is, (25) implicates (26) in the same way that 'Smith is not happy' implicates 'Smith is unhappy.' If these claims are correct, then a chastity-objector would clearly not want to assert (25), because that assertion would implicate (26). And (26) also embodies a value that she rejects.

Given the phenomenon of negative strengthening, we can make sense of the fact that 'A is not chaste' appears also to endorse the values that chastity-objectors reject. Very roughly, when negative strengthening occurs, the denial of a value claim implicates a value on the opposite end of the same evaluative scale. But the objector in question is not someone who merely rejects one side of that evaluative scale — she rejects the entire scale! And so, this speaker will be reluctant to utter even the denial of the original value claim.

The story just told provides reason to reject VC, because it explains the objector's reluctance in a way that is perfectly consistent with the claim that the evaluative content (rejected by the chastity-objector) is part of the truth conditions of (24).[29] VC must therefore be rejected.

between something she most likely accepts — (25*) — and something she outright rejects — namely, (26). Moreover, we can note that (26) is actually the more common reading of (25), which suggests that someone who finds *chaste* objectionable would withhold uttering (25) without further clarifications.

29. To see this, consider an analogous story involving 'right'. Suppose you're a Mackie-style error-theorist about rightness. And by the same token, you also reject the existence of wrongness. Due to your skepticism about rightness, you would be reluctant to utter things like

(A) Going to war is right.

But you would also be reluctant to apply truth-conditional negation to (A) by asserting

(B) Going to war is not right.

After all, typical utterances of this sentence strongly imply

(C) Going to war is wrong.

by virtue of negative strengthening. But your reluctance to apply truth-conditional negation to (A) can be given the same kind of explanation that I've provided with regard to (24). And this account in no way impugns the fairly

It should be noted that we could change Väyrynen's example so that we focus on a different thick concept aside from chaste. But this will not reinstate the type of claim he wants to make. To be sure, with regard to some thick concepts, it is less obvious what implicature gets generated through negative strengthening. For example, since there's no such word as 'unblasphemous', it's not wholly clear what would be implicated by 'A is not blasphemous'. However, the phenomenon of negative strengthening is not limited to expressions that can be prefixed with 'un'. For example, 'good' cannot be combined with that prefix, but an utterance of 'A is not good' tends to implicate that A is bad. And it's plausible that something similar is true for 'blasphemous'. An utterance of

(27) The story is not blasphemous.

seems to imply something like

(28) The story is somewhat reverent.[30]

common view that (A) has an evaluative truth-condition. By analogy, the story I've told about *chaste* is also consistent with claiming that (24) has an evaluative truth-condition.

30. The adverb 'somewhat' is important, due to an asymmetry in the way positive and negative adjectives exhibit negative strengthening. Linguists typically acknowledge that, although 'A is not happy' clearly implicates 'A is unhappy', the analogous implicature is not generated by 'A is not unhappy'. That is,

(i) A is not unhappy.

does not straightforwardly implicate

(ii) A is happy.

However, it is often claimed that (i) implicates something like

(iii) A is somewhat happy (though not quite as happy as the word 'happy' would suggest).

Levinson (2000, 145) and Blutner (2004, 500–1) hold this type of view. Analogous claims hold for negative thick terms like 'blasphemous' and 'unchaste'. The central point here is that, even though there is an asymmetry as described, the objectors will still be opposed to weaker claims like (28).

And someone who thinks *blasphemous* is objectionable may not want to utter (27), for fear of being saddled with something like (28). Perhaps this explains Oscar Wilde's reluctance to utter (27) when faced with the attorney's aggressive questioning: "Did you or did you not consider the story blasphemous?" (Wilde and Carson 1895).[31]

Of course, we may not always have an antonym, like 'reverent', corresponding to each potentially objectionable thick term. But this is merely a limitation of our language, not of my appeal to negative strengthening. Some languages are even more limited. As Joseph Greenberg observes, certain African, Amerind, and Oceanic languages have no word for 'bad'. Nonetheless, speakers in these languages can convey that something is bad by negating their term for 'good' (1966, 52). Their way of expressing that A is not good implicates that A is bad, even though they have no word for 'bad'. So, negative strengthening can occur even if we don't have the appropriate antonym to express the relevant implicature.

