

KNOWLEDGE AS A THICK CONCEPT:

NEW LIGHT ON THE GETTIER AND VALUE PROBLEMS

A Dissertation
Presented to the Faculty of the Graduate School
of Cornell University
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Brent G. Kyle

August 2011

© 2011 Brent G. Kyle

KNOWLEDGE AS A THICK CONCEPT: NEW LIGHT ON THE GETTIER AND VALUE PROBLEMS

Brent G. Kyle, Ph.D.

Cornell University 2011

Abstract: I argue that *knowledge* is a particular kind of concept known as a thick concept. Examples of thick concepts include *courage*, *generosity*, *loyalty*, *brutality*, and so forth. These concepts are commonly said to combine both evaluation and description, and one of the main goals of this dissertation is to provide a new account of how a thick concept combines these elements. It is argued that thick concepts are semantically evaluative, and that they combine evaluation and description in a way similar to that of concepts like *good qua dancer* and *bad qua violinist*. In the final chapter, this basic idea is applied to the concept of knowledge. There, the concept of knowledge is analyzed in terms of *good qua belief*. The upshot is that two central epistemological problems—the Gettier Problem and the Value Problem—can be resolved.

BIOGRAPHICAL SKETCH

Brent G. Kyle attained his B.A. from Tabor College in 1999 majoring in Philosophy and English. In 2003, he received an M.A.R. from Yale University specializing in Philosophy of Religion. He began his Ph.D. at Cornell University in 2005 and completed the degree during the summer of 2011. He is currently teaching Philosophy at the University of New Brunswick, Fredericton.

In memory of my brother,

Bryan R. Kyle

ACKNOWLEDGEMENTS

Many people contributed directly or indirectly to the thoughts expressed in this dissertation. Several of them particularly influenced my thinking by providing especially insightful questions, objections, and comments. These include Matti Eklund, Nicholas Sturgeon, Carl Ginet, Nicholas Silins, and Pekka Väyrynen. My work also benefitted greatly from a Dissertation Writing Group, consisting of Andrew Alwood, Scott O'Connor, Kristen Inglis, and Colin McLear. My gratitude also goes to the Society for the Humanities and the Humanities Council at Cornell University, who generously funded the Dissertation Writing Group. Thanks also to the Cornell Graduate School and the Sage School of Philosophy for funding my travel to two conferences—one on meta-ethics at Texas Tech University, and the other on thick concepts at the University of Kent, UK. The audiences at these two conferences provided invaluable feedback on my work on thick concepts—and for that I am grateful.

I must also thank Virginia Insley, who always made her linguistic intuitions available upon request, and who inspired me to keep writing throughout the frustrating parts of this dissertation. And finally, I owe my utmost appreciation to my parents—Richard and Joyce Kyle, who supported me throughout graduate school—and to my brother, Bryan Kyle, who died in 2002, and to whom I dedicate this dissertation.

KNOWLEDGE AS A THICK CONCEPT:

NEW LIGHT ON THE GETTIER AND VALUE PROBLEMS

by
Brent G. Kyle

Table of Contents

Introduction	1
1 How Are Thick Terms Evaluative?	5
2 Knowledge as a Thick Concept: Explaining Why the Gettier Problem Arises	36
3 How to Analyze a Thick Concept	67
4 Knowledge as Exemplary Belief	103
Conclusion	140
Bibliography	144

Introduction

Over the span of four chapters I argue that *knowledge* is a particular kind of concept known as a thick concept, and that this has important implications for two epistemological problems—the Gettier Problem and the Value Problem.¹

The Gettier Problem is the task of stating necessary and sufficient conditions for knowledge in light of a certain type of example found in Edmund Gettier’s famous paper “Is Justified True Belief Knowledge?” (1963). The Value Problem, on the other hand, is the task of explaining the greater value of knowledge over that of lesser qualities, like mere true belief. Together, these problems are difficulties in understanding the nature and value of knowledge. And I will argue that an inquiry into *knowledge* as a thick concept can shed new light on why these problems arise and how they can be resolved.

But what are thick concepts? These concepts have primarily emerged as a topic of interest in ethics, less so in epistemology. Paradigmatic examples of the thick include ethical concepts like *courage*, *generosity*, *perversion*, *honesty*, *loyalty*, *murder*, and so on. These are typically contrasted with thin concepts (e.g., *good*, *bad*, *right*, and *wrong*), and they are thought to be unique in that their application involves both an evaluation plus a description. For example, if I call someone ‘courageous’, then I have not only evaluated that person (as being good in some way) but I’ve also described her (as being able to face risk).

Ethicists often assume that the evaluation of a thick term is a feature of its semantic content and is therefore not pragmatically associated with the thick term (e.g. via presupposition or

¹ Throughout this dissertation, whenever I refer to a concept (rather than a term or property) I italicize the relevant expression (for example, on this particular page, all italicized words refer to concepts). However, there are a few places where I use italics for providing emphasis, as they are traditionally used. In such places, the context should make it clear that I merely intend to provide emphasis, and not to refer to a concept. Also, in several places throughout this dissertation the symbols ‘#’ and ‘?’ appear before sentences. In such places, these symbols are meant to signal that the sentence in question is linguistically anomalous. In particular, ‘#’ signals that the sentence is highly odd, while ‘?’ signals that the sentence is at least somewhat odd.

implicature). But this assumption has been called into doubt by a number of theorists, including Simon Blackburn, Bruce Brower, and Pekka Väyrynen. In the first chapter of this dissertation—titled “How Are Thick Terms Evaluative?”—I provide a positive argument in favor of this semantic view of thick evaluation, and then defend it against Pekka Väyrynen’s most recent attack involving objectionable thick concepts. This chapter provides an important foundation for the dissertation, since I rely on a semantic view of thick evaluation at various places throughout (esp. chapter 3).

In the second chapter—titled “Knowledge as a Thick Concept”—I move on to the question of what it takes for a concept to count as thick. But here we move beyond mere ethical concepts to an epistemic concept—namely *knowledge*. Essentially, I argue that *knowledge* is a thick concept, and I draw out two implications from this fact. The first is that the Gettier problem is a specific instance of a general problem about thick conceptual analysis. And the second implication is that the analysis of any thick concept (even ethical ones) can run into a problem similar to that of the Gettier problem. The upshot is that we can now explain why the Gettier Problem arises in the first place—it arises because *knowledge* is a thick concept, and a Gettier-like problem is just what we should expect from the analysis of a thick concept.

The third chapter—“How to Analyze a Thick Concept”—focuses on the question of how thick concepts “hold together” both evaluation and description. For example, what explains the fact that I both evaluate and describe someone when I call her ‘courageous’? Theorists have tried to explain this in a number of different ways. But I critique the main contenders and then argue for a new account. My proposal involves a general schema for analyzing thick concepts which I call the Qua Schema. The important part of the Qua Schema is that it incorporates certain notions like being good qua F. Being good qua F is simply a matter of being an exemplar with respect to Fs (i.e. being worthy of imitation in the relevant respects). For example, on my account the courageous

person is a risk-confronter who is good qua risk-confronter (i.e. she's an exemplary risk-confronter). Being a risk-confronter is clearly descriptive, and being good qua risk-confronter is evaluative. And given the way in which these two elements relate, my account tells us how the concept of courage "holds together" an evaluation and a description. Similar things can be said for other thick concepts.

The fourth and final chapter—titled "Knowledge as Exemplary Belief"—uses the previous two chapters as a guide for how to account for knowledge. If *knowledge* is a thick concept (as shown in chapter 2), and if thick concepts can be analyzed according to the Qua Schema (as shown in chapter 3), then it should be possible to analyze knowledge according to the Qua Schema. And this is precisely what I do in chapter 4. In particular, I argue for an account of knowledge that conforms to the Qua Schema, one according to which knowledge is exemplary belief. I then show that such a view provides a new approach to both the Gettier and Value problems. My account of knowledge as exemplary belief can be stated as follows: S knows P just in case S believes P and this belief is good qua belief to the appropriate degree. Of course, the key notion of being good qua belief needs to be properly illuminated before this account will have any purchase. But, as I argue, there are general principles about what it is to be good qua F, and these can be applied to the special case of being good qua belief. As a result, it can be shown that being good qua belief entails both (i) that the belief is true and (ii) that it's not a matter of luck that the belief is true.

Claim (ii) provides the needed resource for addressing the Gettier Problem, since Gettier-style cases are ones in which the subject acquires a true belief as a matter of luck. Moreover, my analysis allows us to explain the greater value of knowledge over that of lesser qualities, like mere true belief. Very roughly, if a person has knowledge, then her belief is good qua belief. And being

good qua belief is an additional value that is absent whenever a person has a true belief without knowledge.

In sum, my dissertation approaches two important problems by first considering the thickness of *knowledge*. My account of a thick concept, together with the claim that *knowledge* is thick, yields important insights into the nature and value of knowledge.

1. How Are Thick Terms Evaluative?

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 metaethical 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 1998; 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).

But 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

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.

¹ Two assumptions ought to be addressed briefly. First, the Semantic View assumes that there is a suitable way of distinguishing between evaluative and non-evaluative content (see Jackson 1998, 120, for opposition). 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). In principle, it should be possible to draw a distinction among 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. (It's also controversial as to how exactly evaluative content is to be understood. But since my argument does not depend on the existence of evaluative content, it therefore does not rely on any particular conception of what evaluative content is.)

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. And I shall argue that we need to make use of conceptual entailment in order to adequately account for the relationship between thick terms and their evaluations.

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*.² 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.

Positive thick terms seem to bear a close relationship to *good in a way*, and negative thick terms appear similarly related to *bad in a way*. But it's not obvious what exactly this relationship is. At the very least, it seems that the typical utterance involving a thick term 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 only to the weaker claim that she is good or bad in *some way or other*.

² Tappolet (2004, 210ff) also thinks there's a tight connection between positive thick concepts and *good in a way*, although she uses the phrase 'good pro tanto' to express this latter concept.

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 uttered S1) normally gives the hearer a defeasible reason to conclude that S2 is true. For example, if I say "The soup is warm" this normally gives you reason to conclude that the soup is not hot. 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 cancellable.

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 "The soup is warm, but not hot," and this does not sound odd. But nothing similar holds true for the relationship between (1) and (2). Consider the following attempt:

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

This sounds awkward and redundant.³ So, it seems unlikely that the connection between (1) and (2) is reinforceable.

The second main feature of conversational implicature is cancellability (Grice 1989, 44). In typical contexts, the fact that I utter “The soup is warm” provides you reason to conclude that the soup is not hot. But this can be canceled—in other contexts you may not have good reason to conclude this. Suppose I say “The soup is warm; in fact, it’s hot.” In this case, my addition of “in fact, it’s hot” 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 cancellable nor reinforceable. We therefore cannot explain this connection by appeal to conversational implicature.

II.3. The Semantic View of ‘Generous’. Although I will consider some 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

³ The oddity of (3) cannot be explained by noting that the second conjunct needs a follow-up clause stating the way in which Nancy is good (e.g. “...she’s good in a way *in that she’s loyal*”). After all, the second conjunct can be stated all alone without sounding the least bit odd—that is, “She is good in a way” sounds perfectly normal without a follow-up clause.

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). More precisely, 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 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

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 the SV-Generous, the second part of (3) is conceptually entailed by what's stated in 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 (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 these are felicitous in their own right, it's plausible 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.

'generous' and 'good in a way'. If I am right that SV-Generous provides the best explanation of that data, then such theorists will have an argument against their view.

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 in which ‘a way’ occurs:

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.⁵ Relative to different contexts, their utterances

⁵ 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 view of context sensitivity, according to which they can both say something true relative to their own context. (See Stanley and Szabó

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.⁶

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:

(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.

⁶ This model is probably too simplistic to handle more complicated quantifier expressions. But it will suffice for current purposes. See Stanley and Szabó (2000, 248-58) for a critique of this simple model and a development of a different semantic model.

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 selective class of contextual domains each containing a way in which that person is good.⁷

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 only predicts 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,⁸ 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

⁷ This view might be trivially true if we include 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 relevant contexts include those with restricted domains (as well as unrestricted ones).

⁸ This is partly because the contexts are innumerable, but also because there is bound to be uncertainty about whether (3)-(5) are felicitous within certain contexts. For example, are (3)-(5) felicitous whenever Scrooge is the speaker? Or are they only felicitous when Scrooge utters them in a community of people who also disvalue generosity? The answers are not completely clear, but a detailed account will depend on how these questions are answered.

and only those contexts in which the speaker is not presumed to disvalue generosity.⁹ No doubt, this example is oversimplified, but the basic strategy should be clear: first, we find a distinctive feature of all the contexts in which (3)-(5) are 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.

⁹ Alternatively, if those sentences are always odd except for when they are uttered within a community of "Scrooges," then the relevant class ought to include all contexts except for those where it is common ground that generosity is disvalued.

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 S2 does not carry I as an implicature.¹⁰ Nothing similar is true for conceptual entailment. Conventional implicata 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 implicata are said to be triggered by words like 'but', 'so', 'even', 'still', 'damn', and 'therefore', as well as appositives and parenthetical constructions.

¹⁰ R.M. Hare seems to believe that the evaluations associated with thick terms are detachable. With regard to the thick term 'courageous' he admits that we do not actually have a "single evaluatively-neutral word" that is truth-conditionally equivalent to 'courageous', but he claims that "we *could* have such a word" (Hare 1963, 188-89). No doubt he would say the same about 'generous'. Also, one way to interpret Blackburn's example involving 'fat' and 'fat↓' is that the former is the detached counterpart of the latter (1992, 290; 1998, 95). Blackburn is not explicit, but perhaps he has something like conventional implicature in mind as an explanation of how thick terms are associated with evaluations.

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 implicata cannot 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 implicata are, in general, not reinforceable.¹² And this is for good reason. Notice that many paradigmatic examples of conventional implicata 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 implicata stated explicitly in the rightmost clause.¹³

But each of (9a-d) sounds normal. Thus, it appears that many conventional implicata are reinforceable. And this means that CI-Generous does not by itself predict the infelicity of (3).¹⁴

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

¹² To my knowledge, Wayne Davis (2010) is the only linguist who even suggests that he believes that conventional implicata in general cannot be reinforced. Strictly speaking, however, what he says is consistent with my claim that some conventional implicata can be reinforced. Potts (2007, 668) claims that the implicata triggered by some particular constructions are not reinforceable, but he shies away from the general claim. See footnote 14 for discussion on those constructions.

¹³ It is possible to quibble over whether the implicata 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, (12b) could be rephrased as “Even Bill passed the test, and that was surprising/unexpected” without creating infelicity.

Proponents of CI-Generous might insist on the reinforceability of conventional implicata 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 implicata 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 (10) should also be felicitous:

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

The boldfaced word in (10) 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 (10) still seems awkward. The felicity of (9a-d) therefore cannot be explained away as suggested.¹⁵

¹⁴ It might be contended that, although some conventional implicata are reinforceable, there are well-defined types of conventional implicata 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 implicata—“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 (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.

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).

¹⁵ One cannot respond merely by noting that the first part of (10) is most naturally read as being about financial institutions. After all, a similar claim can be made for (9d)—the first clause of that sentence is most naturally read as expressing one’s hatred for mowing the lawn (and not for the lawn *per se*).

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):

(11) 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 (11) 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 (11) 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 device¹⁶) is necessary for its felicity—replacing ‘but’ with ‘and’ makes (11) infelicitous. This concession/affirmation structure, according to Horn, explains why sentences like (11) 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.

Does rhetorical opposition explain why the above examples of conventional implicature in (9a-d) are felicitous? At best, it would only account for the felicity of (9a), though it leaves (9b-d) untouched. The concession/affirmation structure required for rhetorical opposition is not present

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).

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 implicata are reinforceable. I do not claim that *all* are reinforceable, but only *some*. 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

(12) 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:

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

(13) clearly implies P in some sense, but it does not take P for granted—after all, (13) might convey P as completely new information. Still, there is a legitimate sense in which (13) seems to background P.¹⁹ The basic idea is that a sentence, such as (13), can convey a number of different propositions, some of which are the main message or point of the utterance (i.e. the foreground entailments),

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

¹⁹ Kadmon (2001, 13) notes that it might be more appropriate to say that P is conveyed “on the side” rather than as background. That slight alteration would not change my arguments that follow.

while others are backgrounded (Levinson 1983). The main message of (13) 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 (12) or (13). But it is worth noting that presuppositions may occur in some contexts and disappear in others. For example, in (12) 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 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 (12) and (13)—appearing respectively as (14) and (15) below—both seem to presuppose that Smith drank Pabst:

- (14) Smith does not regret drinking Pabst.
- (15) 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

- (16) 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) (and of (16)) is something like 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.²⁰

And since the combination of (1) along 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'.²¹ 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 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):

²⁰ The suggestion that (1) may presuppose either Q or R is a possibility that I have loosely modeled after some claims that Väyrynen makes with regard to 'lewd' in his manuscript "Thick Concepts: Where's Evaluation?" (forthcoming). He does not explicitly accept a presupposition account, but this view is implicit in his discussion.

²¹ The attempted justification for this claim would most likely appeal to the behavior of objectionable thick concepts in negations, modals, disjunction, and conditionals (see Väyrynen, forthcoming). 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.

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.²² 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. As we’ve seen, 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?

²² Of course, we could imagine Sue as someone who disvalues generosity, much like Scrooge, in which case her utterance may not sound odd. But that’s beside the point. The point is that, when we do not make any such suppositions about Sue, her utterance seems odd—and that’s what is hard to explain given that neither Q nor R is presupposed by the first part of her utterance of (5).

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) ? Nancy is generous {and/but} she is good in a way.

was taken to be somewhat odd. Indeed, it seems redundant. But similar claims seem true regarding other positive thick terms, such as ‘kind’ and ‘courageous’:

- (17) ? Jones is kind {and/but} he’s good in a way.
 (18) ? Smith is courageous {and/but} he’s good in a way.

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

- (19) ? That was a murder {and/but} it was bad in a way.
 (20) ? That is brutal {and/but} it’s bad in a way.

(17)-(20) 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, (17)-(20) 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) # Nancy is generous and she’s not good in any way

is highly odd, even though each of its conjuncts is perfectly normal. But similar claims seem true regarding ‘kind’, ‘courageous’, ‘murder’, and ‘brutal’:

- (21) # Jones is kind and he’s not good in any way.
 (22) # Smith is courageous and she’s not good in any way.
 (23) # That was a murder and it was not bad in any way.
 (24) # That is brutal and it’s not bad in any way.

(21)-(24) 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, (21)-(24) 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 good representative sample of the class of thick terms, both positive and negative. We therefore have good reason to think that many other thick terms are likely to exhibit similar behavior. In particular, it would be no surprise if ‘loyal’, ‘chaste’, ‘prudent’, ‘wise’, ‘honest’, etc, were also semantically associated with a positive evaluative concept such as *good in a way*. Nor would it be surprising if ‘lewd’, ‘rude’, ‘lie’, ‘dishonest’, ‘unkind’, etc, were semantically associated with a negative evaluative concept such as *bad in a way*. 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 concepts, though I’ve left open the possibility that it might not extend to all thick concepts. This possibility is left open for one main reason. There is no consensus about whether slurs (e.g. ‘kraut’, ‘tart’, etc.) 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.²³ 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 (19)-(20) and (23)-(24). But it’s doubtful that similar data will be available

²³ In *Freedom and Reason*, Hare calls them “secondarily evaluative words” rather than thick terms (1963, 25). The latter terminology did not become vogue in ethics until Bernard Williams’ *Ethics and the Limits of Philosophy* (1985). Still, in Hare’s later works (1995, 325; 1997, 54 and 61) it becomes clear that he takes the class of secondarily evaluative words to be identical to the class of thick terms.

for all slurs. For example, the sentence “Jan is a kraut and he’s not bad in any way” strikes me as infelicitous, but its infelicity seems wholly attributable to the first conjunct “Jan is a kraut.” Thus, whether or not we have reason to extend the Semantic View to all thick terms will depend greatly 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 seem to establish this.

At most, slurs lie at the periphery of the class of thick concepts; they are not paradigms. But, as I’ve argued in this section, the Semantic View is highly plausible with regard to the paradigmatically thick. Now that I have presented my case for the Semantic View, we can consider how well this view can stand 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 (Gibbard 1992, 279; Eklund 2011, 30;

Väyrynen 2009, 440ff), although some might dispute whether these are really objectionable.²⁴

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.²⁵ 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 a sentence of the following form:

(25) 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 (25), namely

(26) A is not chaste.

Their unwillingness to assert (26) is initially puzzling, since it's plausible that chastity-objectors should take (26) to be true. It looks as if these negations also endorse the kind of values rejected by the chastity-objector. And initially this might be taken to suggest that the evaluative content of 'chaste' is not part of its truth-conditions but is instead one of its presuppositions.²⁶

Väyrynen thinks the reluctance of chastity-objectors to assert (26) presents a problem for the Semantic View. In particular, Väyrynen advances the following claim, which I shall call "VC":

²⁴ There has been much discussion on whether *sexual perversion* embodies values that ought to be rejected (Nagel 1969; Priest 1997; Williams 2003; Williamson 2004). Other supposedly objectionable thick concepts have not received as much attention.

²⁵ 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 did, 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 whole-hearted believers who wish to avoid 'fornicate', 'defile', etc.

²⁶ A presupposition view is explored in Väyrynen (2009, 453ff), and such an account is implicit in his manuscript "Thick Concepts: Where's Evaluation?" (forthcoming). See section II.4.b. of this paper for more discussion.

VC: If those who reject the evaluative content associated with *chaste* are not willing to apply truth-conditional negation to (25), then, ceteris paribus, that evaluative content is not part of the truth conditions of (25).²⁷

The basic thought behind Väyrynen's claim is this: If what the chastity-objector rejects is located within the truth-conditions of (25), then truth-conditional negation of (25) should be a perfectly acceptable way of expressing her disagreement (Väyrynen 2009, 448).

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 (25). In other words, they are not willing to assert (26). But this datum, in combination with VC, entails that (ceteris paribus) the evaluative content is not part of the truth-conditions of (25).²⁸ 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, p. 449). So, Väyrynen's argument threatens to pose a more general problem for the Semantic View.

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*.²⁹ But Väyrynen overlooks the possibility that a speaker's reluctance to assert (26) might have something to do with her reluctance to be saddled with a claim about A's being unchaste. In particular, notice that

(26) A is not chaste.

²⁷ 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).

²⁸ 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.

²⁹ 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 (25) and (26) should therefore not be taken seriously as data for this discussion.

clearly seems to imply

(27) 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,” then that will likely convey the stronger claim that Smith is unhappy.

Linguists typically say that claims like (26) and (27) are not truth-conditionally equivalent, but that (26) implies (27) by virtue of conversational implicature (Horn 1989, 331ff; Levinson 2000, 127ff).

Let’s assume this is true for present purposes.³⁰ The relationship between (26) and (27) appears to exhibit negative strengthening. That is, (26) implicates (27) 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 (26) because that assertion would implicate (27). And (27) 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.

³⁰ 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, (26) is equivalent to the following:

(26*) 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 (26) and (27). On this “contrary” reading of ‘not’, (26) is actually equivalent to (27). If this sort of view is correct, then we can explain the speaker’s reluctance to utter (26) by citing the fact that this utterance is ambiguous between something she most likely accepts—(26*)—and something she outright rejects—namely (27). Moreover, we can note that (27) is actually the more common reading of (26), which suggests that someone who finds *chaste* objectionable would withhold uttering (26) without further clarifications.

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 (25).³¹ VC must therefore be rejected.

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

(28) The story is not blasphemous.

seems to imply something like

(29) The story is somewhat reverent.³²

³¹ 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 (25). And this account in no way impugns the fairly common view that (A) has an evaluative truth-condition. By analogy, the story I've told about *chaste* is also consistent with claiming that (25) has an evaluative truth-condition.

³² 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).

And someone who thinks *blasphemous* is objectionable may not want to utter (28) for fear of being saddled with something like (29). Perhaps this explains Oscar Wilde's reluctance to utter (28) when faced with the attorney's aggressive questioning: "Did you or did you not consider the story blasphemous?" (Wilde and Carson 1895).³³

To be sure, 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 by drawing attention to a completely different class of speakers—those who whole-heartedly accept the values associated with a thick concept. With regard to *chaste*, we can call them "chastity-accepters." Chastity-accepters will likely grant that there are neutral cases involving people who are neither chaste nor unchaste. And with regard to these neutral cases, the chastity-accepter will also want to avoid uttering claims like (26) for fear of implicating (27). So, my view explains the reluctance of chastity-accepters in the same way as it does for chastity-objectors—they both want to avoid implicating (27). But this might seem problematic, since the reluctance of each speaker seems importantly different. However, let me point out that this objection overlooks that fact that my explanation can grant that each speaker rejects (27) for very different reasons. The chastity-objector rejects (27) because she thinks that the whole scale from chaste to unchaste ought

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 (29).

³³ For another possible explanation of Wilde's reluctance, see footnote 25.

to be rejected. On the other hand, the chastity-accepter rejects (27) because he thinks that, while many people are unchaste, the particular individual in question cannot be classified as such. In short, the fact that these two speakers reject the same implicature for vastly different reasons is sufficient to explain the difference between their stances towards (26).

An additional worry 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,

- (30) Smith is not chaste, but neither is he unchaste.
- (31) 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 (30) and (31)? No. We can understand the follow-up clauses in (30) and (31) as elements that cancel the conversational implicature from “Smith is not chaste” to “Smith is unchaste.” The cancelation is obvious in (30), since that implicature is explicitly denied by the second part of (30). In (31) the implicature is not explicitly denied, but it’s plausible that the follow-up clause provides enough reason to doubt that the speaker intends to convey that Smith is unchaste. So, the problematic implicature is not generated by (31). Thus, my appeal to negative strengthening allows for the acceptability of (30) and (31) to chastity-objectors.³⁴

³⁴ Väyrynen tries to explain the fact that chastity-objectors find sentences like (31) 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 (31) can be seen as a case in which negative strengthening is canceled, then I see no reason to postulate that (31) is metalinguistic. Moreover, it’s worth pointing out that (31) 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 (31). If we replace ‘not chaste’ in (31) with ‘unchaste’, the result is not the least bit unintelligible:

(31’) 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 (31) to be metalinguistic. Of course, Horn’s incorporation test is not uncontroversial. Geurts (1998, 280) is one critic.

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.³⁵ Väyrynen provides the following example (2009, 448):

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

According to Väyrynen, a chastity-objectors reluctance to assert (32) would be initially puzzling, since these people should take (32) 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 (32), 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 (32) 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

³⁵ A different variety of conditional is where the objectionable thick term appears solely in the consequent of the conditional. But in this case there's no automatic reason to think that chastity-objectors should find the conditional to be true.

would expect with regard to (26). 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 (32), 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 (32) in many contexts seems to imply

(33) 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 (32) and (33) is a type of conversational implicature known as “clausal implicature”.³⁶ 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, 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 (33) in certain contexts, since they fail to believe that abstinence may be chaste. And since (32) conversationally implicates (33), they would be reluctant to utter (32) in those contexts.³⁷

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”),

³⁶ My statement of (33) is modeled after Levinson’s way of stating clausal implicature (2000, 36-7). The occurrence of ‘may’ is meant to signal epistemic possibility (Levinson, 2000, 109). Gazdar (1979, 50, 60) provides a slightly different gloss.

³⁷ Of course, the contexts in which objectors are reluctant to utter (32) would need to “line up” with those in which they would take issue with (33), but I see no immediate reason to think that won’t be the case.

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.

In summary, we’ve seen in this section that there are initially puzzling data regarding the linguistic behavior of speakers who find certain thick concepts objectionable. But I’ve argued that this does not license a rejection of the Semantic View, because other widely-accepted resources can be employed in explaining that data. My goal in this section has not been to explain *all* data arising from the behavior of objectionable thick concepts. Rather, I have only aimed to show that Väyrynen’s objection overlooks an important (and, I might say, promising) reply available to the proponent of the Semantic View. In particular, it’s plausible that a speaker’s reluctance to utter sentences involving objectionable thick terms is explainable through pragmatic relations like negative strengthening and clausal implicature. And an appeal to such relations is perfectly-well available to the proponent of the Semantic View.

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. 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.

2. Knowledge as a Thick Concept: Explaining Why the Gettier Problem Arises

I. Introduction.

The Gettier problem has stymied epistemologists. After Edmund Gettier proved the insufficiency of the tripartite analysis of knowledge, his paper was followed by a barrage of reformulations, most of which were returned with potential counterexamples.¹ A succession of reformulations, and more and more counterexamples, continued for over thirty years. However, within the past decade the general impetus to solve the Gettier problem has significantly receded (with only a few exceptions²). But the abatement of this project is not due to any consensus over one particular analysis—quite the opposite. Despite the great amounts of time and detail invested in the project, no uncontroversial analysis has emerged. Indeed, some even believe the number of inadequate analyses indicates that no correct one is forthcoming.³

Whether or not the Gettier problem is resolvable, we still must face an important question: Why does this problem arise in the first place? So far, philosophers have seen the Gettier problem as either a problem peculiar to the concept of knowledge,⁴ or else an instance of a general problem about conceptual analysis.⁵ But I would like to steer a middle course. I think the problem arises because *knowledge* is a particular kind of concept known as a thick concept, and a Gettier-like problem is just what we should expect from attempts at analyzing a thick concept.⁶ These claims unfold as follows.

¹ See Gettier (1963). And see Shope (1983 and 2004) for a survey of responses.

² For example, see Greco (2003). And Pritchard (2005) seems like he might also be defending an analysis in the traditional sense.

³ For example, see Williamson (2000, 30).

⁴ This is strongly suggested by the way in which epistemologists tend to approach the Gettier problem. See Shope (1983 and 2004) for a survey. No one, as far as I know, has suggested that very similar problems might arise for a broader class of concepts of which *knowledge* is a member.

⁵ For this type of suggestion, see Laurence and Margolis (1999, 15) and Williamson (2000, 30-3).

⁶ Throughout this paper, whenever I refer to a concept rather than, say, a term or a property, I italicize the relevant expression.

Section II aims at developing my central claim that *knowledge* is thick. But what are thick concepts? Very roughly, these are hybrid concepts in that they have both evaluative and non-evaluative content.⁷ Paradigmatic examples of the thick include *honest*, *perverse*, *generous*, *courageous*, *murder*, *insightful*, *stupid*, *clever*, and so on. They are naturally contrasted with thin concepts like *good*, *bad*, *right*, *wrong*, *better*, *worse*, *ought*, *obligation*, and *permission*.⁸ But, as we'll see in section II.3, not just any hybrid concept counts as thick. It is very easy to manufacture hybrid concepts by conjoining evaluative and non-evaluative contents in an ad hoc fashion (e.g., *good and red*). But this maneuver does not generate a thick concept. These ad hoc combinations motivate a further condition on thickness (explained in section II.4). And in section II.5, I show that *knowledge* satisfies this further condition for reasons having to do with Gettier-style cases.

Section III develops a general problem for analyzing thick concepts by appealing to this further condition on thickness. I argue that this condition generates a problem for the analyses of all thick concepts, and that the Gettier problem is an instance of this general problem. Thus, the Gettier problem arises because *knowledge* is thick. For present purposes, I do not take a stand on whether the Gettier problem, or its general counterpart, is resolvable. My primary aim is to bring these problems into better light.

The above claims also have important consequences in ethics. If the Gettier problem is an instance of a general problem about the analyses of all thick concepts, then we can expect that the analysis of any given thick concept will be subject to a problem similar to that of the Gettier problem. And this includes some of our prized ethical concepts. To the extent that the Gettier

⁷ I use the label 'evaluative' broadly enough to include, not only concepts like *good*, but also normative concepts like *ought*. It should be noted, however, that my usage of this label does not presuppose any particular way of drawing the distinction between evaluative and non-evaluative contents. For example, I remain neutral on whether theories like expressivism and prescriptivism provide adequate accounts of the evaluative.

⁸ Some theorists (such as Elgin (2008, 372) and Simmons (1997, 148)) allow wholly non-evaluative concepts like *red*, *grass*, and *green* to count as thin. This diverges from the primary usage of 'thin concept', where the thin is taken to be a subclass of the evaluative. I shall use the notion of a thin concept in the latter way, as picking out a subclass of evaluative concepts.

problem should trouble epistemologists, very similar problems should also trouble ethicists who take interest in thick concepts like *generous*, *honest*, *perverse*, *cruel*, *courageous*, and so forth. I outline a Gettier-like problem for *cruel* in section III.

Before moving on to my argument for the thickness of *knowledge*, let me note that this conclusion conflicts directly with the dominant view about *knowledge*. A number of philosophers have recently claimed that both *knowledge* and *justification* are thin concepts.⁹ I find it quite plausible that *justification* is thin, as I'll explain below, but I think their claim about *knowledge* is mistaken. Since these epistemologists have not bothered to explain or argue for their classification of *knowledge* as thin, I will not address them directly in this paper. Whatever their reasons, I will show that their classification is importantly mistaken.¹⁰ *Knowledge* is thick, and this fact allows us to classify and extend one of the main problems about that concept—the Gettier problem.

II. The Thickness of Knowledge.

II.1. Two Conditions. In order to tell whether the concept *knowledge* is thick, let's consider whether the term 'knowledge' expresses a thick concept. We need conditions that are jointly sufficient for a term to count as expressing a thick concept. Let's begin with a common view in the thick concepts literature. For any term 'Tk', 'Tk' expresses a thick concept if

- (c1) 'Tk' expresses a concept with evaluative content; and
- (c2) 'Tk' expresses a concept that also has (significant) non-evaluative content.

⁹ See Battaly (2008, 435), Axtell and Carter (2008, 427), Elgin (2008, 371), Thomas (2008, 363), and Väyrynen (2008, 392).

¹⁰ Why have they classified *knowledge* as thin? I am ultimately unsure, but here are three possibilities. One is that they have latched on to the generality of thin concepts and, thinking that *knowledge* is more general than, say, *intelligent*, they have thereby classified the former as thin and the latter as thick. Still, these claims about relative generality do not guarantee the correct verdict. A concept might be more general than other thick concepts, and yet still be thick. (By analogy, a person might have greater height than other short people, and still be short.) A second possibility is that they have accepted the characterization of thin concepts noted in footnote 8 and they have also assumed that *knowledge* is wholly non-evaluative. If this is their view, then we'll see more directly why they are mistaken in section II.2. A third possibility is that they do not think that *knowledge* has enough non-evaluative content for it to count as thick. If this is their thought, then I address that too in section II.2.

These conditions represent a very common way of spelling out the hybrid nature of thick concepts.¹¹ For example, it's plausible that 'murder' expresses a concept with both evaluative content (e.g., *wrong*) and non-evaluative content (e.g., *deliberate killing*).

I ultimately do not think (c1) and (c2) are sufficient for thickness. (I explain why in section II.3, and then supplement (c1) and (c2) with a third condition.) Still, these two conditions give us a nice place to start. They state that thick terms have both evaluative and non-evaluative content associated with them. But what does it mean for a term to "have" these contents? In other words, how exactly are thick terms associated with these contents?

This is a highly controversial issue in the thick concepts literature, and rather than resolving it here I will try to remain neutral between two main ways of approaching it.¹² One prevalent option is to say that thick concepts can be analyzed in terms of both evaluative and non-evaluative content. Typically what this means is that there are non-circular, necessary, and sufficient conditions for application of the thick concept, conditions that express the associated evaluative and non-evaluative contents. Although many theorists accept this view, there are also many who oppose it by claiming that thick concepts cannot be analyzed.¹³ Because one of my main goals is to cast light on the issue of thick conceptual analysis, it is fitting to employ a notion that would be amenable to both positions. So, we should not interpret (c1) and (c2) as saying that *Tk* is analyzable in terms of evaluative and non-evaluative contents.

¹¹ For example, see Väyrynen (2008, 390; 2009, 439), Gibbard (1992, 273), Sreenivasan (2001, 3), and Payne (2005, 91).

¹² The two views discussed in this paper might be aptly classified as semantic theories of the relation between thick terms and their evaluative contents. But some people reject semantic theories and instead hold that the evaluative contents are only pragmatically associated with thick terms (e.g. see Väyrynen (2009)). Although I cannot presently justify my assumption that a semantic theory is correct, this view is argued for in detail in my paper "How Are Thick Terms Evaluative?"

¹³ Those who accept the analyzability of thick concepts include Gibbard (1992), Hurka and Elstein (2009), Tappolet (2004), Burton (1992), and Payne (2005). Those who reject it include Dancy (1995), Platts (1979, 244), Elgin (2008, 372-3), and Brewer (2009, 187). Also, Elgin (2005, 343) seems to interpret Bernard Williams as rejecting the analyzability of thick concepts.

A slightly different option is to understand (c1) and (c2) as saying that *Tk* conceptually entails both the evaluative and non-evaluative contents. One worry about this option is that some forms of conceptual entailment are not restrictive enough for this interpretation to be plausible. For example, on some epistemic views of conceptual entailment, all thin concepts will conceptually entail a number of trivial non-evaluative concepts, like the concept *self-identical*.¹⁴ But this means that all thin terms will likely satisfy both (c1) and (c2). Moreover, if Judith Thomson is correct in claiming that everything is good in some way or other, then it looks like all non-evaluative concepts will conceptually entail the evaluative concept *good in a way*.¹⁵ This means that all non-evaluative terms would also satisfy (c1) and (c2). But it would be bizarre if all thin terms and all non-evaluative terms satisfied (c1) and (c2), since these conditions are supposed to capture what's most distinctive about the thick. So, it looks like we need a more restrictive way of interpreting these conditions.

When ethicists propose (c1) and (c2) as conditions that capture what's distinctive about thick terms, they seem to have in mind that the evaluative and non-evaluative contents in question are the ones that would be featured within 'an account' of the thick concept. In other words, they seem to understand (c1) and (c2) as claiming that *Tk* is to be accounted for in terms of an evaluative content and a non-evaluative content. I propose that we interpret (c1) and (c2) in this way.

But what is this accounting-for relation supposed to be? What we should have in mind with regard to the accounting-for relation is similar to a notion employed by Susan Hurley. On Hurley's view, in order to account for X in terms of Y we must provide "an account" framed in terms of Y "of what we mean, understand, and intend ourselves to be talking or thinking about, when we talk or think about X."¹⁶ What this suggests is that the accounting-for relation provides a restrictive

¹⁴ According to one epistemic view, *P* is a conceptual truth if it is necessarily true that whoever fully grasps the thought *that P* will assent to it. The above-mentioned example would satisfy this schema, if anything does.

¹⁵ Thomson (2008, 10).

¹⁶ Hurley (1989, 10). On her view, the accounting-for relation does not necessarily involve conceptual analysis, conceptual priority, or conceptual independence.

subclass of conceptual entailments, a subclass that includes the entailments that would constitute an account of the thick concept but that excludes the trivial entailments mentioned in the previous paragraph. This is admittedly vague, but for present purposes we need not spell this out in detail. An intuitive notion of what would be included within an account of a thick concept will suffice for present purposes.

However, it is worth being explicit about one potential feature of this notion of an account. We must be open to the possibility that the relevant sort of account could be circular. (This is one way in which a full-blown account may differ from a conceptual analysis). For example, it is possible that X is to be accounted for in terms of Y, while also that Y is to be accounted for in terms of X. The reason we should leave open the possibility of a circular account is because some theorists, like Jonathan Dancy, deal with individual thick concepts in precisely this way.¹⁷ Dancy suggests that we can account for each thick concept in terms of an object having certain non-evaluative features “in the right way.” For instance, an act is *lewd* because it displays sexuality in “the right way.” But once Dancy explains what he means by “the right way” it becomes clear that the sort of account he has in mind is circular. According to Dancy, the right way associated with *lewd* “can only be captured” by saying that it is the way that merits “a response merited by the lewdness.”¹⁸ Even though this is circular, Dancy is still accounting for *lewd*. And this means that our accounting-for relation should not automatically rule out circularity.¹⁹

In short, I suggest we interpret (c1) and (c2) in terms of the accounting-for relation. ‘Tk’ expresses a thick concept if

¹⁷ Wiggins (1987) also has a similar view, but it’s not completely clear whether Wiggins has thick concepts in mind or just evaluative concepts more generally.

¹⁸ Dancy (1995, 276).

¹⁹ Of course, the account in question can be informative, even if circular. Dancy’s method certainly requires us to say interesting things about each thick concept—lewdness concerns the way sexuality is displayed, and courage “concerns fear and danger” (1995, 277). But, in connection with this, it should be made explicit that the accounting-for relation need not involve conceptual priority or asymmetrical dependence. On Dancy’s view, for instance, the “right way” associated with *lewd* is neither prior to nor independent of *lewd*. To borrow a term from Susan Hurley (1989, 11), we might say these two notions are “interdependent.”

(c1) ‘Tk’ expresses a concept that is partly accounted for in terms of an evaluative content.

(c2) ‘Tk’ expresses a concept that is partly accounted for in terms of a (significant) non-evaluative content.

I will explain the parenthetical “significance” qualifier in the next section. Let us now consider whether ‘knowledge’ satisfies (c1) and (c2).

II.2. Does ‘Knowledge’ Satisfy (c1) and (c2)? There’s a very simple argument in favor of the claim that ‘knowledge’ satisfies these conditions.

Traditionally it has been claimed that *knowledge* entails *justified true belief* and that the former is to be accounted for (at least partly) in terms of the latter. If that’s right, then it’s likely that ‘knowledge’ satisfies both conditions. Since *justified* is plausibly evaluative, it would follow that *knowledge* is to be partially accounted for in terms of an evaluative content. So, under these assumptions, ‘knowledge’ satisfies (c1). Moreover, *true belief* seems to be a (significant) non-evaluative content. So, ‘knowledge’ appears also to satisfy (c2). Given these assumptions, it looks quite plain that ‘knowledge’ satisfies both conditions.

Of course, things are not quite so simple. Although I think this argument ultimately succeeds, it nonetheless makes a number of assumptions. So, allow me to clarify and substantiate five of them in what follows.

The first assumption is that the argument relies substantively on a few entailments that have been called into doubt—namely that *knowledge* entails *truth*, *belief*, and *justification*.²⁰ Space prohibits a detailed defense of these entailments, and so, for the purposes of this paper, I will assume that they hold. This assumption represents the majority perspective in epistemology. In particular, theorists

²⁰ For a recent challenge to the knowledge-truth entailment see Hazlett (forthcoming). For a challenge to the knowledge-belief entailment, see Lewis (1996, 556). The most controversial is the entailment from knowledge to justification. For dissenters, see Audi (1995 and 2001), Kornblith (2009), and Alston (1989). It’s worth pointing out that these latter dissenters must face an important challenge presented by Timothy Williamson (2007, 111-12).

who aim to solve the Gettier problem typically assume these entailments. And since a major goal of this paper is to shed light on the Gettier problem, it is fitting that we also take them on board.

Second, in addition to assuming these entailments, the above argument also assumes that *knowledge* is to be accounted for in terms of *belief*. This is the traditional view, but there have been recent attempts at reversing this picture. Williamson in particular has recently proposed that *belief* is to be accounted for in terms of *knowledge*. On his view, “to believe *p* is to treat *p* as if one knew *p*.”²¹ The main point to emphasize here is that these two views are not incompatible. As noted, the accounting-for relation does not exclude the possibility of circular accounts. So, it’s possible that *knowledge* is to be accounted for in terms of *belief*, even if *belief* is to be accounted for in the way Williamson suggests.

Third, and most importantly, the argument also assumes that *justification* is evaluative. As noted at the outset, most people take *justification* to be a thin concept. And the claim that *justification* is evaluative just follows from their claim. But is *justification* really a thin concept? I think this is quite plausible, provided we take ‘S’s belief in *p* is justified’ to be conceptually equivalent to one of the following:

- ‘S has good reason to believe *p*’
- ‘S’s belief in *p* is permissible’
- ‘S’s grounds for believing *p* are adequate’
- ‘It’s okay for S to believe *p* (on those grounds)’
- ‘It is better that S believes *p* than not’.

These are common ways of understanding justification. And the concepts *good*, *permissible*, *adequate*, *okay* and *better* are paradigmatic thin concepts. So, if we conceptualize justification in one of these ways, then it’s plausible that *justification* is thin.²² Of course, the above ways of understanding

²¹ Williamson (2000, 46).

²² Of course, the terms ‘justified’, ‘good’, ‘permissible’, ‘okay’, ‘better’, and ‘adequate’ are likely to be context sensitive. They may express different concepts in an ethics conversation from what they would express in an epistemology conversation. Still, this fact does not preclude the rather intuitive claim that I’m making. It’s plausible that the various concepts expressed within each context are thin evaluative concepts, whether ethical, epistemic, legal, etc. It

justification are compatible with a great variety of technical theories of justification (e.g. in terms of reliable belief formation or internal mental states). These technical specifications can be seen as providing the base on which justification supervenes.²³ Thus, I take it as plausible that *justification* is a thin concept. I will rely on this claim later in the paper. But for now, suffice it to say that it entails our desired claim that *justification* is an evaluative concept.

And fourth, the argument also assumes that *true belief* is non-evaluative. One might initially worry that *true belief* is an evaluative concept, because many true-F concepts are evaluative (e.g., *true friend*). However, it is implausible to regard *true belief* similarly. The bare notion of true belief that we're operating with admits of a definition that reveals no evaluative dimension whatsoever—a person S has a true belief that P if and only if it is the case that P and S believes that P. This definition is all we need for present purposes, and it contains no traces of evaluation. So, the concept *true belief* is plausibly non-evaluative.²⁴

And finally, there is an issue about the parenthetical significance qualifier that appears in (c2). Occasionally ethicists emphasize that a thick concept's non-evaluative content is "significant" in a certain sense.²⁵ If we go along with this, then we might worry that, even if *true belief* is non-evaluative, it may not be significant in the relevant sense. But what is meant by 'significant'?

is hard to see how this claim can be rejected. One source of opposition may come from those who distinguish thick from thin by virtue of their relative specificity and generality. Such theorists might claim that *epistemically justified* is not thin because it is more specific than concepts like *all things considered justified*. But, as I already mentioned in footnote 10, these claims about relative specificity do not guarantee the correct verdict. A concept might be more specific than other thin concepts, and yet still be thin.

²³ For a reliabilist view of this sort, see Goldman (1986, 22-7). For an internalist theory of this sort, see Conee and Feldman (2001).

²⁴ Of course, this claim is compatible with the widely accepted claim that true beliefs are highly valuable. Consider an analogy: the concept *diamond* is non-evaluative, even though diamonds are highly valuable.

²⁵ For the most straightforward use of this significance qualifier, see Väyrynen (2008, 390-91; 2009, 439). Some theorists worry that paradigmatic thin terms (e.g. 'ought') have non-evaluative content and that these would mistakenly count as thick unless we add the significance qualifier. For instance, if *ought* conceptually entails *can*, then we might worry that *ought* is going to count as thick. But the significance qualifier would allow us to say that 'ought' does not have enough non-evaluative content, and that's why it doesn't count as thick. To my mind, however, this worry does not adequately motivate the significance-qualifier. I've claimed that a distinguishing feature of thick concepts is that they are to be accounted for in terms of a non-evaluative content. But, whatever the relation is between *ought* and *can*, it's not at

The significance qualifier most likely stems from a theory about the relation between thin and thick concepts. Some ethicists believe that thin and thick lie on opposite ends of a continuum of evaluative concepts, differing from each other in terms of how much non-evaluative content they have (if any).²⁶ On this view, a concept must have enough non-evaluative content if it's to belong on the "thicker end" of the continuum. Having significant non-evaluative content is therefore a matter of having a sufficient amount of non-evaluative content. It's not at all clear how these ethicists plan to measure amounts of content, and I cannot attempt to clarify the issue here. But I think the basic worry can be addressed without delving into the matter.