An objection may arise regarding whether my appeal to negative strengthening over-generalizes. Väyrynen correctly notes that there are some contexts in which chastity-objectors would be willing to assert 'A is not chaste'. For example,

(29) Smith is not chaste, but neither is he unchaste.

(30) Smith is not chaste; the mere fact that he's dedicated to not being sexually provocative does not make him good in any way.

Does my appeal to negative strengthening incorrectly predict that the chastity-objector would be reluctant to assert (29) and (30)? No. We can understand the follow-up clauses in (29) and (30) as elements that cancel the conversational implicature from 'Smith is not chaste' to 'Smith is unchaste.' The cancelation is obvious in (29), since that implicature is explicitly denied by the second part of (29). In (30), the implicature is not explicitly denied, but it's plausible that the follow-up

31. For another possible explanation of Wilde's reluctance, see footnote 24.

clause provides enough reason to doubt that the speaker intends to convey that Smith is unchaste. So, the problematic implicature is not generated by (30). Thus, my appeal to negative strengthening allows for the acceptability of (29) and (30) to chastity-objectors.[32]

Thus, it looks quite plausible that a speaker's reluctance to utter sentences like 'A is not chaste' in typical contexts is explainable through negative strengthening. And if that's so, then VC must be rejected, and the argument against the Semantic View fails.

Or is there an easy way in which Väyrynen can modify his argument? That argument focuses solely on the unwillingness of objectors to utter negations. But he also claims that "[c]onditionals exhibit the same phenomenon". In particular, he thinks that objectors are typically unwilling to utter indicative conditionals whose antecedents contain objectionable thick terms. Väyrynen provides the following example (2009, 448):

(31) If abstinence from extramarital sex is chaste, then so is refraining from desiring extramarital sex.

According to Väyrynen, a chastity-objector's reluctance to assert (31) would be initially puzzling, since these people should take (31) to be true due to a false antecedent. So, perhaps Väyrynen's argument can be refocused on conditionals instead of negation.

In reply, let me first point out that the data involving conditionals is much less secure than that of negation. By this, I mean that (i) there are significantly fewer contexts in which chastity-objectors would exhibit any reluctance at all, and (ii) their reluctance in those contexts would be much weaker. To illustrate (i), we can easily imagine a chastity-objector playing devil's advocate with an overly pious interlocutor. In this case, the objector might feel perfectly comfortable uttering (31), and she wouldn't need to utter any follow-up clauses to qualify her statement (*e.g.*, '…not that I believe in chastity'). Regarding (ii), we can imagine a chastity-objector who foregoes the opportunity to utter (31) and instead replies, 'I wouldn't put it that way, but I guess that seems plausible.' This type of response illustrates a sort of reluctance that is much weaker than what we would expect with regard to (25). For these reasons, I think it is clear that the data involving conditionals is less secure than that of negation.

Nevertheless, in cases where objectors are reluctant to utter conditionals like (31), how can their reluctance be explained by the Semantic View? Since these conditionals do not involve negation, we obviously cannot appeal to negative strengthening. Nonetheless, a structurally similar explanation is available. In particular, an utterance of (31) in many contexts seems to imply

(32) Abstinence from extramarital sex may be chaste, or it may not.

According to many linguists (*e.g.*, Gazdar 1979, 59–62; Levinson 1983, 137; 2000, 108–9), the relationship between sentences like (31) and (32) is a type of conversational implicature known as "clausal implicature". By uttering the conditional 'if p, then q' (rather than the stronger alternative 'since p, q'), the speaker conveys epistemic uncertainty about whether the antecedent is true. The conditional clausally implicates that the antecedent 'p' may or may not be true. For instance, according to Stephen Levinson, an utterance of 'If there is life on Mars,