The basic worry is that, although *knowledge* is to be accounted for in terms of some non-evaluative content (namely *true belief*), this is not enough of a non-evaluative content for us to correctly place *knowledge* on the thick end of the continuum. However, suffice it to say that this worry will probably over-generate by excluding many paradigmatic thick concepts. For example, there appears to be no more reason to deny that *true belief* is enough of a non-evaluative content than there is to deny that *deliberate killing* is enough.²⁷ But many will agree that *deliberate killing* is the non-evaluative content associated with *murder*.²⁸ And *murder* is a paradigmatic thick concept. So *knowledge* appears to be on par with a paradigmatic thick concept, and it therefore looks implausible to deny that *knowledge* has enough non-evaluative content.

all clear that *ought* is to be accounted for in terms of *can*. And this is required for 'ought' to satisfy (c2). Thus, I think more needs to be said to motivate this significance-qualifier.

²⁶ See Väyrynen (2008, 391). Williams (1995, 234) and Brewer (2009, 185-86) also seem to have this sort of view in mind. The basic idea behind the continuum view most likely originates in Scheffler (1987, 417).

²⁷ There are two points to make here. First off, if one thinks the non-evaluative content of *murder* should be built up to something more substantial, like *deliberate human killing*, then it should also be noted that there are cases to be made for building up the non-evaluative content of *knowledge* to something like *reliable true belief*, *safe belief*, or *sensitive belief*.

And second, one might think that deflationism about the truth-predicate threatens to reduce the amount of content of true belief. I actually don't see why this would be the case. Consider that 'true belief', as I'm construing it, need not be defined by using the truth-predicate: S has a true belief that p just in case S has a belief that p and it is the case that p. Alternatively, "it is the case that" could drop out of the definition without problems. Either way, it's not at all clear why this would mean that *true belief* lacks a sufficient amount of content.

²⁸ See Suter (1973, 361), Anscombe (1979, 74), and Priest (1997, 360).

In short, it's plausible that the concept expressed by 'knowledge' is to be accounted for in terms of a significant non-evaluative content—*true belief*. This means that 'knowledge' satisfies (c2). We've also seen that 'knowledge' plausibly satisfies (c1). So, there is at least a strong case for the claim that 'knowledge' satisfies the conditions that are commonly seen as sufficient for thickness.

But I will show in the next section that (c1) and (c2) are not quite sufficient for thickness. Although these conditions are a nice place to start, many ethicists have ultimately had a more restrictive conception of thickness in mind. Over the next three sections I add a further condition to (c1) and (c2) and then argue that 'knowledge' also satisfies that condition. This further condition provides a connection with the Gettier problem.

II.3. Ad Hoc Conjunctions. Imagine that 'gred' expresses the concept *good and red*, and that 'wronday' expresses the concept *wrong and performed on Monday*. For obvious reasons, we can call these ad hoc conjunctions. These terms satisfy both (c1) and (c2), but they are not thick. Thick concepts cannot be manufactured simply by conjoining a thin concept with a wholly non-evaluative concept. Catherine Elgin has explicitly excluded ad hoc conjunctions from the thick.²⁹ And there seems to be good reason for this exclusion. 'Gred' and 'wronday' seem importantly different from paradigmatic thick terms like 'murder', 'courageous', 'generous', 'lewd', and so on. For example, while 'gred' expresses the concept *good and red*, it's quite implausible that 'courageous' expresses a concept like *good and risk-confronter*. A foolhardy person might be good for reasons having nothing to do with her being a risk-confronter. But this wouldn't mean she's courageous. Metaphorically speaking, the evaluative and non-evaluative contents of a paradigmatic thick term seem more entwined than those of an ad hoc conjunction.

²⁹ See Elgin (2008, 372).

Moreover, there are considerable theoretical motivations for excluding ad hoc conjunctions from the thick. As noted earlier, there is much dispute over whether thick concepts are analyzable. And the dominant positions on either side of this debate will want to exclude ad hoc conjunctions. Consider those who claim that thick concepts are unanalyzable. This view would be an obvious nonstarter if ad hoc conjunctions were thick. It's clear that the concepts expressed by ad hoc conjunctions are analyzable (e.g., x is gred if and only if x is good and x is red). So, these theorists must exclude ad hoc conjunctions from the thick in order to maintain cogency.

Others believe thick concepts can be analyzed. These ethicists typically offer certain schemata for analyzing thick concepts, schemata that display the interconnections between the evaluative and non-evaluative contents. But ad hoc conjunctions cannot be analyzed in accord with these schemata. For example, it's often claimed that each thick concept can be analyzed in terms of an evaluative and a non-evaluative content, the former of which obtains "in virtue of" or "on account of" the latter.³⁰ But it is obvious that ad hoc conjunctions do not admit of this style of analysis. My car is gred even though it's not good in virtue of being red. Thus, ad hoc conjunctions are problematic for both views; there are theoretical reasons to exclude them from the thick.

In the next section I provide a third condition on thickness. Since we now want to exclude ad hoc conjunctions from the thick, we are not merely seeking a condition that, together with (c1) and (c2), is sufficient for thickness. We also want a necessary condition.

II.4. A Third Condition. How are ad hoc conjunctions to be excluded from the thick? One option would be simply to pick one of the above positions, and then to use its stance towards thick

³⁰ Hurka and Elstein (2009) have a view like this. Tappolet (2004) and Burton (1992) use the 'in virtue of' relation. Payne (2005) oscillates between 'in virtue of' and 'is a reason for'. Blackburn (1984, 148) provides the 'on account of' relation, though he is talking mainly about pejoratives, and it is not clear whether this can be extended to an account of thick terms. Gibbard's view (1992) is more complicated than these others, but the same claim will apply to his view as well.

conceptual analysis as a way of ruling out ad hoc conjunctions. But this would be dialectically problematic. In that case, the proponents of the opposing view would strangely be forced into claiming that none of the paradigmatic thick concepts are actually thick.³¹ Is there a better way of excluding ad hoc conjunctions? I think so. Here I'll try to lay out a third condition on thickness while remaining neutral on the question of whether thick concepts are analyzable. Let's consider an important point on which all the above mentioned theorists can agree.

Both parties can grant that thick concepts are to be accounted for in terms of other contents. In particular, it's plausible that each thick concept can be accounted for in terms of at least one thin concept. For example, most theorists would say that *murder* is to be accounted for in terms of *wrong*. Others, such as Dancy, might say that *murder* is to be accounted for as a killing that merits a certain response.³² On this view, the relevant thin concept would be *merit*. Henceforth, let Tn stand for the thin concept that's used within the account of a given concept.³³ So the claim here is that, if Tk is a thick concept, then there is some thin concept Tn such that Tk is to be accounted for in terms of Tn .

Moreover, it's plausible that each thick concept is to be accounted for in terms of a cluster of non-evaluative concepts. For example, *murder* is clearly to be accounted for in terms of *killing*, *deliberate*, and *deliberate killing*. But notice that some of the non-evaluative concepts within the relevant cluster will be logically stronger than others. For example, *deliberate killing* is stronger than *killing*. So, let's focus on the strongest of the group, that is, on the non-evaluative concept that is stronger than any other in that cluster. Let N stand in for that concept. More generally, let's say

³¹ There are two other options. One is to state a disjunctive condition as necessary for thickness, where each of the above positions is represented in one of the disjuncts. But this is clearly unattractive, as it does not provide a unified condition on thickness. The second is to require that a thick concept's evaluative content obtain in virtue of the non-evaluative content. This could be claimed without holding that thick concepts are analyzable. But I find this view problematic because it looks like it will lead to the emptiness of many intuitively non-empty thick concepts. For a detailed discussion on this, see my manuscript "How to Analyze a Thick Concept."

³² Here I'm conjecturing based on how he accounts for *lewd*. See Dancy (1995, 275-76).

³³ If there is more than one such thin concept, then let Tn stand for their conjunction.

that *N* is a non-evaluative concept that exhausts the non-evaluative portion of the relevant account (whether or not we're accounting for a thick concept).

Of course, ad hoc conjunctions have similar concepts associated with them. *Gred* has the concept *good* as its corresponding *Tn* and the concept *red* as its corresponding *N*. Consider some provisional examples:³⁴

<u>Target Concept</u>	<u><i>Tn</i></u>	<u><i>N</i></u>
<i>cruel</i>	<i>bad</i>	<i>inducement of pain</i>
<i>murder</i>	<i>wrong</i>	<i>deliberate killing</i>
<i>courageous</i>	<i>good</i>	<i>person who confronts risk</i>
<i>gred</i>	<i>good</i>	<i>red</i>
<i>wrongday</i>	<i>wrong</i>	<i>performed on Monday</i>

Since ad hoc conjunctions and thick concepts are treated the same, this appeal to *N* and *Tn* does not yet provide the difference we're seeking. How then do thick concepts differ from ad hoc conjunctive concepts?

Here's my proposal. I want to claim that, for each thick concept *Tk*, the relevant *N* can be strengthened in a way that's compatible with *Tn* but incompatible with *Tk* itself. But nothing similar is true for ad hoc conjunctive concepts. In other words, I take the following to be necessary for a term to express a thick concept; and together with (c1) and (c2), I take it to be sufficient. For each thick term '*Tk*',

- (c3) There is an additional concept *C* such that (i) *C* is logically stronger than *N*, (ii) compatible with *Tn*, but (iii) incompatible with *Tk*.

(Recall that *Tn* is the thin concept featured within the account of *Tk*, and that *N* is a non-evaluative concept that exhausts the non-evaluative part of that account.)

Since this condition is complex, let me illustrate it with a specific example involving *courageous*. For this thick concept, the relevant *N* was said to be *person who confronts risk*. But now

³⁴ The following examples are merely provisional, but minor alterations will make little difference to what I go on to say. Dancy, for example, would most likely wish to change each *Tn* to a concept like *merit* or *appropriate*, but that wouldn't make any important difference to my main claims.

consider a further concept that is logically stronger than *person who confronts risk*—namely *person who foolishly confronts risk*. This latter concept is clearly compatible with *good*, since a foolish risk-confronter could be good for reasons completely unrelated to how she faces risk—perhaps she’s generous. Moreover, that strengthened concept is clearly incompatible with *courageous*. So, ‘courageous’ satisfies (c3).

But what’s the basic idea behind this condition? The existence of a concept *C*, as described, tells us that there’s a particular kind of case where *N* obtains but the thick concept itself does not obtain. (After all, *C* entails *N* but excludes that thick concept.) However, given that *C* entails all of that thick concept’s non-evaluative content, we’re now left to wonder why *C* excludes the thick concept in the first place. Initially we would expect that this is because *C* excludes its associated thin concept *Tn*. But in fact *C* does not exclude *Tn*. So *N* and *Tn* do not tell us why *C* excludes that thick concept. The upshot is that there must be something more, or something else, to a thick concept than merely the relevant *N* and *Tn*. Whatever else there is to a thick concept, it must exclude the obtaining of the relevant *C*.

But of course there is nothing more to an ad hoc conjunction than merely the conjunction of *N* and *Tn*. Ad hoc conjunctions do not correspond to a concept *C* as described. Take ‘gred’ for example. The concept *red mitten* is no candidate for *C*. That concept is stronger than *red* and compatible with *good*; but it’s clearly not incompatible with *gred* (i.e. with *good and red*). So, *red mitten* is no candidate for *C* since it does not satisfy clause (iii). This is just one example, but I trust the idea is fairly clear. For each ad hoc conjunction, there is no concept *C* such that *C* satisfies all of (i), (ii), and (iii).³⁵ Ad hoc conjunctions do not satisfy the condition in question.

³⁵ There’s actually a proof to show this (For ease of formulation, I’ll slide over the difference between concepts and terms). Let ‘A’ be an ad hoc conjunction expressing the conjunction of a thin concept *Tn* and a non-evaluative concept *N*. From this, it follows that

(1) For any *x*, *x* is A if and only if *x* is N and *x* is *Tn*.

Now suppose for reductio that ‘A’ satisfies condition (c3). That is, suppose there’s a concept *C* such that,

Condition (c3) is neutral on whether thick concepts are analyzable, which means that both of the major views can accept this condition. But it also respects the sense, which many theorists have had, that thick conceptual analysis is a very difficult task. More will be said about this in section III. For now, suffice it to say that condition (c3) tells us there's something more, or something else, to a thick concept than merely a non-evaluative content conjoined with a thin concept. Whatever else there is to a given thick concept, it must exclude the obtaining of the relevant *C*. Nothing similar can be said for ad hoc conjunctions. There is nothing more to an ad hoc conjunctive concept than merely a non-evaluative content conjoined with a thin concept.

But do the paradigmatic thick terms satisfy (c3)? I think the answer is 'Yes'. We've already seen an example with *courageous*, so let's consider two of the other thick concepts mentioned above—*murder* and *cruel*. Recall that the relevant N's were *deliberate killing* and *inducement of pain* respectively. So, concept *C* with regard to these non-evaluative contents might be,

<u><i>Tk</i></u>	<u><i>C</i></u>
<i>cruel</i>	<i>inducement of pain aimed at helping the recipient.</i>
<i>murder</i>	<i>self-defense killing of an aggressor in order to save one's own life.</i>

These clearly satisfy (i), since they are stronger than the relevant N's. And they also satisfy (ii), since they are compatible with the relevant thin concept mentioned in the above conditionals. For example, a surgeon may induce pain in order to help her patient, even though this might be bad because her way of doing it involves lying or breaking a promise. And a self-defense killing in order

(2) *C* is logically stronger than *N*, compatible with *Tn*, but incompatible with *A*.

From these claims, we get a contradiction. Given (2), we're entitled to assume we have a case in which

(3) *b* is *C* and *b* is *Tn*.

Moreover,

(4) For any *x*, if *x* is *C*, then *x* is *N*. [from (2)]

(5) For any *x*, if *x* is *C*, then *x* is not *A*. [from (2)]

(6) *b* is not *A*. [from (3) and (5)]

(7) *b* is *N*. [from (3) and (4)]

(8) *b* is not *Tn*. [from (1), (6), and (7)]

(9) *b* is *Tn*. [from (3)]

Since (8) and (9) contradict each other, we can reject (2). That is, we can reject the claim that the ad hoc conjunction 'A' satisfies condition (c3).

to save one's own life could be wrong because it's done in a way that causes excess suffering to the aggressor (while other less painful means were readily available). The basic idea is that the relevant thin concept can apply for reasons quite independent of the obtaining of C.

Now consider clause (iii) of condition (c3). We can see that the above mentioned C's are incompatible with the thick concept in question. If an act of inducing pain is aimed at helping the recipient (e.g., a surgical incision), then this would disqualify it from counting as cruel. And if one kills an aggressor in order to save one's own life, then this would disqualify it from counting as murder. In short, these paradigmatic thick terms appear to satisfy (c3), and I take it that the strategy I've illustrated above can be generalized to all paradigmatic thick terms.³⁶

Before concluding this section, let me make one final note. I've so far claimed that (c3) is necessary for a term to express a thick concept. And I've also claimed that (c1), (c2), and (c3) are jointly sufficient. But some ethicists might reject this sufficiency claim. They have in mind a slightly more restrictive conception of the thick. Consider some expressions we've not discussed: 'kraut', 'tart', 'good human being', and 'obligation to repay a debt of twenty-dollars'. Do these express thick concepts? Some ethicists say 'No'.³⁷ So, we might wonder whether these questionable cases present problems for my claim that (c1)-(c3) are sufficient.

³⁶ One might object that in each case I have misidentified the relevant N. Perhaps the above considerations just go to show that the negation of a part of C should be integrated into N. For example, one might claim that *inducement of pain not aimed at helping the recipient* is actually the appropriate N corresponding to *cruel*. There are two things to say in response to this. First, even if this point is correct in this particular case, it does not mean that there is no C as described by condition (c3). Rather, it only means that my particular example does not work. But perhaps other examples could be provided. And second, the objection will most likely require that we stretch the accounting-for relation too far. Recall that N is supposed to exhaust the non-evaluative portion of the account of Tk. And this objection claims that the negation of the part of C additional to N is part of that account. In many cases, however, it will seem counterintuitive that Tk is to be accounted for in terms of this expanded content. In fact, I find it dubious that *cruel* is to be accounted for in terms of *inducement of pain not aimed at helping the recipient*.

³⁷ Dancy (1995, 264) and Gibbard (2003, 300-1) suggest that thick concepts ought to be distinguished from pejorative expressions (e.g., 'kraut', 'tart', etc.). Hare (1963, 188-9), on the other hand, tends to lump them in with the thick. Andrew Payne (2005, 91) claims that 'obligation to repay twenty dollars' and 'good human being' do not count as thick.

I should note that I don't see good reason to exclude 'good human being' from the thick. Payne attempts to exclude it from the thick by claiming that the obtaining of its non-evaluative content does not "serve as the reason for" the obtaining of the evaluative content. He supports this by claiming that

This is a complicated issue, and I'm ultimately unsure whether these terms satisfy (c1)-(c3). If they do satisfy these conditions, and if these expressions are not genuinely thick, then some further conditions will need to be proposed in addition to (c1)-(c3). It's beyond the scope of this paper to delve into the matter, but I will say this much. If any further conditions need to be added so as to deal with the above questionable cases, then it's not likely to make any difference for my main thesis. That is, we're primarily interested in whether 'knowledge' expresses a thick concept, and what implications this might have. But it's quite unlikely that 'knowledge' would fail to satisfy those further conditions (whatever they might be). The reason is that 'knowledge' just seems nothing like the above questionable cases. And whatever conditions we add so as to exclude them from the thick, it's likely that these conditions would not also preclude 'knowledge'.

The more interesting question is whether 'knowledge' satisfies (c3). That's something I've not yet shown. I think it clearly does satisfy this condition, and my grounds for thinking so finally connect our discussion up with the Gettier problem.

II.5. Does 'Knowledge' Satisfy (c3)? Condition (c3) tells us that each thick concept Tk has corresponding to it a strengthened version of N that's compatible with Tn but incompatible with Tk . I will address the question of whether 'knowledge' satisfies (c3) by considering whether there's a further concept that satisfies the three clauses within that condition. As I'll show in this section, it's

A good human being is not good in virtue of being a human being, but in virtue of being a particular kind of human being (Payne (2005, 91-3)).

However, a similar claim can be made for many paradigmatic thick terms. Recall that 'murder' seems to express the evaluative content *wrong* and the non-evaluative content *deliberate killing*. Now consider an analogous claim to what Payne makes above:

A murder is not wrong in virtue of being a deliberate killing, but in virtue of being a particular kind of deliberate killing.

This seems just as plausibly true as Payne's claim. Just as there are many human beings who are not good, there are also many deliberate killings that are not wrong (e.g. killings out of self-defense). Thus, if Payne's requirement is successful in excluding 'good human being', then it seems very unclear how this won't end up excluding many paradigmatic thick terms. We should be open to the possibility that, in at least some contexts of utterance, 'good human being' does express a thick concept.

clear that there is such a concept, one that has recently come under the name ‘veritic luck.’³⁸ We’ll also see that Gettier’s important contribution was that veritic luck is compatible with justification.

Although the name ‘veritic luck’ is new to epistemology, the concept is long familiar.

Consider a dated example from A.J. Ayer, 1958:

[A] superstitious person who had inadvertently walked under a ladder might be convinced as a result that he was about to suffer some misfortune; and he might in fact be right.³⁹

Or consider three more cases from Bertrand Russell, 1948:

There is the man who looks at a clock which is not going, though he thinks it is, and who happens to look at it at the moment when it is right; this man acquires a true belief as to the time of day....

There is the man who believes, truly, that the last name of the Prime Minister in 1906 began with a B, but who believes this because he thinks that Balfour was Prime Minister then, whereas in fact it was Campbell Bannerman.

There is the lucky optimist who, having bought a ticket for a lottery, has an unshakable conviction that he will win, and, being lucky, does win.⁴⁰

These represent a larger body of cases that I’ll call cases of veritic luck. They always involve a person who forms a belief on specious or unreliable grounds, such as a hunch, a wish, an inference from a false belief, a non-veridical experience, superstition, or so forth. We’re then told that, as luck would have it, the belief turns out to be true. The epistemology literature is filled with cases like Ayer’s and Russell’s, each of which involves a belief of the same general type. I take *veritic luck* to be the concept that classifies the true beliefs within this particular set of cases.

Although I’m not wedded to any particular definition of veritic luck, here’s the common view: an agent’s belief is veritically lucky just in case it’s a matter of luck that the agent’s belief is

³⁸ The name comes from Mylan Engel (1992, 59) and has been adopted recently by Pritchard (2005, 145).

³⁹ Ayer (1958, 29).

⁴⁰ Russell (1948, 154-5).

true.⁴¹ As suggested by this definition and by the above cases, veritic luck does not include cases where it's merely a matter of luck that the proposition is true. For example, it would not include a case where I luckily win the lottery and thereby come to believe that I've won as a result of seeing the winning ticket. Rather, veritic luck is about it being lucky that I have a true belief, not merely lucky that the proposition is true.

Given this notion of veritic luck, we can now ask whether the concept *veritically lucky belief* satisfies clause (i). In other words, is the following true?

(i') The concept *veritically lucky belief* is logically stronger than the relevant *N* associated with *knowledge*.

Recall that *N* is supposed to exhaust the non-evaluative portion of the account in question. And if we take the relevant *N* to be the concept *true belief*, then (i') is clearly true. Since veritic luck is the property had by the true beliefs in the relevant cases (like Ayer's and Russell's), it's obvious that all cases of veritic luck are cases of true belief. But the opposite does not hold, since there are clearly many true beliefs that are not veritically lucky.

In a moment I will consider whether the relevant *N* might be something stronger than *true belief*, such as *safe belief* or *sensitive belief*. But for now, let's assume *true belief* is the relevant *N*. Under this assumption, it's clear that (i') is true.

Let's now skip to clause (iii) of (c3). It's quite plausible that the concept *veritically lucky belief* satisfies this as well. That is,

(iii') The concept *veritically lucky belief* is incompatible with the concept *knowledge*.

⁴¹ Pritchard (2005, 146) and Engel (1992, 67). It should be noted that Engel includes a reference to the subject's "evidential situation" in his definition, but it's not clear how this makes his view importantly different from the one I've mentioned here.

This is a form of the anti-luck platitude advanced by many epistemologists—“knowledge excludes luck”.⁴² In Ayer’s case, for example, it would not be correct to say that the superstitious person knows that misfortune will befall him. And in Russell’s cases it’s also not true that the person in question has knowledge of the relevant proposition. So, let’s take (iii’) on board.

Does this concept also satisfy clause (ii) of (c3)? In other words, is the following true?

(ii’) The concept *veritically lucky belief* is compatible with the relevant thin concept associated with *knowledge*.

Indeed it is. And this is essentially what Gettier taught us. As noted at the outset, it’s plausible that *justification* is a thin concept. So this is the thin concept associated with *knowledge*. But is *justification* compatible with veritic luck?

The novel insight of a Gettier case is that the concept *veritically lucky belief* is compatible with *justification*. Gettier showed us that a belief can be veritically lucky as well as justified. To see this, recall one of his original cases. Smith has strong evidence for the proposition *p* that Jones owns a Ford. According to Gettier,

Smith’s evidence might be that Jones has at all times in the past within Smith’s memory owned a car, and always a Ford, and that Jones has just offered Smith a ride while driving a Ford.

Now suppose that Smith has another friend, Brown, of whose whereabouts Smith is “totally ignorant.” From *p*, Smith infers *q* that either Jones owns a Ford or Brown is in Barcelona. Smith has correctly inferred *q* from a proposition for which he has “strong evidence,” and is therefore “completely justified” in believing *q*. But as it turns out, *p* is false—“Jones does not own a Ford, but is at present driving a rented car.” Moreover, “by the sheerest coincidence, and entirely unknown to

⁴² Of course, the platitude needs to be restricted. As Pritchard and Engel have each pointed out, some forms of luck are not excluded by knowledge. But they do think veritic luck is the main form that is excluded.

Smith,” Brown happens to be in Barcelona.⁴³ So, Smith has a true belief in *q*, indeed, a veritically lucky true belief. And that belief is justified. Veritic luck is therefore compatible with justification.

The case I’ve used here is representative of Gettier-style cases. The post-Gettier literature provides a swath of similar examples, though I won’t reproduce them here.⁴⁴ But I will add that, in general, I take them to be cases where the subject has a veritically lucky belief. And I take the insight of Gettier’s cases to be that veritically lucky belief is compatible with justification.

Epistemologists already knew that veritic luck was incompatible with knowledge. This was precisely the point behind Ayer’s and Russell’s examples. But Gettier showed us that adding justification to the account of knowledge is not enough to exclude veritic luck.⁴⁵ In sum, Gettier cases show us that (ii’) is true.

Thus far, I’ve argued that, if *true belief* is the relevant *N* associated with knowledge, then *veritically lucky belief* satisfies clauses (i)–(iii). But what if I’m wrong in assuming that *true belief* is the relevant *N*? There are a number of potential (though controversial) entailments of knowledge that should also be considered, ones that are stronger than *true belief*. I will discuss several of these for the remainder of this section.

Let’s first consider the concept *safe belief*. For reasons that will emerge in a moment, we should have Pritchard’s definition of safety in mind. It states,

S’s belief is safe *iff* in nearly all (if not all) near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true.⁴⁶

Defined in this way, *safe belief* is indeed stronger than *true belief*. What if it turns out that *safe belief* is the relevant *N*?

⁴³ Gettier (1963, 232–33).

⁴⁴ See Shope (1983 and 2004) for other examples.

⁴⁵ This way of construing Gettier cases is also suggested by Pritchard (2005, 148), Engel (1992, 72), and Matthias Steup (2008). Consider a quote from Steup: “The role of the justification condition is to ensure that the analysans does not mistakenly identify as knowledge a belief that is true because of epistemic luck. The lesson to be learned from the Gettier problem is that the justification condition by itself cannot ensure this.”

⁴⁶ Pritchard (2007, 283).

In this case, a slight modification of each of (i')-(iii') will work well enough for our purposes.

For example, (i') can be replaced with

(i*) The concept *veritically lucky safe belief* is logically stronger than the relevant *N* associated with *knowledge*.

A veritically lucky safe belief is a safe belief that occurs in a veritic luck case.

But one might think there's a problem with (i*). Pritchard has argued that safety excludes veritic luck. This would mean that *veritically lucky safe belief* is not a coherent concept. However, it's important that we take note of the scope of Pritchard's claim. His claim is only that, for any contingent proposition *p*, a safe belief in *p* excludes veritically lucky belief in *p*. But it's crucial that he restricts his claim to contingent propositions. After all, it's fairly obvious that there could be veritically lucky safe beliefs in necessary truths.⁴⁷ For example, I could form the belief that the square root of 15876 is 126 without knowing anything about mathematics but purely on a hunch. And as luck would have it, I've formed a true belief. Indeed, it's true in all the possible worlds in which I hold that belief. So, by Pritchard's definition of safety, my belief is safe. And indeed, it intuitively seems to be a matter of luck that I have a true belief.⁴⁸ So, my belief is a veritically lucky safe belief, which means that this concept is coherent after all. And, of course, *veritically lucky safe belief* is logically stronger than *safe belief*. Thus, (i*) will work as a feasible alternative to (i') in the event that *safe belief* is the relevant *N*.

Moreover, if *safe belief* is the relevant *N*, then the following conditions would work as alternatives to (ii') and (iii') respectively:

⁴⁷ It's worth noting that Pritchard defines luck in such a way that a belief in a necessarily true proposition could not be lucky. Although I suggest the opposite here, this is not a problem for my argument. Pritchard is working with a rather technical notion of luck, and it would need to be altered or extended in order to deal with beliefs in necessary truths. Even Pritchard grants that his account would need to be extended so as to cover this sort of luck, but he elects to leave it "to a later occasion" (2007, 280).

⁴⁸ We can get similar results with logically contingent propositions. Suppose I become really paranoid and suspicious of my parents and irrationally form the belief that I am not the biological son of Richard Kyle. And suppose that, through some strange twist of fate, my belief is true. I am actually adopted and am the biological son of someone else. Presumably, there are no nearby possible worlds in which I had different biological parents from those I actually have. So, there's no nearby possible worlds in which my belief is false. So, my belief is safe, but still veritically lucky.

(ii*) The concept *veritically lucky safe belief* is compatible with the relevant thin concept associated with *knowledge*.

(iii*) The concept *veritically lucky safe belief* is incompatible with the concept *knowledge*.

First consider (ii*). We can easily revise Gettier's case so that Smith's justified true belief is also safe. Suppose that instead of inferring q from p , he infers r —that either Jones owns a Ford or the square root of 15876 is 126. Moreover, suppose Smith has no clue about mathematics and is merely pulling the second disjunct out of thin air. Given this slightly revised Gettier case, Smith's belief in r is safe, since it's true in all possible worlds where he has that belief. And it's also veritically lucky. After all, there's nothing in this revision that should prompt us to treat it differently from the original case. Thus, (ii*) also appears to be true.

Now consider (iii*). We've already seen that *veritically lucky belief* is incompatible with knowledge. Insofar as this is so, *veritically lucky safe belief* will also be incompatible with knowledge. So, (iii*) also appears to be true. In short, (i*)-(iii*) would work as feasible alternatives to (i')-(iii'). And if *safe belief* is the relevant N , then *veritically lucky safe belief* can be the relevant C .

No doubt, the epistemological literature has provided many other concepts that could be potential substitutes for N . *Sensitive belief* and *reliable true belief* are just two examples.⁴⁹ What if one of these other concepts is the relevant N ? As we've seen just above, the framework for thinking about whether 'knowledge' satisfies (c3) stems from Gettier counterexamples. Those examples strongly suggest that 'knowledge' satisfies (c3) when we take N to be something rather minimal, such as *true belief*, or even something more robust, such as *safe belief*. But there have been other attempts at building up *true belief* by adding on further non-evaluative concepts (e.g., *sensitivity*, *statistical reliability*, etc). Without considering each of these individually, there are two general points to make.

⁴⁹ I'm assuming that defeasibility theories of knowledge (such as from Ginet (1988) and Klein (1971)) do not provide a potential substitute for N (in addition to true belief). The reason is that their fourth condition would not be wholly non-evaluative, given that these conditions incorporate justification.

First, for each of these additions to true belief, there tends to be much controversy over whether it would provide even a necessary condition for knowledge. Of all the proposed non-evaluative entailments of knowledge, *true belief* is by far the most widely accepted. Further expansions upon *true belief* tend to provoke much dispute, and those disputes are, in effect, disputes about whether such add-ons can be part of the relevant N. But second, even if some add-on is acceptable as a necessary condition on knowledge, there is no reason to think that a Gettier-style counterexample could not be drawn up. I'm not claiming that Gettier-style counterexamples are insurmountable, but I am doubtful that they can be handled merely by building up the non-evaluative content of knowledge. The history of failed attempts solidifies my doubt even further.⁵⁰

But there's one last possibility. What about the concept *true belief that's not veritically lucky*? Could this be the relevant N? I don't think so. Certainly it's true that *knowledge* necessarily entails this concept. Nevertheless, I think it is clear that *knowledge* is not to be accounted for in terms of *true belief that's not veritically lucky*.⁵¹ To be sure, many epistemologists are keen on providing an account of knowledge that excludes luck. But the standard approach is to account for *knowledge* in terms of some additional element (e.g., safety, a no-defeater condition, etc.), and then to claim that this additional element provides the basis for excluding luck. The reason for this standard approach is that it explains something that we very much want to explain—namely, the anti-luck platitude that knowledge excludes luck. But if we accounted for *knowledge* directly in terms of *being not veritically lucky*, then this direct approach would simply assert the anti-luck platitude without explaining it. Such an account would thereby defeat our whole purpose in providing an account of knowledge that

⁵⁰ See Shope (1983 and 2004) for a survey of many different attempts.

⁵¹ The view under consideration here should not be confused with Peter Unger's early view according to which a person knows p only if it's "not at all accidental that the man is right about its being the case that p" (1968, 161). Despite its initial appearance, Unger's view is actually more of a reliability type view, since this condition amounts to there being something in one's situation that guarantees, or makes it highly probable, that one wouldn't be wrong.

excludes luck. I therefore do not think it is plausible to account for *knowledge* directly in terms of *true belief that's not veritically lucky*.

What we've seen in this section is that there clearly seems to be a concept corresponding to 'knowledge' that satisfies clauses (i)-(iii) of (c3). That concept is *veritically lucky belief* (or some slight expansion upon it, such as *veritically lucky safe belief*). This means that 'knowledge' satisfies (c3) as well as the whole set of conditions that I've taken to be jointly sufficient for thickness. I conclude, then, that 'knowledge' expresses a thick concept.

The concept of knowledge has been misclassified as thin. We've seen that this central concept is actually thick. But this misclassification is not petty. As we'll see in the final section, the considerations I've put forward in favor of the thickness of *knowledge* strongly suggest some claims of importance for both epistemology and ethics.

III. A New Look at the Gettier Problem.

Given the conclusion that 'knowledge' is thick and given my arguments in favor of it, two important claims follow. We've just seen that Gettier counterexamples are at the heart of the fact that 'knowledge' satisfies (c3). I now wish to draw out some implications from (c3) that connect more closely to the Gettier problem. The first is that there is a general problem of analyzing thick concepts, and the second is that the Gettier problem is a specific instance of this more general problem. Let's look at them in turn.

III.1. The Problem of Thick Conceptual Analysis. In attempting to establish our main conclusion—that 'knowledge' expresses a thick concept—I've also provided grounds for accepting that (c3) is necessary for a term to express a thick concept. That is, we've seen reason to accept that

If a term 'Tk' expresses a thick concept, then there is an additional concept *C* that (i) is logically stronger than *N*, (ii) compatible with *Tn*, but (iii) incompatible with *Tk*.

Recall that Tn is the thin concept featured within the account of Tk , and N is a non-evaluative concept that exhausts the non-evaluative content within that account. This condition gives us a theoretically neutral way of excluding ad hoc conjunctions, like ‘gred’ and ‘wronday’, from the thick. Unlike in the case of paradigmatic thick terms, there is no concept C corresponding to ad hoc conjunctions that satisfies clauses (i), (ii), and (iii). This condition also respects the sense (which many have had) that thick conceptual analysis can be a difficult if not impossible task. Let’s now see how this plays out in more detail. I’ll argue that there’s a general problem of thick conceptual analysis, a problem that could arise for any thick concept.

If the above claim is true, then, for any thick concept Tk , there is a particular task involved in stating an analysis of Tk —the task of explaining why Tk excludes C . The basic idea is this: since there is a concept C (as described above) corresponding to each thick concept, the analysis of Tk must somehow exclude C . The task, then, is to state non-circular, necessary, and sufficient conditions that explain why Tk excludes C . This can be difficult because, even if one uncovers the relevant N and Tn , the mere conjunction of these two concepts will not do the trick (as suggested in (i) and (ii)). Thus, one must state a further condition, presumably a supplement to the mere conjunction of N and Tn , so as to generate an analysis that can explain why Tk excludes C .

It may help to see how this problem arises for a given thick concept. So let’s consider an example involving the thick ethical concept *cruel*. With regard to this thick concept, we can plausibly suppose that the relevant N has something to do with causing pain⁵² and that Tn is a thin concept like *bad* or *wrong*. Now let’s ask whether the conjunction of two such concepts is sufficient for cruelty:

For any act x , x is cruel if (i) x is performed by an agent who sees that doing x will cause someone else pain and (ii) x is bad (or wrong).

⁵² Here we can understand ‘pain’ in the broad sense—not merely a raw physical sensation, but also a broader form of distress including psychological and emotional pain.

Is this true? Not likely. It's not hard to imagine bad actions that knowingly induce another's pain, but that are not cruel. For example,

Dismissive Surgeon: Smith gave his consent to undergo surgery on the condition that his surgeon, Dr. Jones, agrees not to use a certain incision instrument X during the surgery. (Smith is superstitious and, for some strange reason, strongly insists that X not be used). Knowing that the surgery can be performed perfectly well with an alternative incision instrument, Dr. Jones promised, and even signed a contract stating, that she won't use X. However, once surgery time came around Dr. Jones didn't want to comply with Smith's strange request. So, she decided to use X during the surgery despite Smith's insistence. Of course, Smith was completely unaware of which instrument Dr. Jones used. And Dr. Jones knew that her use of X would cause pain to Smith, just as any surgical incision would. But there were no unusual complications. Smith recovered normally, enduring the pain from the incision, but knowing nothing about Dr. Jones' use of X during the surgery.

Intuitively, Dr. Jones' act of using instrument X was not cruel. However, it is true that she knew her use of X would cause pain to Smith. Moreover, Dr. Jones' use of X was quite plausibly bad (or wrong), since she broke a promise, violated a contract, and voided Smith's consent to undergo that very surgery. But then it looks like conditions (i) and (ii) both obtain, which means they are not sufficient for cruelty.⁵³

Thus, with regard to the thick concept *cruel*, the problem I'm outlining might manifest itself as the task of stating a condition, in addition to (i) and (ii), that generates an analysis of *cruelty* that can explain why helpful inducements of pain are not cruel. It is by no means obvious what this condition would be, and I shall not try to state it here. The present goal is simply to illustrate how the problem I've been describing can arise for a thick ethical concept.

Henceforth, I will refer to this general problem as the problem of Thick Conceptual Analysis (the TCA problem, for short). This is the task of stating non-circular, necessary and sufficient

⁵³ One might claim that the non-evaluative content of *cruel* (i.e. clause (i)) has been misidentified. Not merely does a cruel act knowingly induce pain, but it must be intended to induce pain. But this condition is surely too strong to be a feasible candidate for N. Children can act cruelly to one another without intending to induce pain. Their intentions might be less malicious. Perhaps they aim to attain popularity or improve self-image, while being indifferent to the pain that's induced. But this doesn't diminish the fact that their treatments of one another are cruel.

conditions for *Tk* that can explain why *Tk* excludes *C*. Judging from the literature, the additional condition that theorists are most likely to propose is one that includes a substantive explanatory relation between the relevant thin concept and the relevant non-evaluative content—for example a condition stating that *Tn* applies in virtue of the fact that *N* applies. Whether or not this amounts to an adequate response remains to be seen.⁵⁴

In this section, we’ve seen that the TCA problem arises from the fact that (c3) is required for thickness. It is therefore a general problem that can arise for any thick concept, including important ethical ones, like *cruel*.⁵⁵ Let’s now return to the thick concept of knowledge.

III.2. The Gettier Problem. Does the TCA problem arise with respect to *knowledge*? It does and it must. After all, if *knowledge* is a thick concept and if all thick concepts are subject to the TCA problem, then *knowledge* is also subject to the TCA problem. Here I’d like to show that the Gettier problem is a specific instance of the TCA problem.

We’ve already seen that the type of case Gettier proposed is at the heart of the fact that ‘knowledge’ satisfies (c3). That is, Gettier revealed that veritic luck is compatible with the thin concept associated with ‘knowledge’, namely *justification*. Still, the Gettier problem is supposed to be something more general than the particular cases Gettier himself provided. What then is the Gettier problem?

I take the Gettier problem to be the task of stating non-circular, necessary, and sufficient conditions for *knowledge* that can explain why *knowledge* excludes the kind of belief described in

⁵⁴ As I argue in my manuscript “How to Analyze a Thick Concept” this response is not adequate, though I offer a new position that I take to be quite plausible.

⁵⁵ Is the TCA problem resolvable? A resolution to this problem would likely be a general schema that shows promise in mapping out conditions for the application of each thick concept. This is a matter I won’t attempt to address here (though I do address it in my manuscript “How to Analyze a Thick Concept”).

Gettier-style cases.⁵⁶ Of course, this plainly looks like a specific version of the TCA problem—the task of stating non-circular, necessary and sufficient conditions for *Tk* that can explain why *Tk* excludes *C*.

But is this correct? Is the Gettier problem a special instance of the TCA problem? The main step in establishing an affirmative answer is to see that veritically lucky belief is the peculiar kind of belief at issue in Gettier-style cases. This is a claim made by Pritchard and strongly suggested by Mylan Engel.⁵⁷ And it's also supported by the fact that cases like Ayer's and Russell's intuitively resemble typical Gettier cases. In fact, it's common for epistemologists to treat Russell's stopped-clock example as a Gettier-style case (even though it pre-dated Gettier).⁵⁸ For these reasons, I take it to be plausible that the Gettier problem is an instance of the TCA problem.

If this is correct, then it would answer our initial question: Why does the Gettier problem arise in the first place? On my view, the Gettier problem arises because *knowledge* is thick, and a Gettier-like problem is just what we should expect from a thick concept.

Philosophers have typically understood the Gettier problem as either something peculiar to the concept of knowledge or else a symptom of a general problem about conceptual analysis. But I've steered a middle course by arguing that it's actually an instance of a problem about *thick* conceptual analysis. The Gettier problem is not something peculiar to the concept of knowledge, since other thick concepts are subject to very similar problems. Research into the Gettier problem should now take a new direction, one that inquires into the nature of thick concepts more generally.

Moreover, if my characterization of the Gettier problem is correct, then we need not explain this problem by appealing to a general difficulty about conceptual analysis. For example, some

⁵⁶ For this sort of characterization, see Sturgeon (1993, 157). Ginet (1988, 105) also classifies it this way, only without the reference to Gettier-style cases. Shope (2004, 289) mentions this as one of the two main ways the Gettier Problem has been understood.

⁵⁷ See Pritchard (2005, 148-9) and Engel (1992, 64-72)

⁵⁸ For example, see Pritchard (2005, 148) and Plantinga (1993, 33). Of course, for Russell's case to count as a Gettier-style example we must assume that the subject in this case has a justified belief. And that would be additional to what Russell himself claims.

philosophers believe that most concepts are unanalyzable on the grounds that few of them (if any) admit of uncontroversial analyses.⁵⁹ But if my argument is correct, then we need not rely on this general skepticism about conceptual analysis in order to explain why the Gettier problem arises. Indeed, for all I've claimed, the TCA problem and the Gettier problem might be resolvable, and these very important thick concepts might turn out to be analyzable.

III.3. Conclusion. Knowledge is importantly thick. I've proposed three conditions on thickness, and I've argued that *knowledge* satisfies all of them. Although theorists have classified *knowledge* as thin, their classification is importantly mistaken. As I've argued, the thickness of *knowledge* allows us to explain why the analysis of that concept is troubled by the Gettier problem.

Along the way, we saw that thick concepts differ importantly from ad hoc conjunctions, like *good and red*. And I proposed a theoretically neutral condition that explains this difference, a condition that connected our discussion with the Gettier problem. Gettier's counterexamples are at the heart of the fact that *knowledge* satisfies this condition. This secured my case for the conclusion that *knowledge* is thick, but it also provided the basis for two related implications.

The first implication was that all thick concepts are subject to a problem of thick conceptual analysis (the TCA problem). And the second was that the Gettier problem is a specific instance of this more general problem. These allowed us to explain why the Gettier problem arises in the first place, and they also revealed that very similar problems would arise for other thick concepts, including ethical ones like *cruelty*. The Gettier problem is widely thought to be an incredibly difficult problem. It now turns out that more of us have this problem than one might initially think.

⁵⁹ For a discussion on this, see Laurence and Margolis (1999, 14-16).

3. How to Analyze a Thick Concept

I. Introduction.

It is becoming more and more common for ethicists to distinguish between two types of evaluative concepts—the thin and the thick. Thin concepts typically include *good*, *bad*, *right*, and *wrong*.

Paradigmatic thick concepts include *courageous*, *murder*, *lewd*, *rude*, *generous*, *cruel*, *loyal*, and many more.

As ethicists often point out, it is distinctive of these thick concepts that they “somehow hold together” an evaluation plus a description.¹ If I call you ‘courageous’, then I’ve not only evaluated you, but I’ve also described you in a specific way (e.g. as being one who faces risk). Other expressions, like ‘good’ and ‘red’, do not typically exhibit this feature. So, it is initially puzzling how both evaluation and description can be associated with a single concept. Just how does a thick concept hold together an evaluation plus a description? This is our main question.

Our main question presents a challenge to those who sit on either side of an important divide in ethics. On one side of that divide, there are Separatists—those who believe that thick concepts can be analyzed into evaluative contents and non-evaluative contents. On the other side, there are Non-Separatists—those who believe thick concepts cannot be analyzed at all.² As the title of this paper suggests, my sympathies lie with Separatism. Separatists do not close off any avenues for answering our main question. They have all the Non-Separatist’s resources for answering it,³ but

¹ For example, see Dancy’s opening statement (1995, 263).

² For examples of Non-Separatism, see Platts (1979), Lovibond (1983), and Dancy (1995). Elgin (2005) also interprets Bernard Williams as holding this view. As I construe this view, their denial is specifically about thick conceptual analysis, not conceptual analysis in general. I shall have nothing to say here about general skepticism regarding the possibility of conceptual analysis. For people who accept the analyzability of thick concepts, see Tappolet (2004), Payne (2005), Burton (1992), Gibbard (1992), and Hurka and Elstein (2009). As I’ve laid out the distinction, the two positions of Separatism and Non-Separatism are not exhaustive, contrary to what the labels suggest. It is possible to hold, as Frank Jackson does (1998), that thick concepts can be analyzed, but not into distinct types of descriptive and evaluative contents. However, since Jackson does distinguish between descriptive and evaluative terms, the Separatist view that I go on to propose can likely be framed in a way that would be amenable to Jackson’s view.

³ For example, two main avenues have been vaguely associated with Non-Separatism. First, there’s the option of appropriating a speech act view in order to answer our main question. On this view, the deployment of thick concepts, in the right circumstances, counts simultaneously as a speech act of evaluating and as describing. Alan

they also have the option of regarding evaluation and description as two components within the analysis of a thick concept—one component is its evaluative content, while the other is its non-evaluative content. This is by no means an argument for Separatism; it is merely a reason for taking Separatism as a starting point of inquiry into our main question. The issue of how well it compares to Non-Separatism will be addressed at the end of our search.

The overall goal of this paper is to support a particular form of Separatism by proposing and defending a schema for analyzing thick concepts (what I call the Qua Schema). This schema is intended to be general enough so that virtually any thick concept could in principle be analyzed in accordance with it.⁴ For contrastive purposes, the discussion of the Qua Schema (in section III) will be situated between two discussions of rival schemata (in sections II and IV). The goal of sections II-IV is to explain the Qua Schema and to show that alternative schemata have troubles answering our main question that the Qua Schema can overcome. Later, in section V, I supply my positive argument in favor of the Qua Schema. There, I argue that there is good reason to accept the Qua Schema, and, as I point out in section VI, even the Non-Separatist has good reason to accept it.

II. The Conjunctive Schema.

II.1. Conjunctive Schema Explained. Merely appealing to Separatism does not by itself answer our main question. Separatism only tells us that thick concepts can be analyzed in terms of evaluative and non-evaluative contents, but it does not explain how these two contents relate to one another.

Thomas vaguely suggests this sort of view in a footnote (2008, 347), though no one to my knowledge has developed it. Or second, there's Dancy's view (1995, 276-77). This will be discussed in the final section.

⁴ Two clarifications about this seemingly strong thesis are in order: First off, the 'in principle' is important. It might be that we don't have the vocabulary to provide analyses for each thick concept. For example, we may not have a term to express the concept itself, or we may not have the vocabulary to state its evaluative or non-evaluative contents. But second, it's important to note that what I count as a thick concept might be somewhat more restrictive than what other theorists will count. For instance, I do not consider slurs, epithets, and derogatory words as expressing thick concepts. This restrictive use is in-line with other theorists, such as Jonathan Dancy (1995, 264) and Allan Gibbard (2003, 300ff), who are hesitant to count these terms as thick.