32. Väyrynen tries to explain the fact that chastity-objectors find sentences like (30) to be acceptable by claiming that these sentences are instances of metalinguistic negation (2009, 449). See Horn (1989, 377) for an account of metalinguistic negation. However, if (30) can be seen as a case in which negative strengthening is canceled, then I see no reason to postulate that (30) is metalinguistic. Moreover, it's worth pointing out that (30) fails Horn's incorporation test for metalinguistic negation (1989, 392ff). The negation in 'The king of France is not happy, because there is no king of France' cannot be incorporated. That is, when 'not happy' is replaced with 'unhappy', the result is unintelligible: # 'The king of France is unhappy, because there is no king of France.' But nothing similar is true for (30). If we replace 'not chaste' in (30) with 'unchaste', the result is not the least bit unintelligible:

(30′) Smith is unchaste; the mere fact that he's dedicated to not being sexually provocative does not make him good in any way.

So, this disparity is prima facie evidence against taking (30) to be metalinguistic. Of course, Horn's incorporation test is not uncontroversial. Geurts (1998, 280) is one critic.

the NASA budget will be spared' clausally implicates 'There may or may not be life on Mars' (2000, 36). Assuming this is correct, we are once again in a position to explain the chastity-objector's reluctance by way of conversational implicature. Chastity-objectors are likely to take issue with (32) in certain contexts, since they fail to believe that abstinence may be chaste. And since (31) conversationally implicates (32), they would be reluctant to utter (31) in those contexts.[33]

Thus, Väyrynen's strategy would fare no better if he focused on conditionals instead of negation. I should also note that the type of explanation just provided (vis-à-vis clausal implicature) can also be applied to disjunctive statements (*e.g.*, 'Either Smith is chaste, or he's keeping secrets'), belief reports (*e.g.*, 'The Pope believes Smith is chaste'), as well as modal statements (*e.g.*, 'It's possible that Smith is not chaste') (Levinson 1983, 136–7; 2000, 108–11). For reasons similar to those mentioned earlier, I believe the data with regard to these statements is much less secure. But, in contexts where there is reluctance, clausal implicature is a perfectly viable explanation. Disjunctions, modals, and belief reports also implicate that the speaker is uncertain (*e.g.*, about whether Smith is chaste), and this is something the chastity-objector would want to avoid in certain contexts.

So far, I have argued that VC is false and that it cannot be salvaged through appropriate modification. Proponents of the Semantic View can appeal to pragmatic mechanisms — like negative strengthening and clausal implicature — to explain an objector's reluctance to utter certain sentences involving objectionable thick terms. However, even if VC is mistaken, it might be challenged that the proponent of the Semantic View is here appealing to an explanation that is inferior to rival explanations because it seems less unified and less simple.[34] Space prohibits a comparison of explanations with regard to simplicity and unity, but let me briefly say why the proponent of the Semantic View does not falter in these respects.

Despite initial appearances, the proposed way of explaining the reluctance of objectors has a unified base. Both clausal implicature and negative strengthening are forms of conversational implicature. These inferences are therefore based on the addressee's assumption that the speaker is following certain principles of conversation, such as Grice's Cooperative Principle (Grice 1989, 26). So, the unified base underlying the proposed explanation is this: Objectors are reluctant to utter sentences like (25) and (31), as well as certain modals, disjunctions, and belief reports, because such utterances would conversationally implicate claims that objectors want to avoid. There is no clear sense in which this explanation lacks unity, and so I think the charge is unwarranted.

It is also misguided to claim that the proposed explanation lacks simplicity. In appealing to negative strengthening and clausal implicature, the proponent of the Semantic View is not postulating anything that rival views would not already postulate for more general reasons.[35] For example, it is widely accepted that 'not happy' conversationally implicates 'unhappy'. And, once this view is granted, it is extremely hard to deny that this same relation also holds between 'not chaste' and 'unchaste'. So, even those who reject the Semantic View would likely accept the particular pragmatic relations appealed to by the proponent of the Semantic View. Therefore, the proposed explanation involves no further postulates, and is no less simple, than these rival views.[36]

33. Of course, the contexts in which objectors are reluctant to utter (31) would need to "line up" with those in which they would take issue with (32), but I see no immediate reason to think that won't be the case.