Could there be a very simple way of spelling out this relation? Perhaps these contents are combined by mere conjunction. This brings us to the first alternative account—The Conjunctive Schema:

Conjunctive Schema:

For any x and any thick concept T , x falls under T iff x is N and x is E .⁵

In general, I will assume that there's a non-evaluative expression that is substitutable for 'is N '. This expression should be seen as one that exhausts the non-evaluative content of T . Similarly, 'is E ' should be seen as an expression that exhausts the evaluative content of T .⁶

I think it makes an important difference as to what we use for an evaluative expression. And this is something I will pay great attention to when I construct my positive account. But in general, whenever the variable ' E ' appears in this paper, let it stand in for a paradigmatic thin term (e.g. 'good', 'bad', 'right', 'wrong', 'better', 'worse', 'best', 'worst', 'ought', etc.). That thick concepts are analyzable in thin terms is a common assumption accepted by most Separatists,⁷ although, as we'll see, my positive account (i.e. the Qua Schema) diverges from this view in important ways.

II.2. Conjunctive Schema Rejected. Does the Conjunctive Schema provide a plausible answer to our main question? I think the answer is 'No'. The reason is that mere conjunction allows for both contents to obtain coincidentally, but thick concepts do not tolerate such coincidence. Let me illustrate this by considering a conjunctive analysis of the thick concept *cruel*. According to the Conjunctive Schema, the analysis of *cruel* would go something like this:

⁵ This sort of schema is sometimes attributed to traditional non-cognitivists, such as R.M. Hare and C.L. Stevenson. Such non-cognitivists might be unhappy with the way I've stated the schema since it implicitly assumes that ' a is E ' is truth conditional. Their preferred statement of the conjunctive schema (as well as the schemas below) might be slightly different. For example, they might state the conjunctive schema as follows: ' a is T ' just means ' a is N and a is E ' (where ' T ' is a thick term rather than a thick concept). Alternatively, the following might also be suitable: For any x and any thick concept T , x falls under T iff x is N and it is correct to apply ' E ' to x .

⁶ One qualification here: The terms we use to express the relation between the evaluative and non-evaluative content will typically be non-evaluative terms (e.g. 'and', 'only if', 'on account of', etc.). Call these the connective expressions. So, to be more precise, I should say the expression substituted for 'is N ' exhausts the non-evaluative content of T minus the content expressed by the relevant connective expression.

⁷ See Burton (1992), Payne (2005), and Hurka and Elstein (2009). To my knowledge, Tappolet (2004) is the only Separatist who rejects this assumption.

For any act x , x is cruel iff (i) x is performed by an agent who knows that doing x induces someone else's pain and (ii) x is bad (or wrong).

Assuming this provides an adequate portrayal of a conjunctive analysis of *cruel*, we can easily see a problem. It's not hard to imagine acts that knowingly induce another's pain, but which are bad (or wrong) for completely unrelated reasons. Take the following case:

Dismissive Surgeon: Smith gave his consent to undergo surgery on the condition that his surgeon, Dr. Jones, agrees not to use a certain incision instrument X during the surgery. (Smith is superstitious and, for some strange reason, strongly insists that X not be used). Knowing that the surgery can be performed perfectly well with an alternative instrument, Dr. Jones promised, and even signed a contract stating, that she won't use X. However, once surgery time came around Dr. Jones plainly did not want to comply with Smith's strange request; she decided to use X during the surgery despite Smith's insistence. Of course, Smith was completely unaware of which instrument Dr. Jones used. And Dr. Jones knew that her use of X would cause pain to Smith, just as any surgical incision would. But there were no unusual complications. Smith recovered normally, enduring the pain from the incision, but knowing nothing about Dr. Jones' use of X during the surgery.

Intuitively, Dr. Jones' act of using instrument X was not cruel; after all, she was only performing a routine surgery. However, she did know that her use of X would cause pain to Smith. Moreover, her use of X was bad (or wrong), since she broke a promise, violated a contract, and voided Smith's consent to undergo that very surgery. But, then it looks like conditions (i) and (ii) both obtain, which means the analysis mapped out by the Conjunctive Schema is mistaken. Thus, *cruel* appears not to satisfy the Conjunctive Schema, and I take it that this type of problem is quite general.

One might claim that the non-evaluative content of *cruel* (i.e. clause (i)) has been misidentified. Not merely does a cruel act knowingly induce pain, but it must be intended to induce pain. But this condition is surely too strong. Children often act cruelly to one another without intending to induce pain. Their intentions might be less malicious; perhaps they aim to attain popularity or improve self-image, while being indifferent to the pain caused to others. But this doesn't diminish the fact that their treatments of one another are cruel.

Perhaps then the non-evaluative condition should include a statement to the effect that the agent in question lacks the aim of helping the pained recipient. But this response does not handle the Dismissive Surgeon case. We need not imagine that Dr. Jones was aiming to help Smith in the first place. Perhaps Dr. Jones has no aims whatsoever regarding Smith's well-being; she may only aim to do her job without getting fired. If this were the case, it might diminish our view of Dr. Jones' character, but it clearly would not make her use of X an act of cruelty.

In short, what we've seen here is that at least one paradigmatic thick concept, namely *cruel*, appears not to satisfy the Conjunctive Schema. And the kind of problem that has emerged looks like it will emerge for just about any thick concept. This schema therefore does not answer our main question. It does not provide a general answer to the question of how thick concepts combine evaluative and non-evaluative contents. The Conjunctive Schema is very likely to fall prey to cases of coincidence (like that of the Dismissive Surgeon) because it merely incorporates the relation of conjunction while limiting the substitutable evaluative expressions 'E' to the paradigmatic thin terms.

The above considerations present what I'll call the linkage problem. This is basically a special instance of our main question: How does a thick concept "hold together" an evaluation plus a description? The only difference is that the linkage problem is specific to the conceptual analyses of thick concepts. It is the problem of formulating a schema that looks capable of mapping out necessary and sufficient conditions for thick concepts generally, and that incorporates an appropriate relation between their evaluative and non-evaluative contents. The failure of the Conjunctive Schema teaches us something about what "appropriate" means here: the appropriate relation must give us confidence that cases of coincidence (like that of the Dismissive Surgeon) are not likely to arise. So, the most obvious next step is to seek a relation above and beyond mere conjunction.

II.3. An Explanatory Connection? Many theorists have been aware of the need to state a substantive relation between a thick concept's evaluative and non-evaluative contents. The most common view is an explanatory view according to which the thick concept's evaluative content obtains in virtue of the obtaining of its non-evaluative content. I will discuss this explanatory view later on, but it's worth being upfront about an important question. This view incorporates an explanation of the obtaining evaluative content. But why should we expect this explanation to appear within the analysis of the thick concept? This is not typically the case for other types of concepts. An adequate analysis of the concept *vixen* need not say what it is in virtue of which the animal in question is a fox. So, why would we expect a satisfactory analysis of *cruelty* to tell us what it is in virtue of which the act is bad?

It might be thought that the linkage problem forces an explanatory account onto us. But this is mistaken. There are other ways of forging a substantive connection between evaluative and non-evaluative content. Our language is rich with complex evaluative expressions that are partly composed of non-evaluative terms. Consider some complex evaluative expressions in which 'good' occurs as a constituent: 'good at dancing', 'good for Bob', 'good with children', 'good as a coach', 'good to Sue', 'good qua teacher', and so on. My account makes use of these resources in order to forge the needed connection between the evaluative and non-evaluative contents of a thick concept. But the connection is semantic, not explanatory. And the answer as to what explains the obtaining of the evaluative content need not appear anywhere within the analysis.

III. The Qua Schema.

III.1. Qua Schema Explained. My preferred view is what I call "the Qua Schema." On this view, certain restrictions are placed on what type of evaluative expression is to be featured within the account:

The Qua Schema:

For any x and any positive thick concept T :

x falls under T iff x is N and x is good qua N_{SUB} .

For any x and any negative thick concept T :

x falls under T iff x is N and x is bad qua N_{SUB} .⁸

The variable ‘ E ’, which stands in for a paradigmatic thin term, does not appear anywhere in this schema. The relevant evaluative expressions are instead ‘good qua N_{SUB} ’ and ‘bad qua N_{SUB} ’.

What do these formulae represent? Let’s begin with N_{SUB} . Basically, N_{SUB} depends on what N is. So far, we’ve been taking N to be the totality of the thick concept’s non-evaluative content. By contrast, I take N_{SUB} to be a content that is conceptually entailed by N . So, for example, if we suppose that *deliberate killing* is the total non-evaluative content associated with *murder*, then *cause of death* might be the relevant N_{SUB} associated with *murder*. It’s worth noting that the account allows for the possibility that N and N_{SUB} are identical, although in many cases they will be distinct.

Here are some provisional examples for what the relevant N_{SUB} might be for a given thick concept:

⁸ Let me provide two caveats that won’t make much difference for the present discussion: First, this basic account should probably be modified slightly in order to accommodate the difference between gradable and non-gradable thick concepts. Some thick concepts, like *cruel*, *generous*, and *munificent*, are gradable in that they allow for their items to be more or less cruel, generous, munificent, etc. Others, like *murder* and *knows*, are not gradable. This difference can be accounted for as a difference in the degree associated with the expression ‘good [bad] qua N_{SUB} ’. For non-gradable thick concepts, the evaluative clause can state “there is **a particular degree** d such that x is good [bad] qua N_{SUB} to **at least** degree d .” Among gradable thick concepts, however, there is some variation. For some of them, such as *generous*, the evaluative clause can state “there is **some degree or other** d such that x is good qua N_{SUB} to degree d ”. For other gradable thick concepts, such as *munificent*, the evaluative clause could state “there is **some high degree or other** d such that x is good qua N_{SUB} to degree d ”. The basic point here is that the Qua Schema can take on a couple of different forms depending on how we interpret the degree quantifier.

A second caveat: There’s a qualification stemming from the possibility of “multi-valent” thick concepts, ones that have both positive and negative evaluative content. Two potential candidates might be *kinky* and *quirky*. I am not convinced by any of the linguistic evidence that there are any multi-valent thick concepts. Still, it wouldn’t be preposterous if there were such concepts, and my account should be able to handle them if there were. One way for the Qua Schema to account for them would proceed along the lines of dividing up the non-evaluative contents into at least two distinct contents N_{SUB1} and N_{SUB2} and incorporating both of the evaluative terms ‘bad qua N_{SUB1} ’ and ‘good qua N_{SUB2} ’. The account would go something like this:

For any x and any multi-valent thick concept T , x falls under T iff x is an N and x is good qua N_{SUB1} and x is bad qua N_{SUB2} .

Generous: N_{SUB} = giver
Courageous: N_{SUB} = risk-confronter
Loyal: N_{SUB} = commitment-keeper
Honest: N_{SUB} = assertion
Murder: N_{SUB} = cause of death

Rude: N_{SUB} = social disregard
Cruel: N_{SUB} = inducement of pain
Perverse: N_{SUB} = abnormality
Lewd: N_{SUB} = sexual display
Chaste: N_{SUB} = sexual-abstainer

These are merely provisional examples; serious conceptual analysis might yield different results.

Generally speaking, what does ‘good qua K’ mean? To say that something is good qua K is, roughly, to say that it’s a K and it is good as such, or alternatively, that it is good as a K.⁹ What these paraphrases suggest (and what most dictionary entries will attest) is that ‘qua K’ is supposed to mean ‘in the capacity of being a K’. Thus, when we say that something is good qua K, we mean roughly that it is good in the capacity of being a K. This, I believe, is what Judith Thomson has correctly latched onto when she characterizes being good qua K as “being a model, exemplar, paradigm, or good specimen of a K.”¹⁰ The basic idea behind this exemplar-based account is that the property of being good qua K applies to all and only those Ks that are worthy of being copied, imitated, or modeled after in the relevant K respects.¹¹ Conversely, the property of being bad qua K applies to all and only those Ks that are inappropriate to copy, imitate, or model after in the relevant K respects. This exemplar-based account provides the basis for understanding the evaluative contents that are featured with the Qua Schema.

Let me now illustrate the kind of analysis this schema might provide by focusing on a provisional analysis of the concept *murder*:

For any act x, x is a murder iff x is a deliberate killing and x is bad qua cause of death.

⁹ We should guard against one deviant usage of ‘good qua K’. The locution ‘good qua K’ is not to be read as ‘good in that it’s a K’ where being a K is taken to be a good-making property. If something is good qua toaster, we wouldn’t normally want to say that it’s good in that it’s a toaster, nor that being a toaster is a good-making property. Moreover, if something is good in that it promotes overall happiness, where promoting overall happiness is taken to be a good-making property, we don’t normally mean that the item is good qua promotion of overall happiness. So, we must guard against this usage.

¹⁰ Thomson (2008, 19).

¹¹ For more detail on this exemplar-based account, see my paper “Knowledge as Exemplary Belief.” As I note there, the exemplar-based account of being good-qua-K should not be seen as a conceptual analysis of *good-qua-K*, since the account quickly becomes circular once the needed detail is provided.

This is merely a provisional example but it illustrates the kind of simple analysis that can be achieved under the Qua Schema. There are two initial problems that may arise for this analysis, as well as other instances of the Qua Schema. So let's consider these two problems before turning to the advantages of the Qua Schema.

The first potential problem is that it is often unclear as to what will count as being good/bad qua N_{SUB} . Focusing on the above analysis, it might be claimed that it is unclear what kinds of events will count as bad qua cause of death. Admittedly, there is some ambiguity or context-sensitivity involved with this expression. But, upon reflection, I think it's easy to see that there is a determinate meaning attached to it. More precisely, there are cases in which this expression clearly would apply, and others in which it clearly would not apply. For example, suppose that at a ripe-old age Smith dies peacefully in his sleep of "natural causes." Depending on how we spell out the details, it clearly could be true in this case that Smith's cause of death is not bad qua cause of death. Alternatively, suppose that a physician killed Smith in his sleep (e.g. via lethal injection) without Smith's consent, but that Smith's death was just as peaceful as in the earlier case. This second case is notably different from the first. In particular, it seems true that Smith's cause of death in this second case is bad qua cause of death. Thus, there appear to be clear cases in which the relevant expression will (or will not) apply. Although these examples do not resolve all worries one might have with the above analysis,¹² they do serve to allay the initial worry. Even if the expressions 'good qua N_{SUB} ' and 'bad qua N_{SUB} ' seem initially to be unclear, they can be seen as having determinate meanings once the right cases are considered.

A second objection may state that, in some cases, the appropriate reading of 'good qua N_{SUB} ' cannot be obtained unless one already has an antecedent grasp of the target concept. For example,

¹² This is far from a full-blown defense of the above analysis of *murder*. The point here is merely to illustrate one instance of the Qua Schema, and to discuss some initial problems that may arise.

it's not clear that one can grasp the appropriate reading of 'bad qua cause of death' without first grasping the target concept *murder*. But then it looks like the Qua Schema gets the conceptual priority relations all wrong. In particular, *bad qua cause of death* does not appear to be conceptually prior to *murder*, contrary to what is suggested by the Qua Schema. In response, let me first note that there's no obvious reason why the Qua Schema should be committed to the view that *good/bad qua* N_{SUB} is conceptually prior to the relevant thick concept. At least in some cases, the Qua Schema allows for the possibility that *good/bad qua* N_{SUB} is identical to (and therefore not prior to) the relevant thick concept. But, second of all, even if we grant this priority claim, we can see that the above objection overlooks an important distinction between (i) failing to grasp a concept and (ii) failing to grasp that a term expresses that concept. And the above objection mistakenly presupposes that one's failure to grasp that 'bad qua cause of death' expresses the appropriate concept entails that one fails to grasp the concept in question. But this is clearly mistaken. By analogy, I may fail to know that the German word 'gelb' expresses the concept *yellow*, but this obviously does not entail that I fail to grasp the concept *yellow*. Thus, I think the objection holds no sway. In general, I see no reason to doubt that people who fully grasp the concept *murder* will also grasp the concept *bad qua cause of death*. But grasping the latter concept doesn't guarantee that they'll always know which concept is expressed by the locution 'bad qua cause of death'.

III.2. Linkage Problem Revisited. Let's now see how the Qua Schema deals with the linkage problem. Adequately addressing the linkage problem requires stating a schema that incorporates the appropriate relation between a thick concept's evaluative and non-evaluative content. I said that the appropriate relation must give us confidence that cases of coincidence are not likely to arise, and that, in order to accomplish this, we at least need a relation above and beyond mere conjunction.

As we've seen, the Qua Schema incorporates a substantive relation between T's evaluative and non-evaluative contents. On the Qua Schema the evaluative content is *good qua* N_{SUB} . And if something is good qua N_{SUB} , then this guarantees that at least some feature of T's non-evaluative content obtains—namely N_{SUB} . And the connection goes in the other direction as well: T's non-evaluative content conceptually entails something, N_{SUB} , which is expressed by (and entailed by) our evaluative expression 'good qua N_{SUB} '. So, we have a substantive relation above and beyond mere conjunction. On the Qua Schema, the evaluative and non-evaluative contents are “intertwined” (albeit not in the strong sense that the Non-Separatist proposes).

Does this relation give us confidence that cases of coincidence will not arise for the Qua Schema? Let's first look at this issue more generally, and then we'll focus specifically on the Dismissive Surgeon case.

Generally speaking, cases of coincidence are likely to arise for the Conjunctive Schema because (i) it only incorporates the relation of conjunction and (ii) it limits the substitutable evaluative expressions to the class of paradigmatic thin terms, terms that are very general and can apply for a wide variety of different reasons. The explanatory accounts that we will examine in the next section reject (i). The Qua Schema is unique in that it rejects (ii). The Qua Schema limits the relevant evaluative expressions to the ones that satisfy the formulas 'good qua N_{SUB} ' and 'bad qua N_{SUB} '. These are not thin terms, and they only apply for very specific reasons having to do with what N_{SUB} is. So, their range of application is much narrower than that of paradigmatic thin terms. And their application always has something to do with the obtaining of the relevant non-evaluative content. How does this feature come to bear on the Dismissive Surgeon case?

The Qua Schema provides a plausible way of diagnosing that case. Let's reflect on why we are reluctant to apply the concept *cruel* to Dr. Jones' action. It's pretty clear that her act of using instrument X is bad in some way or other. After all, in using that instrument she is breaking a

promise, violating a contract, and voiding Smith's consent to undergo that very surgery. Still, merely being bad in some way or other is not enough for an act to count as cruel. Intuitively, the badness of a cruel act must have something to do with its causing pain. And the badness of Dr. Jones' action really has nothing to do with this.¹³ With this in mind, I would provide the following diagnosis of the Dismissive Surgeon case: Although Dr. Jones' act is an inducement of pain, it's not bad *as such*. That is, it's not bad as an inducement of pain. Or, equivalently, it's not bad qua inducement of pain.¹⁴ This, I think, is why Dr. Jones' action does not count as cruel.

An action can be bad in a number of different ways (as was Dr. Jones' action) without being bad qua inducement of pain. It is perfectly coherent to say that June is bad qua piano teacher, but not bad qua pianist. A similar thing is true for Dr. Jones' action. This action was bad in a number of different ways (e.g. qua promise-breaking) though it was not bad qua inducement of pain.

This intuitive diagnosis falls directly out of the Qua Schema, since 'bad qua inducement of pain' is an instance of 'bad qua N_{SUB}'. The concept *inducement of pain* seems a good candidate for (at least part of) the non-evaluative content of *cruel*. So this diagnosis is one that falls directly out of the Qua Schema.

Thus, it looks like the Qua Schema can deal with the linkage problem. It involves a substantive relation, above and beyond mere conjunction, and it gives us confidence that cases of coincidence will not likely arise. Thus, the linkage problem does not force us into employing an explanatory connection within our schema. The semantic connection of the Qua Schema will suit just fine. Of course, there are alternative schemata that do employ an explanatory connection.

¹³ This can be seen if we consider variant cases. For example, imagine that Dr. Jones decided not to use X and instead used an alternative instrument. In such a variant it is hard to see her as acting badly at all. Alternatively, imagine that Dr. Jones used instrument X but, miraculously, was able to perform the surgery without inducing any pain whatsoever. In this variant, Dr. Jones would be acting badly in the same way as the original, though without inducing any pain.

¹⁴ Note that I'm not denying the obvious claim that being an inducement of pain is a bad-making property. This is the deviant usage noted in footnote 9.

Indeed, this is the most common view. But, as we'll see, these accounts either fail to deal with the linkage problem, or they fall prey to an additional problem of unwanted emptiness.

IV. Explanatory Accounts.

IV.1. Connective Schema Explained. According to the Connective Schema, there is a substantive relation between the evaluative and non-evaluative contents of a thick concept. The obtaining of the evaluative content is explained by, or depends on, the obtaining of the non-evaluative content. For example, some theorists require that *a*'s being E holds "in virtue of" or "on account of" its being N.¹⁵

Although such a relation is very often employed, no one within this context has stated in detail what the relation is supposed to involve. In other contexts, this sort of relation can involve a universally generalized conditional of the form "For any y, if y is N, then y is E." There is reason to think that this is (at least) what theorists in the thick concept literature have in mind. In particular, this sort of conditional is often implicitly assumed to be a necessary condition for the application of some individual thick concepts, such as *sexual perversion*.¹⁶ So, for present purposes, I will incorporate such a conditional within the Connective Schema:¹⁷

Connective Schema:

For any x and any thick concept T, x falls under T iff x is N and for any y, if y is N, then y is E.

¹⁵ For example, Andrew Payne (2005) oscillates between 'in virtue of' and 'is a reason for'. Simon Blackburn (1984, 148) provides the 'on account of' relation. See also Hurka and Elstein (2009).

¹⁶ See Priest (1997), Primoratz (1997), and Williams (2003)

¹⁷ Here we need not understand the conditional as being conceptually or necessarily true. Moreover, this conditional is probably not strong enough to characterize these dependence relations (e.g. "in virtue of", "on account of", etc). But this won't make any difference for the criticism I go on to advance, since I shall be arguing that any schema that incorporates such a bridge principle is too strong as it is.

Let's call the embedded conditional a "bridge principle" since it bridges the evaluative and non-evaluative contents. By requiring such a bridge principle, the Connective Schema shows some promise in resolving the linkage problem.

IV.2. Connective Schema Rejected. But a very different problem arises for this schema—the problem of emptiness. It is often suggested that thick concepts employed within "hypertraditional" societies might turn out to be empty.¹⁸ But the problem I have in mind is much worse. If the Connective Schema is correct, then the worry about emptiness is also likely to arise with regard to the thick concepts employed by "modern" or reflective societies. And some of these thick concepts, like *murder*, are clearly not empty. So, under the Connective Schema it looks like we'll get a lot of unwanted emptiness. How does this problem arise in detail?

It looks like *murder* will probably turn out empty under the Connective Schema.

Traditionally, the concept *murder* is said to involve a particular non-evaluative content—*deliberate human killing*¹⁹—as well as a specific sort of moral evaluative content—*wrong*.²⁰ Theorists are not explicit about how these contents are related, but let's ask whether they're related in the way suggested by the Connective Schema. Applying that schema to the above contents, the following would be the bridge principle associated with the concept *murder*:

- (i) For any y, if y is a deliberate human killing, then y is wrong.

¹⁸ The term "hypertraditional" comes from Williams (1985, 142). He uses it to refer to societies that are "maximally homogenous and minimally given to general reflection." Gibbard (2003, 298-99) provides a portrayal of the problem of emptiness regarding hypertraditional thick concepts. There, he seems willing to accept the claim that *chastity* is empty.

¹⁹ See Suter (1973, 361) and Anscombe (1979, 74). Priest also mentions a fairly intuitive definition of murder, incorporating the evaluative content *morally wrong* and the non-evaluative content *human killing* (1997, 360).

²⁰ Anscombe, to my knowledge, is the only person who rejects the entailment from murder to wrongness. See (1979, 73-4). Even if she's right, there is bound to be some sort of moral badness associated with murder, and my argument below could proceed similarly.

But if (i) is necessary for the application of the concept *murder*, then that concept is surely empty. As any moral philosopher will point out, there is good reason to think (i) is false. There are many possible (indeed actual) cases of deliberate human killing that do not seem to be wrong. As just one example, consider killings out of self-defense where the only way the agent can save her own life is to kill the attacker. Suppose we agree there is at least one case involving a permissible deliberate human killing. Then, given how we've specified the content of *murder*, this concept will turn out to be empty under the Connective Schema.

In response, the proponent of the Connective Schema needs to claim that the traditional way of specifying the contents of *murder* is mistaken, and she'll need to show that its contents can be specified in a way that allows us to avoid the above consequence. The possibility of such a defense cannot be dismissed outright, but it is worth spending some time to see that the problem I'm pointing to cannot be easily fixed. As we'll see, it's not at all clear that the problem of unwanted emptiness can be avoided by the Connective Schema.

Such a defense would involve stating a bridge principle that fulfills two distinct requirements: (a) the bridge principle can have no counterexamples and (b) it must satisfy the Connective Schema. But, as I shall indicate below, it is very doubtful that there is any such bridge principle. It is very easy for a bridge principle to fail to satisfy at least one of these requirements.

Let's first consider (a). For the proponent of the Connective Schema to meet the present challenge, she will need to provide a non-evaluative expression '*N**' (not coextensive with 'deliberate human killing') such that '*a* is murder' conceptually entails '*a* is *N**'. However, for any case in which this conceptual entailment is plausible, it will be doubtful that our resulting bridge principle avoids counterexamples. For example, suppose we build up our initial non-evaluative expression so as to achieve something more robust—'deliberate human killing not aimed at saving the agent's own life'.

While it's initially somewhat feasible that '*a* is a murder' conceptually entails '*a* is a deliberate human killing not aimed at saving the agent's own life', this specification results in a false bridge principle:

- (ii) For any *y*, if *y* is a deliberate human killing not aimed at saving the agent's own life, then *y* is wrong.

Although this principle plausibly satisfies the Connective Schema, it is likely to be falsified by at least one counterexample. Consider a stockpile of potential counterexamples—active euthanasia, craniotomies, killings aimed at saving the lives of innocent people, just war killings, capital punishment, and so on. This revision therefore shows little promise of helping us avoid the emptiness of *murder*.

Regarding (b), I want to suggest that, even if we were to uncover a bridge principle for which there are no counterexamples, there's no reason to assume that such a principle would satisfy the Connective Schema. To see why, let's suppose that, upon reasoning critically about the morality of killing, we eventually uncover a suitable concept *F* such that *F* does not apply to any permissible, deliberate, human, killings (hereafter PDHK's).²¹ And suppose further that we express this concept within the antecedent of our bridge principle as follows:

- (iii) For any *y*, if *y* is a deliberate human killing and *y* is an *F*, then *y* is wrong.

There are no counterexamples to this bridge principle, since *ex hypothesi* it is falsified by no PDHK's. However, in this case there is no reason to think (iii) satisfies the Connective Schema. That schema requires that the antecedent of our bridge principle expresses all and only the non-evaluative content of *murder*. But, in advance of knowing what *F* is, we cannot assume that the antecedent of (iii) meets this requirement.

²¹ The process of uncovering concept *F* might be something like what Frank Jackson has in mind with "mature folk morality", the state of our ordinary moral conceptions after having been "exposed to debate and critical reflection (or would end up, should we keep at it consistently and not become extinct too soon)." See Jackson (1998, 133).

Let me point out two ways in which the antecedent of (iii) would fail to express all and only the non-evaluative content of *murder*. First, the concept *F* might be partly or wholly evaluative. It might be that, in order to formulate a bridge principle that is not falsified by any PDHK's, we have to incorporate certain evaluative terms (e.g. 'innocent person') within the antecedent. In this case, our specification of *F* would be at least partly evaluative, and (iii) would therefore not satisfy the Connective Schema.²²

Or second, even if we suppose that *F* turns out to be wholly non-evaluative, this does not mean that it's even remotely plausible that '*a* is a murder' conceptually entails '*a* is *F*'. For instance, suppose '*F*' just is the complicated expression 'aimed neither at saving a human's life nor at reducing a human's unnecessary pain'. In this case, we've supplied a wholly non-evaluative specification of *F* which applies to very few (if any) PDHK's. But here it's very implausible that '*a* is murder' conceptually entails '*a* is aimed neither at saving a human's life nor at reducing a human's unnecessary pain'. Not only is it dubious that these two claims are conceptually linked, but it's also implausible that there's an entailment (of any sort) from the former to the latter.²³ Of course, it is no surprise that this sort of maneuver should fail, since we uncovered concept *F* by doing moral reasoning about the relationship between deliberate human killing and wrongness. But there is absolutely no reason to think that this by itself should shed any light on what murder is. We could reason in this way without employing, or even possessing, the concept *murder*. In short, the above

²² Hurka and Elstein (2009, 526) have recently proposed a three-part schema that allows there to be evaluative contents intermingled within the non-evaluative part of the concept. An initial worry is that this view will falter on answering our main question, unless it is explained in general terms how these evaluative and non-evaluative contents are intermingled. Such an explanation is certainly not impossible, but, due to space constraints, a full discussion of what forms it would take must be reserved for another occasion.

²³ That '*a* is murder' does not entail '*a* is aimed neither at saving a human's life nor at reducing a human's unnecessary pain' can be shown by the following case: Suppose you are a hit-man working alongside your partner Sal. Sal is hired to kill an innocent person, Jones, and you are along from the ride. It turns out that the only way Jones can defend himself from Sal is to kill him, which is exactly what Jones tries to do. But, in order to save Sal's life, you decide to kill Jones. Since Jones is innocent in this case (and merely trying to save his own life) it is very plausible that you murdered him, even though your action was aimed at saving Sal's life.

two types of cases show that we cannot assume that arriving at a true bridge principle about the relationship between deliberate human killing and wrongness will give us a bridge principle that satisfies the Connective Schema.

It is a serious challenge to the proponent of the Connective Schema to state the needed bridge principle. It's not at all clear what that bridge principle would be. There are many ways in which a bridge principle might be either false or unable to satisfy the Connective Schema. In short, *murder* is a paradigmatic thick concept that is obviously not empty. But it's not at all clear that the proponent of the Connective Schema can avoid the unwanted conclusion that it is empty.

As we've seen in this section, the worry about emptiness does not merely arise for "hypertraditional" thick concepts, but also for thick concepts readily employed within "modern," reflective societies. The problem of emptiness, as I'm construing it, is the problem of avoiding the implication that very many of our obviously applicable thick concepts are in fact empty. There's no reason to think the Connective Schema can avoid this problem. Let's now ask how the Qua Schema fares against it.

IV.3. Emptiness and the Qua Schema. The Connective Schema runs into the problem of emptiness because this schema claims that the application of an obviously applicable thick concept, *murder*, requires the truth of a particular sort of bridge principle, but we have no reason to think there is any such true bridge principle. Of course, the Qua Schema is not vulnerable to that very same problem, since it does not involve any bridge principle. Is there some other way for the problem of emptiness to arise?

Let me lay out a potential argument for the conclusion that the Qua Schema leads to unwanted emptiness. This argument involves the notion of what Judith Thomson calls a

“goodness-fixing kind”.²⁴ Later, I’ll state Thomson’s account of a goodness-fixing kind, but for the purposes of the Qua Schema we need not commit to any particular definition. An intuitive notion will suffice. Basically, the notion of a goodness-fixing kind is aimed at distinguishing between two types of kinds: For some kinds K1 there intuitively seems to be such properties as being good qua K1. For example, it’s intuitively plausible that there are such things as being good qua dancer, good qua human being, good qua teacher, etc. But for other kinds K2, there just does not seem to be any such thing as being good qua K2. To borrow a few of Thomson’s examples, there is no such thing as being good qua pebble, or good qua smudge, or good qua thumb-movement. So, the basic idea is that the notion of a goodness-fixing kind (however it’s spelled out in detail) should be able to distinguish between the two classes. The K1’s are the goodness-fixing kinds while the K2’s are not. (Similar can be said about the parallel notion of a badness-fixing kind).

Now the problem of unwanted emptiness might arise if it turns out that, for some obviously non-empty thick concept T, its corresponding N_{SUB} is not a goodness/badness fixing kind. The envisioned argument would go as follows. (1) If there is a property of being good qua K, then K is a goodness-fixing kind.²⁵ But (2) there are non-empty positive thick concepts for which the relevant N_{SUB} is not a goodness-fixing kind. Therefore, for some non-empty positive thick concepts, there’s no such property as being good qua N_{SUB} . And this means we’ll get unwanted emptiness under the Qua Schema. (A parallel argument can be run with regard to negative thick concepts and badness-fixing kinds).

As suggested earlier, it looks like premise (1) must be accepted if we are to draw the intuitive distinction between the relevant kinds. Intuitively, we cannot allow that, for every kind K, there is a property of being good qua K. So, the notion of a goodness-fixing kind is used to pick out all and

²⁴ See Thomson (2008, ch.2).

²⁵ This is essentially a claim that Thomson makes (2008, 21).

only those kinds K1 such that there is a property of being good qua K1. And premise (1) just follows from this way of employing that notion. So, it looks like we must accept premise (1).

But what do we say about premise (2)? What are the thick concepts referred to in that premise? It is not immediately clear what they would be, so let me provide a few possibilities. Below is a group of (what I take to be) non-empty thick concepts, along with what might be their corresponding N_{SUB} 's:

Generous: N_{SUB} = giver

Courageous: N_{SUB} = risk-confronter

Loyal: N_{SUB} = commitment-keeper

Honest: N_{SUB} = assertion

Murder: N_{SUB} = cause of death

Rude: N_{SUB} = social disregard

Cruel: N_{SUB} = inducement of pain

This list does not include *lewd*, *perverse*, and *chaste*, since some believe these might be empty.²⁶ The seven examples mentioned above give us an idea of what premise (2) might be saying. These seven are probably not empty. And premise (2) might then be claiming that at least one of the above N_{SUB} 's is not a goodness/badness fixing kind.

But is that true? Ultimately, I see no reason to concur, unless we are being unduly restrictive with our notion of a goodness/badness-fixing kind. If we restrict goodness/badness-fixing kinds to artifact-kinds (e.g. umbrella) then this would clearly exclude all of the above N_{SUB} 's. However, this is unduly restrictive. It clearly makes perfect sense to say that someone is good qua human being, good qua teacher, or good qua dancer, even though these don't involve artifact kinds.

Moreover, it is also worth noting that Thomson's particular notion of a goodness-fixing kind is probably not restrictive enough to exclude the above N_{SUB} 's from counting as goodness/badness-fixing kinds. For Thomson, a kind K is a goodness-fixing kind just in case "what being a K *is* itself" determines the standards a K must meet if it's to be good qua K.²⁷ (Presumably, then, a K would be

²⁶ Gibbard (1992 and 2003) seems willing to claim that *lewd* and *chaste* are empty. Priest (1997), Primoratz (1997), and Slote (1974) argue that *perverse* is empty.

²⁷ Thomson (2008, 83).

a badness-fixing kind just in case “what being a K *is* itself” determines the standards a K must *fail to meet* if it’s to be bad qua K). For example, the kind umbrella is a prime candidate for a goodness-fixing kind, since what being an umbrella *is* itself determines the standards an umbrella must meet, if it’s to be good qua umbrella. Still, among Thomson’s examples of goodness-fixing kinds, there are also many non-artifact kinds, such as beefsteak tomato, tiger, human being, and liar. Moreover, she also includes a number of action-kinds—such as asserting a proposition, answering a question, and pronouncing a word.²⁸ So, the inclusiveness of her notion makes it hard to claim that the above N_{SUB} ’s are not goodness/badness-fixing kinds. I am inclined to be similarly inclusive with regard to the notion of a goodness-fixing kind.

It is worth pausing on the N_{SUB} ’s associated with the above negative thick concepts. One might object to my view by claiming that the sentence ‘*a* is bad qua K’ just means that *a* is a defective K.²⁹ It will then be pointed out that it’s not at all obvious that there is any such thing as a defective cause of death or a defective inducement of pain, etc. These properties look very odd indeed. In response, let me point out that the above synonymy claim is clearly mistaken. Note that there is also no such thing as a defective violinist or a defective dancer. These properties seem very odd as well. But, of course, it’s quite natural to use the locutions ‘bad qua violinist’ and ‘bad qua dancer’. These K’s are clearly badness-fixing kinds. So something has gone awry with the objector’s reasoning. I think the synonymy claim is mistaken. The reading of ‘bad qua K’ as ‘defective K’ breaks down when it comes to non-artifact kinds.³⁰

In short, I see no reason to doubt that these obviously applicable thick concepts are associated with goodness/badness-fixing kinds. So, I see no reason to accept premise (2). Of

²⁸ Thomson (2008, ch.6).

²⁹ Although Thomson doesn’t speak in terms of bad qua K, she seems to take the notion of being a defective K as the “flip side” of being good qua K. See chapter 12. This suggests that she would take being bad qua K as being roughly the same thing as being a defective K.

³⁰ In this regard, it’s worth noting that Thomson struggles to give an adequate account of a defective human being, in contrast to the simple account of a defective toaster (2008, ch. 12).

course, this doesn't mean that no thick concepts will turn out to be empty on the Qua Schema. It only means that it doesn't look like we'll get any unwanted emptiness. Some thick concepts, such as *lewd*, *chaste*, and *perverse*, are not obviously applicable, and so it wouldn't be a problem for the Qua Schema if they turned out to be empty. In fact, the emptiness (or non-emptiness) of these concepts might be an interesting conclusion to draw from the Qua Schema.

In this section we've also learned that the Qua Schema is tightly bound up with the notion of a goodness-fixing kind. For a thick concept to have any application, it must entail a goodness/badness-fixing kind. This connection between the Qua Schema and goodness-fixing kinds will re-emerge later in section V.

IV.4. Revised Connective Schema Explained. The above has been mainly aimed at criticizing the Connective Schema, and then comparing it with the Qua Schema. The Connective Schema falls prey to the problem of emptiness while the Qua Schema appears to come out unscathed. But what about other explanatory accounts? Are there other versions that might avoid the problem of emptiness? Christine Tappolet and Stephan Burton advocate a Revised Connective Schema that at least shows promise of avoiding the problem of emptiness. It goes as follows:

Revised Connective Schema:

For any x and any thick concept T, x falls under T iff x is N and x is E, and x is E in virtue of some particular instance of N-ness.³¹

The important part of this revision has to do with the phrase 'some particular instance of N-ness.' Whereas the original version of the Connective Schema holds that the evaluative content obtains in virtue of N-ness, the Tappolet/Burton view holds that the evaluative content obtains "in virtue of some particular instance of" N-ness.

³¹ See Tappolet (2004, 216) and Burton (1992, 31).

But we are immediately left to wonder whether this change marks any improvement over the two accounts already criticized (i.e. the Conjunctive Schema and the original Connective Schema).

Tappolet does not explain the account, so we will have to look at what Burton says.

Burton suggests that he wants this account to be importantly different from how I construed the original Connective Schema. On his view, a thick concept groups together all the cases in which an act is good (or bad) in virtue of being some particular instance of N-ness. But what does this mean? Burton tries to spell this out by claiming that the concept groups together its instances “simply ignoring the fact that [the] evaluation depends in each case on the various different characteristics and contexts of the various particular instances of [N-ness].”³² Thus, according to Burton the analysis of a thick concept should allow the evaluative content to depend at least partly on certain properties other than N-ness. It therefore seems that Burton wants to avoid being committed to a bridge principle such as,

- (iv) For any y, if y is N, then y is E,

where N-ness is by itself sufficient for the evaluative content to hold. So, Burton at least aspires for a view that is importantly different from the original Connective Schema.

IV.5. Revised Connective Schema Rejected. It is doubtful that this sort of view can deal with the linkage problem. In particular, there’s no reason to think the relation picked out by “in virtue of some particular instance of” differs in extension from the relation of mere conjunction. One would hope this relation is importantly different from mere conjunction, but it’s not at all clear how it is. First off, recall Burton’s explanation of this relation. On his view, a thick concept groups together its instances “simply ignoring the fact that [the] evaluation depends in each case on the various different characteristics and contexts of the various particular instances of [N-ness].” Our first clue

³² Burton (1992, 31).

that something has gone awry here is that a proponent of the Conjunctive Schema could say exactly the same thing about each thick concept.

Now consider that, much like the Conjunctive Schema, Burton's account looks like it will run into the same problem with the Dismissive Surgeon case. Recall Dr. Jones' action. Her usage of instrument X was bad in virtue of some particular instance of being an inducement of pain. It was bad in virtue of the particular instance described within the Dismissive Surgeon case. Certainly, the badness of her action depends at least partly on various properties other than its being an inducement of pain (e.g. its being a promise-breaking, etc.). But this is allowed on the present understanding of Burton's account. His account seems to provide no reason to deny the claim that Dr. Jones' action was bad in virtue of some particular instance of being an inducement of pain. And so it has no means to explain why her behavior does not count as cruel.³³

To summarize: the Burton/Tappolet account aspires to be a view that is not committed to a bridge principle, but this leaves open much doubt as to whether it can handle the linkage problem. The Qua Schema looks preferable to this revised explanatory account.

IV.6. A Residual Worry. It is time to deal with a residual worry regarding the criticisms I've leveled against rival accounts. Throughout this paper, I have rejected three general schemata—the Conjunctive Schema, the Connective Schema, and the Revised Connective Schema. In doing so, I relied on certain assumptions about what non-evaluative expression is to be substituted into each

³³ It should be pointed out that Tappolet differs from Burton in that she wants to avoid analyzing thick concepts in terms of thin concepts like *bad* and *wrong*. Instead, Tappolet analyzes them with "affective concepts", such as *admirable*, *pleasant*, *amusing*, *desirable*, and *pride-worthy*. According to Tappolet, these concepts are conceptually related to certain kinds of attitudes, such as desire and admiration, and they imply (roughly) that such attitudes are the appropriate ones to have towards the thing in question (Tappolet, 211). So, for example, instead of analyzing *courage* in terms of *good*, she analyzes it in terms of the concept *admirable* (Tappolet, 216). However, this does not help her in dealing with the sort of case I've provided. It would make no substantive difference if we were to replace 'bad' in the above paragraph with the appropriate affective term, such as 'offensive', 'appalling', or 'off-putting'. After all, it is easy to imagine a variant case in which Dr. Jones' action is offensive, appalling, off-putting, etc. for reasons having nothing to do with its being a known inducement of pain. So it is doubtful that Tappolet's focus on affective concepts will be useful here.

schema for ‘N’. And under these assumptions I argued that the relevant schemata fall prey to two problems—the linkage problem and the emptiness problem. But one might take the emergence of these problems, not as showing that the relevant schemata are untenable, but as merely showing that I’ve misidentified the relevant non-evaluative contents.

However, this rejoinder is not sufficient as it stands. It is not enough merely to claim that these problems arise because the relevant non-evaluative contents have been misidentified. The proponent of this rejoinder must also provide one of two things—either (i) a correct statement of that non-evaluative content, or (ii) an explanation of why we are unable to supply a correct statement. Let me briefly address these two possibilities in turn.

Regarding (i), I have already argued that there is no obvious way in which the non-evaluative contents of *murder* and *cruel* can be specified in accordance with the rival schemata. But the proponent of (i) might insist that this content can at least be partially specified. For example, Daniel Elstein and Thomas Hurka have together pointed out that it is possible to partially specify the non-evaluative content of a thick concept without providing a full specification. For example, on their view the analysis of *murder* might go something like this:

For any act x, x is a murder iff there are properties A, B, and C of a general type associated with killing, such that x has A, B, and C and anything that has A, B, and C is bad.³⁴

This analysis would be an instance of an explanatory schema. But the analysis leaves it partially unspecified as to what A, B, and C are supposed to be, except that they are somehow associated with killing.

In reply, let me first point out that the analyses associated with this type of schema might be, for all we know, equivalent to the analyses mapped out by the Qua Schema. For example, since the above analysis does not fully specify what A, B, and C are, then there is no reason to think that this

³⁴ See Hurka and Elstein (2009, 521-22) for this general pattern of analysis.

analysis is not necessarily equivalent to a fully-specified analysis of *murder* that conforms to the Qua Schema (such as the one provided in section III.1). But if these two schemata could be equivalent, then what grounds do we have for preferring one over the other as an answer to our main question? The only grounds seem to be the relative simplicity of the analyses mapped out by each schema. And on this score, I think the Qua Schema clearly provides the better answer to our main question. We saw in section IV.2 that it seems likely that A, B, and C will be highly complex, and possibly disjunctive, properties. If that's right, then this would complicate the above analysis quite considerably. But, as we've seen, the Qua Schema seems able to provide very simple analyses—there is no reason to suspect that its analyses will need to be “built up” in the ways required for alternative schemata. If this is right, then it appears that the Qua Schema provides a simpler account of how thick concepts hold together evaluation and description. And this would tip the scales in favor of the Qua Schema.

Let's now consider (ii). The proponent of this option would try to explain the fact that we are now unable to provide a complete specification of the thick concept's non-evaluative content in conformity with one of the rival schemata. For instance, it might be claimed that this specification has not yet been discovered because it is very complex and is partly determined by the end result of our reasoning about that thick concept.³⁵

My reply to (ii) is basically the same as the one I provided with regard to (i). For all we know now, the analyses mapped out by the rival schema, whenever the full specification is provided, might turn out to be equivalent to the analyses mapped out by the Qua Schema. So, it seems that relative simplicity will be the only grounds we have for comparing the Qua Schema with these alternative schemata in terms of how well they answer our main question. And, as this second approach plainly

³⁵ This is a variation of the sort of reply Frank Jackson gives to the Moorean open question argument. See Jackson (1998, 151). This explanation, of course, provides no reason for thinking there is in fact a specification that will vindicate their account. For this, we must rely on their promise of uncovering one.

admits, the analyses mapped out by the rival schema will be “very complex.” So, it looks like simplicity will tip the scales in favor of the Qua Schema once again.

V. An Argument for the Qua Schema.

V.1. Explanatory Power. So far I have assumed that thick concepts are analyzable, and I have argued that the Qua Schema fares well in face of the problems that arose for three rival schemata.

However, the fact that the Qua Schema fares well against these problems does not amount to an argument in its favor (at least not an argument that should persuade the Non-Separatist). In this section, I argue that there is strong reason to accept the Qua Schema, and in the next section I explain why these reasons apply equally to the Non-Separatist.

In brief, I want to claim that the Qua Schema has significant explanatory power, and that this provides strong reason to accept it. The explanatory hypothesis I advance is the hypothesis that each thick concept can in principle be analyzed according to the Qua Schema. And I want to claim that this explains two general features about thick concepts.

First, it explains how thick concepts hold together an evaluation plus a description. In other words, it answers our main question. On the Qua Schema, evaluation is represented as the evaluative content *that a is good/bad qua* N_{SUB} , and description is represented as the non-evaluative content *that a is* N . And these two contents are “held together,” not merely by conjunction, but also by the relationship between N and N_{SUB} . That the Qua Schema answers this challenge counts significantly in its favor. Insofar as an account of thick concepts does not answer the main question, that account will suffer by comparison to ones that do answer it.

The second explanatory point will require considerable development over the next few subsections. I’ll argue that the Qua Schema explains an interesting relationship between thick concepts, on the one hand, and the concepts *good in a way* and *bad in a way*, on the other. As we’ll see,

this relationship is not easy to explain, since many views fail to explain it. But the Qua Schema provides a particularly simple and elegant explanation of this relation. Let's now consider what this relationship is.

V.2. Ways of Being Good and Bad. As I have argued in a previous paper (i.e. "How Are Thick Terms Evaluative?"), there is an interesting relationship between thick concepts, on the one hand, and the concepts *good in a way* and *bad in a way*, on the other. This relationship can be summed up by the following two principles:

- (A) For any positive thick concept *T*, it is conceptually true that, if *x* falls under *T*, then *x* is good in a way.
- (B) For any negative thick concept *T*, it is conceptually true that, if *x* falls under *T*, then *x* is bad in a way.³⁶

Support for these principles comes largely from linguistic data involving the infelicity of sentences like the following:

- # Nancy is generous and she's not good in any way.
- ? Nancy is generous and she's good in a way.

This kind of data is wide spread across all paradigmatic thick concepts. And, as I argued, the relevant instances of (A) and (B) provide the best explanation of this linguistic data.