34. Väyrynen (2012) briefly advances these two charges against my view.

35. Of course, the proponent of the Semantic View is postulating evaluative meanings, which rival views do not postulate. But, as I argued in section II, these are not postulated beyond necessity, because they are needed for explaining the infelicity of sentences like (3)–(5). Moreover, rival views, such as that of Väyrynen (2012), tend to postulate that evaluations project, or are presupposed, which is not something that the Semantic View needs to postulate. Thus, the two views initially appear to be on par with regard to simplicity. But see footnote 36.

36. In fact, it now looks like the charge of lacking simplicity can be turned against rival views, such as Väyrynen's appeal to projection (see his 2012). Why

IV. Conclusion.

This paper has run the gamut of possible views on how thick terms might be associated with evaluations. I have argued for a Semantic View, according to which many (if not all) thick concepts conceptually entail evaluative contents.[37] In section II, it was argued that this view best explains certain data involving thick terms and expressions like 'good in a way' and 'bad in a way'. A number of rival hypotheses were shown unable to explain this data.

The Semantic View, however, has a number of detractors, and I addressed the most recent of them in section III. It was argued that the considerations raised by objectionable thick concepts do not supply a compelling case against the Semantic View.

If I am correct in holding that thick terms bear a semantic relationship to evaluations, then this settles a dispute that is central to a broader debate in ethics. In particular, a number of ethicists have dismissed the importance of thick concepts within ethics, and they have done so by claiming that thick terms are not semantically associated with evaluations. But if my argument in this paper is correct, then this claim is mistaken and therefore does not permit ethicists to overlook the thick. It is highly plausible that thick terms are associated with evaluations in a way similar to how thin terms are commonly thought to be associated with evaluations. Thick terms are semantically evaluative.[38]

should we seek additional resources, like projection — on top of what's already available with negative strengthening and clausal implicature — for explaining the linguistic behavior of objectors? The answer is by no means obvious.

37. Recall that footnote 1 provides a modified version of the Semantic View that does not assume a distinction between evaluative and non-evaluative content. The modified view holds that thick concepts conceptually entail the contents expressed by thin terms, and this modified view is equally supported by the arguments in this paper.

38. An early version of this paper was distributed within a dissertation writing group at Cornell University; many thanks to its participants for their feedback — Scott O'Connor, Kristen Inglis, Colin McLear, and Andrew Alwood. My gratitude especially goes to Matti Eklund for his advice, comments, and willingness to discuss many different versions of this paper. I am also indebted to Nicholas Silins, Carl Ginet, Kent Dunnington, Nicholas Sturgeon, and the referees for *Philosophers' Imprint*.

References

Anscombe, G.E.M. 1958. 'Modern Moral Philosophy'. *Philosophy* 33: 1–19.

Atlas, J.D. and S.C. Levinson. 1981. 'It-Clefts, Informativeness, and Logical Form: Radical Pragmatics (Revised Standard Version)'. In *Radical Pragmatics*, Peter Cole, ed. New York: Academic Press: 1–61.

Bach, K. 1999. 'The Myth of Conventional Implicature'. *Linguistics and Philosophy* 22: 327–66.

Blackburn, S. 1992. 'Through Thick and Thin'. *Proceedings of the Aristotelian Society*, suppl. vol. 66: 285–99.

——. 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.

Blutner, R. 2004. 'Pragmatics and the Lexicon'. In *The Handbook of Pragmatics*, L.R. Horn and G. Ward, eds. Oxford: Blackwell: 488–514.

Brower, B.W. 1988. 'Virtue Concepts and Ethical Realism'. *The Journal of Philosophy* 85: 675–93.

Chierchia, G. and S. McConnell-Ginet. 1990. *Meaning and Grammar: An Introduction to Semantics*. Cambridge, Mass: MIT Press.

Dancy, J. 1995. 'In Defense of Thick Concepts'. *Midwest Studies in Philosophy* 20: 263–79.

Eklund, M. 2011. 'What Are Thick Concepts?' *The Canadian Journal of Philosophy* 41: 25–50.