Here, I would like to focus on a broader question: How are these two principles themselves to be explained? It is worth noting that a schema cannot easily explain these principles without

³⁶ It might be claimed that everything will count as good in some way or other (see Thomson (2008, p. 10) for an argument for this claim). But, if it is conceptually true that everything is good in a way, then (A) might be true for completely trivial reasons having nothing to do with thick concepts. Even wholly non-evaluative concepts would have similar claims true of them. In response, let me note that, if it is conceptually true that everything is good in a way, then the arguments I leveled in support of (A) would actually support a slightly stronger principle, such as

(A*) For any positive thick concept *T*, one is in a position to know (merely by virtue of one's full competence with *T*) that, if *x* falls under *T*, then *x* is good in a way.

More precisely, if it is conceptually true that everything is good in a way, then the best explanation of the sort of linguistic data mentioned just below would be something like (A*) rather than (A). For more on this, see my paper "How Are Thick Terms Evaluative?" This complication will make little substantive difference to the discussion below, so I omit the details.

falling prey to the linkage problem. For example, consider that it would be implausible to analyze thick concepts directly in terms of good in a way or bad in a way. The reason is that it is far too easy for something to count as good in a way or bad in a way. Any attempt to plug ‘good in a way’ or ‘bad in a way’ into an analysis would yield an analysis for which cases of coincidence can easily arise (like that of the Dismissive Surgeon case).³⁷ Thus, the linkage problem suggests that a markedly different route must be taken for explaining (A) and (B). As I’ll now show, the Qua Schema can provide such a route. A very elegant explanation of (A) and (B) falls directly out of the Qua Schema.

V.3. The Qua Schema’s Explanation. Let’s see how the Qua Schema plays an integral role in explaining (A) and (B). Here I’ll focus exclusively on (A), though what I say below will apply similarly to (B).

The following is the first part of our explanation of (A), and it’s taken directly out of the Qua Schema:

Part 1: For any positive thick concept T, it is conceptually true that, if x falls under T, then x is good qua N_{SUB} .

In order to achieve a full explanation of (A) we will need to combine Part 1 with a second claim as follows:

Part 2: It is conceptually true that, if x is good qua N_{SUB} , then x is good in a way.

Is Part 2 plausible? I think so. Here’s the argument: Let us suppose that x is good qua N_{SUB} . From this, it obviously follows that there is such a property as being good qua N_{SUB} . Moreover, according to what was said in section IV.3, if there is any such property as being good qua N_{SUB} , then N_{SUB} must be a goodness-fixing kind. Furthermore, no matter how we spell out the notion of a

³⁷ For example, any of the rival schemas could have plugged ‘good/bad in a way’ directly into their analyses (thereby explaining (A) and (B)), but the linkage problem would have arisen even more easily in that case. For instance, it’s incredibly easy to show that Dr. Jones’ action was bad in some way or other. It’s more difficult to show that it’s morally bad or wrong. Even the original Connective Schema could not have effectively avoided the linkage problem if it substituted ‘bad in a way’ for ‘E’ in the analysis of *murder*. Due to how easy it is to count as bad in a way, all deliberate human killings will turn out to be bad in some way or other, even those that do not count as murders.

goodness-fixing kind in detail, our account will need to sustain an obvious truth—If N_{SUB} is a goodness-fixing kind, then being good qua N_{SUB} is a way of being good.³⁸ And lastly, if being good qua N_{SUB} is a way of being good, then anything that has this property is therefore good in a way. The above premises are all conceptual truths, and from them it follows that Part 2 is true.

Together, parts 1 and 2 can obviously explain (A). But the crucial point is this: since Part 1 comes straight from the Qua Schema, and the argument for Part 2 requires key notions that are bound up with the Qua Schema (e.g. goodness-fixing kind), it looks like the Qua Schema can play an integral role in explaining (A). Similar claims can be made for (B). In the next two sections I will explore the question of whether (A) can be explained without the Qua Schema.

V.4. Two Inadequate Explanations. Let's now look at two alternative ways of explaining (A). As I will show, both of them fail for similar reasons.

First, consider a more general variant on Part 1. Let's suppose the following explanation falls out of an alternative schema in much the same way that the original Part 1 fell out of the Qua Schema:

Part 1': For any positive thick concept T, it is conceptually true that, if x falls under T, then x is good ____ N.

Just as before, we can let 'N' exhaust T's non-evaluative content, and we can let 'good ____ N' stand for any complex expression where 'good' is combined with the relevant non-evaluative phrase modifying 'good' (e.g. 'good at dancing', etc). Given *Part 1'*, the corresponding version of Part 2 would need to state:

³⁸ This certainly doesn't follow from the generic notion of a goodness-fixing kind that I've been using. But the point is this: However we decide to mark out the difference between the goodness-fixing kinds and all of the other kinds, the account will need to sustain an obvious truth like this. One reason we're inclined to affirm that there is such a thing as being good qua teacher is that, if we consider only what it is to be a teacher (bracketing other uses a teacher may have), then there still seems to be a particular way in which a teacher could be good. And one reason we're inclined to deny that there is such a thing as being good qua pebble is that, if we consider only what it is to be a pebble (bracketing other uses a pebble might have), there seems to be no particular way in which a pebble could be good.

Part 2': It is conceptually true that, if x is good ____ N, then x is good in a way.

The problem is that Part 2' is highly implausible. Consider the properties of being good-for-nothing and being good-for-no-one. These properties fit the form of 'being good ____ N', but they are not ways of being good. And they will clearly provide counterexamples to Part 2'.

There's a deeper point lying beneath the surface, a point that does not rely on these particular counterexamples. Even if we find a way of excluding the peculiar properties of being good-for-nothing and good-for-no-one, we're still left without any resources for showing that Part 2' is at all plausible. In advance of knowing what the property of being good ____N is supposed to be, we cannot confidently hold that everything that has this property will be good in a way. The peculiar properties mentioned above serve to highlight this problem.

Christine Tappolet's schema falls prey to a similar problem. She analyzes positive thick concepts in terms of affective concepts, such as *admirable*, *amusing*, *pleasant*, *desirable*, *pride-worthy*, etc. And she then claims that these affective concepts pick out properties that are determinates of good in a way. Thus, her explanation would go as follows: Let 'is A' pick out whatever affective property is associated with T:

*Part 1**: For any positive thick concept T, it is conceptually true that, if x falls under T, then x is A.

*Part 2**: If x is A, then x has a property that is a determinate of being good in a way.

The same problem that arose for Part 2' also arises for Part 2*. Why should we think the property of being A is a determinate of good in a way? In advance of knowing what the property of being A is supposed to be, we cannot be confident that it is a determinate of good in a way. Even if we grant that some affective properties are determinates of good in a way (e.g. being admirable, being

pleasant, etc.), this does not guarantee that the affective properties associated with every positive thick concept are going to be determinates of good in a way.³⁹

Here I've dispelled two alternative ways of trying to explain (A). The problems with those explanations are representative of a broader difficulty about explaining these principles. We need some reason for thinking the evaluative concept featured within our analysis of a thick concept will conceptually entail *good in a way*. But, in advance of knowing what that particular concept is that's associated with each thick concept, we don't seem to have any guarantee that it will conceptually entail *good in a way*. However, unlike the above two explanations, the Qua Schema provides such a guarantee through its use of the property of being good qua N_{SUB} , as well as the related notion of a goodness-fixing kind. If the Qua Schema is correct, then, no matter what thick concept we're dealing with, we're always guaranteed to have a way of explaining its connection with good in a way and bad in a way. Although I've not addressed every possible rival explanation, the general difficulty mentioned seems likely to arise for nearly all schemata. The Qua Schema provides a unique way around this difficulty, and it therefore provides a strong explanation of the claims in question.

V.5. A Simpler Explanation. But there is a simpler explanation that should be mentioned. It is possible to accept Part 1 and Part 2 while remaining neutral on whether the Qua Schema is true. These two explanatory principles fall out of the Qua Schema, but they could be accepted all by themselves, without accepting the Qua Schema. This explanation would provide a simpler way of explaining (A), since, unlike the Qua Schema, it doesn't require us to assume the analyzability of thick concepts. So, it looks like there is a simpler way of explaining (A) without accepting the Qua Schema.

³⁹ Tappolet (2004). A further problem for Tappolet, as noted in section III.4, is that her view cannot deal with the linkage problem.

In reply, let me point out that we've already seen reason for accepting the Qua Schema over this simpler alternative. The Qua Schema explains how a thick concept holds together an evaluation and a description. That is, it answers our main question. But the above simpler explanation of (A) leaves our main question completely unanswered. So, it is still incumbent upon the proponent of this simpler view to answer this challenge. And since the simpler view is obviously compatible with the Qua Schema, one way for its proponent to answer our main question is for him to accept the Qua Schema.

In short, it looks like the Qua Schema has significant explanatory power, more than any of the views considered above. It explains how thick concepts hold together an evaluation and a description. And it also explains the interesting relationship between thick concepts, on the one hand, and the concepts *good in a way* and *bad in a way*, on the other. I take this over all explanatory power to be good grounds for accepting the Qua Schema.

VI. Non-Separatism

This simpler explanation reminds us of the Non-Separatist from whom we parted ways at the beginning of our search. Non-Separatists could accept the simpler explanation of (A) just mentioned, but they must find another way of answering our main question, one that does not presuppose the analyzability of thick concepts.

Non-Separatists have been surprisingly silent on the question of how thick concepts hold together evaluation and description. The only detailed answer, to my knowledge, comes from Jonathan Dancy. Dancy explicitly concerns himself with “the relation between evaluation and description in a thick concept.”⁴⁰ Let's look at how he cashes out this relationship with regard to the thick concept *lewd*. Dancy claims that ‘lewd’ applies to something *x* on account of the fact that *x* has

⁴⁰ Dancy (1995, 275).

certain features F “in the right way.” But what is “the right way”? Once Dancy expounds on this, it becomes clear that his account is not an analysis in the traditional sense. This is because his account is circular. According to Dancy, “the right way” can only be “captured” by saying that it’s the way that merits “a response merited by the lewdness.”⁴¹ Basically, Dancy’s notion of “the right way” can only be cashed out in terms of lewdness itself. The account is therefore circular.

I do not object to Dancy’s account on the grounds of its circularity. But I do want to claim that, as far as answering our main question goes, Dancy’s account is inferior to the Qua Schema. We want to know how *lewd* holds together an evaluation and a description. In addition to explaining how these two aspects relate, an optimal answer would also tell us what these aspects are. But Dancy cannot say much at all about the evaluative aspect. He expounds on the evaluation associated with *lewd* by saying that this concept specifies that certain responses are merited towards the item in question. But when it comes to saying which responses are merited, Dancy has no answer except to say that they’re the responses merited by lewdness. But this leaves an important part of its evaluative aspect mysterious. Is admiration one of the responses merited by lewdness? What about horror? What about apathy? Intuitively none of these responses is merited by lewdness. But Dancy’s account has no resources for telling us why, except to say that they are not in fact merited by lewdness.

The Qua Schema, on the other hand, is able to provide a non-circular answer to our main question. The evaluative aspect associated with *lewd* is plausibly the concept *bad qua sexual display*. And this relates to the non-evaluative content of *lewd* in that *sexual display* is at least part of that non-evaluative content. Given that the Qua Schema has an answer to our main question, and that Dancy’s answer is at best incomplete, I take this to be a significant advantage of the Qua Schema.

⁴¹ Dancy (1995, 276).

Thus, the analyzability approach has an advantage, provided the correct account can be formulated. I've argued that the Qua Schema is that account. Still, the Non-Separatist will not be persuaded by what I've said so far. For he has an argument for thinking the path of analysis is a dead-end.⁴² As I'll show in closing, this argument does not succeed against the Qua Schema.

Without delving fully into the argument, let me note that it assumes something that looks quite dubious under the Qua Schema. Their argument assumes that, if thick concepts were analyzable, then the two components of the analyses would have to be a *descriptive content* and an evaluative content. The former is plainly defined as something that picks out a property, while the latter, by contrast, is something that fails to pick out a property. It is under this assumption (and only under this assumption) that Separatism implies that a thick concept's evaluative content has no effect on whether the thick concept actually applies. And this is often seen as problematic.

Whether or not this consequence is deeply troubling, it certainly does not follow from anything said in this paper. Nowhere have I made the weighty assumption that evaluative content fails to pick out a property. I've managed to avoid this assumption by refraining from calling our two types of content 'evaluative' and 'descriptive', but instead calling them 'evaluative' and 'non-evaluative'. The latter delineations at least allow for the possibility that both contents pick out properties. So, it's compatible with everything I've said in this paper that both components have an effect on whether the thick concept applies.

But this is not likely to impress the Non-Separatist. He may possibly hold that, if there is any such thing as evaluative content, then such content always fails to pick out a property. This is a controversial claim in meta-ethics, and this is no place to meet the issue head on. But it's worth noting that Non-Separatists typically have had paradigmatic thin concepts in mind with regard to

⁴² See McDowell (1981). For those who rely on some form of this argument, see Dancy (1995, 263) and Williams (1985, 141).

this claim. And no matter how plausible this claim is with regard to the thin concepts incorporated by rival schemas, the claim does not seem to carry over to the concepts employed by the Qua Schema. That is, even if we grant that non-error-theorist cognitivism is problematic with regard to thin concepts (e.g. *good*, *bad*, *right*, and *wrong*), this does not mean that such cognitivism with regard to *good qua* N_{SUB} is problematic. To see this, let's suppose N_{SUB} is a goodness-fixing kind. This means that the concept *good qua* N_{SUB} is the same sort of concept as *good qua teacher*, *good qua human being*, and *good qua toaster*. But why would we want to insist that these concepts fail to pick out properties? As ethicists have pointed out, the arguments against non-error-theorist cognitivism about concepts like *right*, *wrong*, *morally good*, and *morally bad* are significantly less persuasive when it comes to concepts like *good qua toaster*.⁴³ For example, *good qua toaster* just does not seem to have the motivational character that is supposedly associated with concepts like *wrong*. It also seems less difficult to fit the property of being good qua toaster into the naturalistic, scientific world-view.

The point here is that non-error-theorist cognitivism seems to be on fairly secure ground with respect to *good qua* N_{SUB} . But then the Non-Separatist's main assumption looks highly dubious. Now that the Qua Schema is a plausible version of Separatism, it's not at all clear why the analysis of a thick concept would need to incorporate an evaluative content that fails to pick out a property. In short, the Non-Separatist's argument does not hold much sway against the Qua Schema, and we're now left to wonder why he doesn't follow us down the path of analysis.

⁴³ See Thomson (2008 ch. 3) and Geach (1956). I trust it is intuitively clear that being good qua toaster does not entail that the thing in question is good simpliciter. I don't insist on the stronger claim, which Geach and Thomson make, that there is no such thing as goodness simpliciter.

4. Knowledge as Exemplary Belief

I. Introduction.

I argue in this paper that knowledge is exemplary belief. Very roughly, a subject knows a proposition just in case her belief in that proposition is good qua belief. By itself, this account is extremely bare—it explicitly states neither a factivity condition nor an anti-luck condition. But, as I argue, these conditions are implicitly bound up with what it is to be good qua belief. The picture that will emerge is that factivity and anti-luck are standards that a belief must meet if it's to be good qua belief.

Section II of this paper is devoted to explaining the general notion of being good qua K. Then, in section III, I apply this general account to the case of belief so as to yield my account of knowledge as goodness-qua-belief. And in section IV I argue that this account of knowledge can resolve two epistemological quandaries—the Gettier Problem and the Value Problem. Before proceeding, however, let me briefly explain what the Gettier and Value problems are and how they relate.

The Gettier Problem is the task of stating informative, necessary, and sufficient conditions for knowledge, ones that explain why knowledge excludes the kind of belief described within Gettier-style cases. The kind of belief in question can be illustrated by Bertrand Russell's stopped-clock example (which actually predated Gettier's 1963 paper):

S has a true belief, *p*, as to the time of day, but only because he is looking at a clock that he thinks is going. In fact, it happens to be stopped (Russell 1948, 154).

The subject in this case believes truly that the time of day is such-and-such. But, intuitively, he does not know this proposition. The central element in this type of case is that the subject has a true belief purely as a matter of luck (e.g. had he checked the clock a few minutes earlier or later, he

would have formed a false belief).¹ This general type of case (which I shall call a “Gettier-style case”) has led epistemologists to believe that knowledge requires, in addition to true belief, a condition that excludes luck (e.g. see Pritchard 2005, 1-6; Dancy 1985, 134). In sections III-IV I will show that my account of knowledge involves an anti-luck condition that can adequately deal with Gettier-style cases.²

Not only does knowledge require something more than mere true belief, but it also seems more valuable than mere true belief. Socrates originally voiced this intuition in Plato’s *Meno*. The Value Problem, in its most basic form, is the task of explaining the greater value of knowledge over that of mere true belief. This problem has proven difficult to solve, since (as Socrates observed) true belief appears to have no less practical value than knowledge. For example, a true belief about how to get to Larissa is just as useful as having knowledge about how to get there—both will get you to that destination (Plato 1961, 97b-c). Because of this intuition, many epistemologists think the value of knowledge consists not merely in the practical value of getting things right. The problem then is to explain what else contributes to the total value of knowledge, aside from the value of true belief.

A number of epistemologists have recently complained that attempted solutions to the Gettier Problem tend to be so gerrymandered and ad hoc that they render the value of knowledge utterly mysterious (Kvanvig 2003, 129-30; Williamson 2000, 31; Zagzebski 2003, 12). The basic desideratum behind their complaint is that any solution to the Gettier Problem should ideally provide a solution to the Value Problem.³ That is, an account of what is required for knowledge (in

¹ A further important feature of most Gettier-style cases is that the true belief in question is justified. Indeed, all of Gettier’s own examples were cases like this (1963). For the purposes of this paper, I will understand Gettier-style cases broadly enough to include cases in which the luckily true belief in question is not justified.

² There are other types of counterexamples to traditional accounts of knowledge, stemming from skeptical and lottery scenarios. Such examples will not be addressed in this paper.

³ Whether a solution to the Value Problem is a requirement for any solution to the Gettier Problem is quite another matter. For a helpful discussion on this, see DePaul 2009.

addition to true belief) would ideally also tell us why knowledge is more valuable than mere true belief. The account proposed in this paper aims to do both.

My account of knowledge requires an explication of the notion of being good qua belief. But before this can be properly understood, the more general notion of being good qua K (for any kind K) needs to be explicated. In the next section, I explain this notion by proposing several principles that will be useful in supporting my account of knowledge in section III.

II. Good Qua K.

What does it mean to say that something is good qua K? According to the most common English usage, to say that something is good qua K just is to say that it's good *as a K*.⁴ We often apply this latter phrase in ordinary contexts—"Joe is good as a dancer", "Sue is good as a violinist, but not as a pianist," etc. And within such contexts this phrase is typically interchangeable with the relevant 'qua' locution—"Joe is good qua dancer," "Sue is good qua violinist, but not qua pianist," etc.

This synonym for 'good qua K' gets us in the right ballpark for understanding that expression, although it tells us nothing illuminating about the property of being good qua K. Can we say something more informative? In what follows, I will provide an account of being good qua K. The account, however, is not meant to be a conceptual analysis in the traditional sense. This is because the account of being good qua K is circular, as we'll see below. Nevertheless, this account is sufficiently informative (despite its circularity), and it will be extremely useful for the theory of knowledge advanced later.

⁴ Neither of these locutions means the same thing as 'good *for a K*'. Thomson makes this mistake (2008, 19), but simple examples can show their non-equivalence: Suppose that philosophers are notoriously bad athletes, but one day I see you playing basketball and notice that you're actually quite good at basketball. In this context, I can truthfully say that you are good for a philosopher (meaning that you're good at basketball for a philosopher). But I would have no grounds to say that you're good as a philosopher or good qua philosopher, at least not in the absence of evidence of your philosophical abilities. Now suppose, alternatively, that philosophers are notoriously good at poker, but I see you playing an absolutely pathetic game of poker (hardly even knowing the rules). In this context, where poker-playing skills are salient, it would be false of me to say that you're good for a philosopher. But, for all that matters, I could truthfully say that you're good as a philosopher or good qua philosopher, given that you have the requisite philosophical abilities.

We can provide an informative account of being good qua K, if we start off with an idea recently proposed by Judith Thomson. According to Thomson, to be good qua K is to be “a model, exemplar, paradigm, or good specimen of a K” (2008, 19). As I understand Thomson, each of these four items is meant to express basically the same thing (and in a moment I shall explain what that is). But, for the sake of simplicity, let’s set models, paradigms, and good specimens to the side, and fix on the notion of an exemplar as the central characterization of being good qua K. Thomson’s underlying claim can be summed up with what I call the GE Principle. For any kind K and any x which is a K,

GE Principle

x is good qua K iff x is an exemplar with respect to Ks.

This principle links goodness qua K with the property of being an exemplar (hence the name “GE Principle”). Occasionally, I will speak of “exemplary Ks” as shorthand for what’s expressed in the explanans.

But what is an exemplar? An exemplar is something that’s worthy of being imitated, copied, patterned after, emulated, etc.⁵ For the sake of simplicity, I will primarily speak of exemplars in terms of imitation, although other notions—such as copying—may appear in certain contexts if they prove more apt. I will also use ‘imitates’ and ‘copies’ to denote a form of resemblance which can be partially specified as follows: A imitates (or copies) B in respect R only if A possesses the properties of B that are associated with R. Thus, for the movie version of *The Scarlet Letter* to be a copy of the book with respect to storyline, the movie must have the same storyline as the book.

This notion of an exemplar immediately brings up another question: In what respect is an exemplar worthy of imitation? Clearly an exemplar need not be worthy of imitation in all respects. If Sue is an exemplary violinist, but a lousy pianist, then she is worthy of imitation with respect to

⁵ There is also a purely non-evaluative use of the term ‘exemplar’, according to which exemplars are the *typical* examples of a given class. But this non-evaluative usage should be disregarded for purposes of this discussion.

violin-playing, but not piano-playing. The mere fact that something is an exemplar with respect to a given kind does not mean that the item is worthy of imitation with respect to *all* of its qualities. Rather, an exemplar with respect to Ks is worthy of imitation in certain particular respects, and those respects are fixed by whatever K is.

What does it mean for a kind K to fix a respect in which an exemplar is worthy of imitation? An adequate answer to this question requires that we first understand what Thomson has called “goodness-fixing kinds.”

For some kinds K there is obviously no such property as being good *qua* K—to use some of Thomson’s examples, there is no such thing as being good *qua* pebble or good *qua* smudge (2008, 22). But, for other kinds K, there *is* such a property as being good *qua* K—e.g. the property of being good *qua* toaster. The latter are goodness-fixing kinds. But what makes it the case that a given kind is goodness-fixing? Thomson plausibly claims that there is such a property as being good *qua* K if and only if “what being a K *is* itself sets the standards that a K has to meet if it is to be good *qua* K.” She provides the following example:

...[B]eing a toaster is being an artifact manufactured to toast, and that itself sets the following standard for being good *qua* toaster: toasting well (2008, 21).

Clearly nothing similar can be said for the kind pebble or the kind smudge. Being a pebble or a smudge does not itself set any standards for being good *qua* pebble or good *qua* smudge. In short, this provides an account of what makes a given kind goodness-fixing. I find this view highly plausible with regard to being good *qua* K,⁶ and I propose that we accept it.

By accepting this view, we can now understand what it means for a given kind K to fix a respect in which an exemplar is worthy of imitation. In brief, K fixes a respect by virtue of setting a standard. I suggest that there is a one-to-one correspondence between respects and standards.

⁶ Thomson also wants to extend this account to the property of being a good K, which is slightly different. That’s a controversial claim to which I am not committed. See Hare (1967, 79) for a critique of that kind of claim.