Foot, P. 1958. 'Moral Arguments.' *Mind* 67: 502–13.

Gazdar, G. 1979. *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.

Geurts, B. 1998. 'The Mechanisms of Denial'. *Language* 74: 274–307.

Gibbard, A. 2003. 'Reasons Thin and Thick'. *The Journal of Philosophy* 100: 288–304.

Greenberg, J.H. 1966. *Language Universals*. The Hague: Mouton and Co.

Grice, P. 1989. *Studies in the Way of Words*. Cambridge, Mass: Harvard University Press.

Hare, R.M. 1963. *Freedom and Reason*. Oxford: Oxford University Press.

Horn, L.R. 1989. *A Natural History of Negation*. Chicago: University of Chicago Press.

——. 1991. 'Given as New: When Redundant Affirmation Isn't'. *Journal of Pragmatics* 15: 313–36.

Huang, Y. 2007. *Pragmatics*. Oxford: Oxford University Press.

Hursthouse, R. 1995. 'Normative Virtue Ethics'. In *How Should One Live?: Essays on the Virtues*, R. Crisp, ed. Oxford: Oxford University Press: 19–36.

Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.

Levinson, S.C. 1983. *Pragmatics*. Cambridge, UK: Cambridge University Press.

——. 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, Mass: MIT Press.

Lovibond, S. 1983. *Realism and Imagination in Ethics.* Minneapolis: University of Minnesota Press.

Lyons, J. 1977. *Semantics: Volume 1*. Cambridge, UK: Cambridge University Press.

McDowell, J. 1981. 'Non-Cognitivism and Rule-Following'. In *Wittgenstein: To Follow a Rule*, S. Holtzman and C. Leich, eds. London: Routledge.

Millgram, E. 1995. 'Inhaltsreiche ethische Begriffe und die Unterscheidung zwischen Tatsachen und Werten'. In *Zum moralischen Denken* vol. 1, C. Fehige and G. Meggle, eds. Frankfurt a. M.: Suhrkamp: 354–88.

Murdoch, I. 1971. *The Sovereignty of Good*. New York: Schocken.

Platts, M. 1979. *Ways of Meaning: An Introduction to a Philosophy of Language*. Cambridge, Mass: MIT Press.

Potts, C. 2007. 'Into the Conventional-Implicature Dimension'. *Philosophy Compass* 2: 665–79.

Richard, M. 2008. *When Truth Gives Out*. Oxford: Oxford University Press.

Sadock, J. 1978. 'On Testing for Conversational Implicature'. In *Syntax and Semantics, vol. 9: Pragmatics,* P. Cole, ed. New York: Academic Press: 281–97.

Stalnaker, R. 1970. 'Pragmatics'. *Synthese* 22: 272–89.

Stanley, J. and Z.G. Szabó. 2000. 'On Quantifier Domain Restriction'. *Mind & Language* 15: 219–61.

Thomson, J.J. 2008 . *Normativity*. Chicago: Open Court.

Väyrynen, P. 2009. 'Objectionable Thick Concepts in Denials'. *Philosophical Perspectives: Ethics* 23: 439–69.

——. 2011. 'Thick Concepts and Variability'. *Philosophers' Imprint* 11: 1–17.

——. 2012. 'Thick Concepts: Where's Evaluation?' Forthcoming in *Oxford Studies in Metaethics* 7, R. Shafer-Landau, ed. Oxford: Oxford University Press.

Ward, G.L. 1988. *The Semantics and Pragmatics of Preposing*. New York: Garland Publishing.

Wilde, O. and E. Carson. 1895. 'Testimony of Oscar Wilde on Cross Examination'. http://www.law.umkc.edu/faculty/projects/ftrials/wilde/wildelibeltranowcross.html. Accessed June 8, 2010.

Williams, B. 1985. *Ethics and the Limits of Philosophy*. Cambridge, Mass: Harvard University Press.

——. 1995. 'Truth in Ethics'. *Ratio* 8: 227–42.

Williamson, T. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.

Wilson, D. 1975. *Presuppositions and Non-Truth-Conditional Semantics*. New York: Academic Press.