More precisely, there is a one-to-one correspondence between the respects in which an exemplary K is worthy of imitation, on the one hand, and the standards that are set by what a K is, on the other.⁷ For example, a standard that is set by the kind toaster is that the members of that kind must toast bread well (if they're to be good qua toaster). And there is a corresponding respect in which an exemplary toaster is worthy of imitation—it's worthy of imitation with respect to whether it toasts bread well. Another example, adopted from Thomson, is this: A standard that a map of England must meet if it's to be good qua map of England is that it must be a "close-enough-to-correct map of England" (2008, 85). So, on my proposal, there is a corresponding respect in which an exemplary map of England is worthy of imitation—it is worthy of imitation with respect to whether it's close-enough-to-correct. So, we get one-to-one correspondence between standards and respects in the way just illustrated. More precisely, there exists a standard such that a K must have property P (for it to be good qua K) if and only if there exists a particular respect R (in which an exemplary K is worthy of imitation) and R is identical to the respect of whether the item has P.

One obvious point about goodness-qua-K is that it comes in degrees. But there may be a number of standards that an item must meet for it to attain a particular degree of goodness qua K. Furthermore, it's natural to assume that being good qua K to degree d1 would require different standards from those required by some other degree d2. So, a given K can be better or worse qua K depending on which standards it meets. For example, the minimum standard required for an item to be good qua toaster is that it toasts bread well. But for it to be good qua toaster to a very high degree (i.e. for it to be excellent qua toaster), it must satisfy other standards too—it must toast other food items well (e.g. waffles), it must also be safe, easy to use, easy to clean, etc. These are all

⁷ Just what are respects? And what are standards? I take a respect to be a set containing a property and its negation. For example, the respect of whether P just is the set containing the properties P and not-P. And in order to imitate x in respect of whether P, one must have whichever of these two properties x has (i.e. either P or not-P). I take a standard for being good qua K to be non-trivial condition that must be met for something to count as good qua K. In this way, standards are in effect non-trivial necessary conditions for being good qua K.

standards that are set by what it is to be a toaster, and just how good qua toaster something is will be determined by which of these standards it meets.

It follows from the above that, if there are n number of standards required for an item to be good qua K to degree d , then there will be n number of respects in which that item would be worthy of imitation (if it were good qua K to degree d). Rather than referring to each of these respects by name, I will abbreviate this by speaking of “the relevant K respects.” But the relevant K respects are to be understood as those that stand in one-to-one correspondence to the standards required for being good qua K to that particular degree d . Moreover, since the standards in question are fixed by what a K is, and since those standards correspond (as suggested) to the relevant K respects, it can also be said that the relevant K respects are fixed by what a K is.

The above claims provide some essential ways of clarifying the GE Principle. We now have a more detailed account of being good qua K : For any kind K , any x which is a K :

Specified GE Principle

x is good qua K to a particular degree iff x is worthy of imitation in the relevant K respect(s).

This particular version of the principle will need to be modified, as I’ll explain in a moment. But we must first pause over an interesting asymmetry within this principle.

The Specified GE Principle reveals an asymmetry that the original GE Principle conceals. In particular, the Specified GE Principle reveals the explanatory priority of being good qua K . To see this, let’s suppose we have the following facts—(1) the fact that Mark is good qua tennis-player, and (2) the fact that Mark is worthy of imitation in the relevant respects. Is one of these facts explanatorily prior to the other? Yes, intuitively the first fact is prior. That is, (1) explains (2), but not vice versa. Mark’s being good qua tennis-player explains the fact that he’s worthy of imitation, but his being worthy of imitation clearly does not explain the fact that he’s good qua tennis-player.

This explanatory priority is crucial because it helps us understand what *kind* of worth is expressed within the Specified GE Principle. In particular, this explanatory priority strongly suggests that the Specified GE Principle involves a form of instrumental worth—that is, the imitation of exemplary Ks has worth as a means to some other valuable end.⁸ To see this, let's revisit the example from above. We said that (1) explains (2), and not vice versa. But *how* exactly does (1) explain (2)? In particular, *how* does Mark's being good qua tennis-player explain the fact that he's worthy of imitation in the relevant respect? The answer is quite plain: If other people imitate Mark, then they can also become good qua tennis-player. Doing so is valuable as a means to the end of making them good qua tennis-player. In other words, imitating Mark has instrumental worth in that it can make the imitator good qua tennis-player.⁹

Generally speaking, the imitation of things that are good qua K is a means to the valuable end in which the imitator is also good qua K. And this is how x's being good qua K explains the fact that x is worthy of imitation in the relevant K respect. This should be made explicit within our final version of the GE Principle. For any kind K, any x which is a K, and any degree d,

Finalized GE Principle

x is good qua K to degree d iff copying x in the relevant K respect(s) has instrumental worth with regard to the end of making the copied version good qua K to degree d.¹⁰

⁸ For the purposes of this paper, I do not need to assume that the end in question must be something with intrinsic (=non-instrumental) value.

⁹ What is meant by "instrumental worth"? There are roughly two different conceptions, the second of which is what we must have in mind for the present discussion. The two conceptions are: (a) x has instrumental worth if and only if x actually brings about the existence of a valuable end. Or (b) x has instrumental worth if and only if x has potential to bring about the existence of a valuable end (see Rønnow-Rasmussen, 2002, 29). We must have (b) in mind for the purposes of this paper. This is because the whole reason for bringing in instrumental worth is to account for how (1) explains (2). But conception (a) does not help us account for that. After all, it could be that no one *actually* imitates Mark in the relevant respect, in which case imitating Mark does not *actually* bring about the valued end of making the imitator good qua tennis-player. So, in this case, imitating Mark would have no instrumental worth on conception (a). This can be easily resolved if we strictly appeal to conception (b). Thus, in this paper, I will assume that instrumental worth is to be understood according to (b).

¹⁰ Counterexamples to this principle would arise if we allowed deceptive masquerading to count as copying in the relevant sense. But in fact it does not. Recall that the relevant sense of copying requires that the one who copies actually takes on the properties of the one copied. By contrast, a deceptive masquerade requires that one give the false appearance of taking on those properties.

As one can see, the Finalized GE Principle is circular. Being good qua K is accounted for in terms of the goal of making the imitator good qua K. Perhaps this circularity could be removed, but that is beyond our present aim. The aim of this section is not to provide a non-circular analysis of the concept *good qua K*. Rather, the aim is to develop an account that tells us something interesting about being good qua K. And this account does so. It tells us that being good qua K is to be worthy of imitation with regard to a certain end—the end of making the imitator good qua K. Thus, the Finalized GE Principle is informative, despite its circularity. Moreover, as I'll now show, this Finalized GE Principle, together with a plausible claim about instrumentality, entails an important claim that will be useful later in this paper.

The Finalized GE Principle links goodness-qua-K with instrumental worth. But we must notice an important point about instrumentality. It is a conceptual truth that an instrumentally worthy item is conducive to a valuable end. In other words, the instrumentally worthy item has potential to bring about that valuable end.¹¹ Thus, we can take the following as a necessary condition for having instrumental worth:

(Instrumentality) For any x and any valuable end E, if x has instrumental worth with regard to E, then x has potential to bring about E.¹²

This helps elucidate the kind of worth that we've been discussing. Assuming this view is correct, we can combine it with the Finalized GE Principle in order to get a claim that will become important later. I shall call it the Instrumental Claim. For any kind K, any x which is a K, and any degree d,

¹¹ Just *how much* potential does an instrumentally worthy item need to have? For present purposes, we need not specify a precise degree of potential. An intuitive gloss will suffice for what I purport to do with this notion in section III. The basic idea is that potential and luck exclude one another. To the extent that it's a matter of luck that process P produces end E, we cannot say that P had potential to bring about E—or, at the very least, it does not *enough* potential for P to count as having instrumental worth with regard to E.

¹² There would be obvious counterexamples to this principle if we were operating with the conception of instrumentality referred to as (a) in footnote 9. Thus, we must be careful that we're dealing with conception (b) from that footnote.

Instrumental Claim:

If x is good qua K to degree d, then copying x in the relevant K respect(s) has potential to bring it about that the copied version is good qua K to degree d.

In section III.2, this claim will become particularly useful in establishing that my account of knowledge excludes luck.

The goal of this section has been to illuminate the general notion of being good qua K. The above general claims will help us better understand the specific property of being good qua belief.

Let's now examine goodness-qua-belief in depth, and see how it fits into my account of knowledge.

III. Knowledge.

On my view, knowledge is exemplary belief. And, as we saw above, this view can be seen as shorthand for an account of knowledge in terms of being good qua belief. In particular, the analysis of knowledge I propose goes like this:

Analysis of Knowledge

S knows P iff S believes P and this belief is good qua belief to at least degree d_T .¹³

I will suggest that this holds true for some particular degree d_T . As I claimed in the above section, the degree to which something is good qua K will depend on which standards for being good qua K have been met by that item. And those standards are set by what being a K is itself. I shall say more in what follows about what degree d_T is.

This analysis of knowledge is extremely simple by comparison to most post-Gettier analyses. At first glance, it seems *all too* simple. After all, it omits the two conditions that are nearly always found within analyses of knowledge—the factivity condition and the anti-luck condition. In the

¹³ The view of belief I have in mind is a rather standard picture: to have a belief in a proposition is to regard that proposition to be true in the right sort of way. The term 'belief' is actually ambiguous between (i) the state a subject is in when she regards a proposition as true, and (ii) the proposition she regards as true. But my usage of this term is strictly meant to express (i), and not (ii).

next two sections, I shall argue that these conditions (although not explicitly stated) are entailed by what it is to be good qua belief to degree d_T .

III.1. Factivity. In order to show that my account involves a factivity condition, I need to introduce what Thomson calls “correctness-fixing kinds.” The notion of a correctness-fixing kind is very similar to the notion of a goodness-fixing kind. For many kinds K_{CF} , there is such a property as being a correct K_{CF} —these are the correctness-fixing kinds. Thomson’s main examples include the properties of being a correct map of England, being a correct performing of Mozart’s C Major Sonata, and being a correct spelling of the word “chiaroscuro” (2008, 86).¹⁴ If a given K is not a correctness-fixing kind, then there is no such property of being a correct K —e.g. there is obviously no such thing as being a correct smudge or a correct pebble.

Notice that belief is a correctness-fixing kind. It makes perfect sense to say that someone has a correct (or incorrect) belief. Moreover, the fact that belief is a correctness-fixing kind follows from a widely-accepted view about the norm of belief— S has a correct belief that P if and only if S has a belief that P and it is true that P . Discussions of this principle typically take it to be a truism (Lynch 2009, 228; Shah and Velleman 2005). Calling it a truism may be an overstatement, since the sufficiency of this principle can be doubted.¹⁵ Nonetheless, even if this principle did not provide a sufficient condition, there would most likely be some alternative principle, phrased in terms of knowledge or justified true belief, that would provide a sufficient condition. And if there is such an alternative principle, then we must grant that there is such a property as being a correct belief. In other words, we must grant that belief is a correctness-fixing kind.

¹⁴ Gideon Rosen (2001) also highlights the fact that many different kinds can be characterized as correct or incorrect, even if they cannot be classified as true or false.

¹⁵ Williamson (2000, 46) accounts for belief as a state in which one treats P as if one knows P . If this were right, then it might be natural to hold a stronger view on the correctness of belief, such as the following: S has a correct belief that P iff S has the belief that P and this belief counts as knowledge.

Although the sufficiency of the above principle could be doubted, I will assume that it at least provides a necessary condition for the correctness of belief. That is, I will assume the truth of

Norm of Belief:

S's belief is a correct belief only if its propositional content is true.

This principle, as well as the claim that belief is a correctness fixing kind, gives us two of the key claims needed to show that there's a degree of goodness-qua-belief that entails truth.

The third claim is a principle about how goodness links up with correctness (hence the name "GC Standard"). For any correctness-fixing kind K_{CF} , and any x which is a K_{CF} , and some degree d^*

GC Standard:

If x is good qua K_{CF} to at least degree d^* , then x is a correct K_{CF} .

Basically, GC Standard tells us that correctness is a standard that a K_{CF} must meet if it's to attain a certain degree of goodness qua K_{CF} . As I'll suggest, this principle holds true for some particular degree d^* . In order to see why we should accept GC Standard, let's consider a potential objection.

Thomson makes some claims that initially seem to present troubles for GC Standard. She describes the following example regarding the correctness-fixing kind map of England:

[A] map of England might be a good one even if it is not the case that *everything* it represents about England is true of England. Its being good enough in other ways—informative, easy to understand, and so on—might outweigh the fact that it gets a small detail wrong (2008, 85).

This case provides a potential counterexample to GC Standard, since it describes how an item can be good qua K_{CF} without being a correct K_{CF} .

However, I want to suggest that this item is not good qua K_{CF} *to the appropriate degree*. And that's why it's not a genuine counterexample to GC Standard. Notice that the map in question would have been better qua map of England had it not gotten that small detail wrong. In particular, if the map had the qualities of being "informative, easy to read, and so on," but did not go wrong with respect to that small detail, then it would have been better qua map of England than the one originally described by Thomson. This, I think, is undeniable even as a general truth—any good

(but incorrect) map would be better qua map if it was correct. And if this is so, then it seems that correctness is a standard that must be met for a map of England to attain a particular degree of goodness qua map of England.

To see this in more detail, suppose for a reductio that correctness is not among the standards $S_1 \dots S_n$ that a map of England must meet if it's to be good qua map of England to the highest possible degree. Since correctness is not among $S_1 \dots S_n$ we can suppose that we have a map of England that meets all of $S_1 \dots S_n$ without being correct. Since this map has met all of $S_1 \dots S_n$, we can also suppose that it's good qua map of England to the highest possible degree.¹⁶ But these assumptions lead to a problem. It seems undeniable that this map would be better qua map of England if it was a correct map. So, if this map were correct, it would be good qua map of England to a degree greater than the greatest possible degree—but that's absurd! So, we must reject our initial assumption. Correctness is among the standards that a map must meet if it's to be good qua map of England to the highest possible degree. It follows that there is some particular degree d^* such that a map of England must be correct if it's to be good qua map of England to d^* . And if that's so, then the case Thomson has described is not a genuine counterexample to GC Standard.

My assumption that an item would be better qua map if it was correct might be questioned on the grounds that, in some cases, achieving correctness can compromise other virtues (e.g. being informative and easy to understand). However, although that may be true in some cases, it is certainly not necessarily true that correctness compromises other virtues. Suppose that map M has all the relevant virtues in a map (e.g. being informative, easy to understand, etc) except that M is incorrect. In particular, let's suppose that M falsely portrays a plot of land as having a small lake. The important point to notice is this: It is possible for M to be correct without compromising any of the other virtues, because there is a possible situation in which that plot of land actually has a small

¹⁶ I am assuming, plausibly, that there is a highest possible degree, and that degrees of goodness-qua-map do not extend upwards for infinite.

lake (exactly as depicted by the map). In that case, M would be correct while still possessing the other virtues in a map. Thus, the objection overlooks the fact that correctness can be attained by changing the world, rather than the map itself, and this change would not compromise other virtues. Similar claims can be made with regard to other correctness-fixing kinds.

The line of argument just sketched can be generalized to fend off any counterexample to GC Standard. This is because correctness always seems to make an item better qua K_{CF} than it would have been if it were an incorrect K_{CF} . Thus, GC Standard looks quite plausible.

Granting the truth of GC Standard, let's return to the specific case of being good qua belief. It can now be shown that there's a particular degree of goodness-qua-belief that entails truth. We've already established that belief is a correctness-fixing kind. That claim, combined with GC Standard, entails that there's a particular degree d^* such that an item is good qua belief to d^* only if it's a correct belief. And by combining this claim with Norm of Belief, we can conclude that an item is good qua belief to degree d^* only if its propositional content is true.

Finally, recall that I account for knowledge in terms of being good qua belief to at least degree d_T . And at the beginning of this section I promised that I would say what degree d_T is supposed to be. It's now time to do so. I stipulate that degree d_T , within my account of knowledge, is identical to degree d^* . This is the natural way of specifying d_T , given that it's highly intuitive that knowledge entails truth, and this specification allows for my account to support this intuition. By combining this specification of d_T with the argument in the preceding paragraph, it follows that my account of knowledge entails truth.¹⁷

¹⁷ It's worth mentioning that nothing I've said so far would rule out the possibility that degree d_T is the highest possible degree of goodness qua belief. But more than likely, d_T is not the highest degree. In particular, there are surely other features (aside from truth) that could make an item better qua belief—if one's belief is epistemically justified to a high degree, that would seem to mark an improvement qua belief. So, even though my argument thus far does not rule out the possibility that d_T is the highest possible degree, there are presumably other factors that could be brought to show that d_T is in fact not the highest degree.

III.2. *Anti-Luck*. It is no great victory that my account of knowledge can support the widely-held view that knowledge entails true belief. After all, it's precisely at this point that the most troublesome counterexamples arise for a theory of knowledge. Recall that Russell's stopped-clock example shows that true belief is not sufficient for knowledge, because it's merely a matter of luck that the subject believes correctly. For my account to handle Gettier-style cases like this, it must be shown that being good qua belief (to degree d_T) entails that it's not a matter of luck that the subject believes correctly.

In order to establish this, I will advance an argument based on the previously established Instrumental Claim and GC Standard. First, recall Instrumental Claim from section II: For any kind K , any x which is a K , and any degree d ,

Instrumental Claim

If x is good qua K to degree d , then copying x in the relevant K respect(s) has potential to bring it about that the copied version is good qua K to degree d .

And second, recall GC Standard from section III.1: For any correctness-fixing kind K_{CF} , any x which is a K_{CF} , and some particular degree d^* ,

GC Standard:

If x is good qua K_{CF} to at least degree d^* , then x is a correct K_{CF} .

From these two claims it follows that, for any correctness-fixing kind K_{CF} , any x which is a K_{CF} , and some particular degree d^* ,

Instrumental Correctness Claim

If x is good qua K_{CF} to degree d^* , then copying x in the relevant K_{CF} respect(s) has potential to bring it about that the copied version is a correct K_{CF} .

I shall assume that this last principle is already established, since it follows from principles established earlier. For the bulk of this section, I will argue for an anti-luck condition by employing the Instrumental Correctness Claim.

How do we get from the Instrumental Correctness Claim to our desired anti-luck conclusion? As I will show, the Instrumental Correctness Claim can be combined with one final

principle in order to yield our desired anti-luck conclusion. However, this final principle requires that we specify a subclass of correctness-fixing kinds—I shall call them “correctness-fixing *control* kinds”. Very roughly, a member of a correctness-fixing control-kind is typically under an agent’s control in the right sort of way. I will explain this notion more fully in a moment, but let’s first consider the final principle that we need in order to get to our desired anti-luck conclusion: For any correctness-fixing control kind K_{CFC} and any x ,

Copied-Luck Claim:

If it’s a matter of luck that x is a correct K_{CFC} , then it’s not the case that copying x in the relevant K_{CFC} respect(s) has potential bring it about that the copied version is a correct K_{CFC} .

The general idea behind the Copied-Luck Claim is this: If x is luckily correct, those who copy x are also leaving it up to luck as to whether they attain correctness; but then, since it’s left up to luck as to whether they are correct, their copying x must not be conducive to correctness.

Cast in this light, the Copied-Luck Claim seems rather intuitive. But its plausibility hinges greatly on what “the relevant K_{CFC} respect(s)” are. So let’s consider one respect that would make the Copied-Luck Claim true if deemed relevant:

R = the respect of whether the K_{CFC} came about in such a way that it would be lucky if it were correct.

If R is among the relevant respects referred to within the Copied-Luck Claim, then the Copied-Luck Claim is obviously true. After all, if x came about in such a way that it would be lucky for it to be correct, then it would also be lucky if any y that copies x in *that* respect turns out to be correct. But then it’s surely not the case that y ’s copying x in that respect has potential to bring about y ’s correctness, since copying x in that respect leaves it up to luck as to whether y is correct.¹⁸ So, it looks like the Copied-Luck Claim will be true provided that R is among the respects referred to within that claim.

¹⁸ See footnote 11 for a brief discussion on the relationship between luck and potential.

Is R among the relevant respects referred to in the Copied-Luck Claim? In what follows I want to suggest that, since we are now dealing with a specific class of control kinds, and since there is a close connection between control and luck, it is very plausible that R is a relevant respect.

Of course, all this depends greatly on what correctness-fixing control kinds K_{CFC} are supposed to be—a point I’ve not yet fully explained. The class of correctness-fixing control kinds just is the intersection of the class of correctness-fixing kinds with the class of control-kinds. Since we already know what correctness-fixing kinds are, the important question is what exactly control kinds are.

The basic idea is that a member of a control-kind is under an agent’s control in the right sort of way. Raising one’s arm, walking, talking, spelling a word, answering a question, and so forth, are all examples of control kinds. More precisely, control kinds are defined as follows:

Control Kinds

Kind K is a control-kind iff nearly all the members of K are attributable to an agent who can exercise non-managerial control over some of their features.

This definition obviously requires clarification of “non-managerial control.” But first off, let me emphasize that the sort of control I have in mind need not be voluntary control—that is, it need not be the sort of control that we can exercise just by willing an action, or by responding to practical reasons, etc. I shall assume that non-voluntary control is possible.¹⁹ But what is non-managerial control? Simply put, I take this to include any form of control that is *not* managerial. Following Pamela Hieronymi (2006), managerial control is the sort of control we typically have over ordinary objects, like books and billiard balls. We exercise managerial control when we perform actions

¹⁹ Hieronymi (2006) argues for this claim, but I find it obviously true. Machines, thermostats, insects, zombies, and many other creatures are intuitively able to control certain elements without volition. An example from Harry Frankfurt might help to justify the intuition:

“Consider the difference between what goes on when a spider moves its legs in making its way along the ground, and what goes on when its legs move in similar patterns and with similar effect because they are manipulated by a boy who has managed to tie strings to them” (1988, 78).

In the first case, the spider seems to exercise control without the robust form of volition under discussion. This non-voluntary control is present in the first case, but missing in the second case. I shall assume that volition and control can come apart in the way illustrated by the first case.

aimed at affecting ordinary objects so that they will accord with our thoughts about them. We can control these ordinary objects only because we know that they will behave according to certain regularities. But managerial control is clearly not the only type of control we can exercise. For example, the way in which we typically control our bodily limbs—raising arms, moving legs, and so on—is intuitively different from managerial control. The way in which we control our bodily limbs, in typical circumstances, is a paradigm of non-managerial control.²⁰

Above, I provided some examples of control kinds. But it may also be helpful to consider some kinds that fail to count as control kinds. Sneezing, convulsing, and flinching are not control kinds because the agent cannot exhibit non-managerial control over these events. Moreover, due to the nature of non-managerial control, most ordinary objects (e.g. a map of England) will fail to be members of control kinds.

It is now easier to see what the class of correctness-fixing control kinds is. It is the intersection of the class of correctness-fixing kinds with the class of control-kinds. Examples include spelling a word, answering a question, asserting a proposition, and many more (I shall later argue that belief is a correctness-fixing control-kind).

Correctness-fixing control-kinds are important, because it is plausible that they fix the following respect as one of the relevant respect(s) in which a K_{CFC} is to be copied:

R = the respect of whether the K_{CFC} came about in such a way that it would be lucky if it were correct.

In other words, it is plausible that what it is to be a K_{CFC} fixes R as one of the relevant K_{CFC} respects mentioned in Copied-Luck Claim. But why think this is so? The reason stems from a close connection between luck and control. In particular, many philosophers agree that luck entails a

²⁰ It's worth noting that that the distinction between managerial/non-managerial control does not necessarily coincide with the distinction between basic and non-basic actions. In typical cases, if I spell the word 'chiaroscuro' orally, and without using any devices (e.g. computers, typewriters, etc), then I've exhibited non-managerial control over that action. But my spelling of that word surely wouldn't count as a basic action, since it requires the performance of many other more basic actions (e.g. uttering the letter 'c', uttering the letter 'h', etc).

certain lack of control on the part of the agent for whom the state of affairs is lucky. In other words, it is widely agreed that a state of affairs S is lucky for a person P only if S is sufficiently far out of P's control (Riggs 2009, 220; Greco 1995, 83; Zimmerman 1993, 231; Statman 1991, 146).

To put this in our present terminology, we can say that

(*) If x's being a correct K_{CFC} is lucky for the agent to whom it's attributed, then x's being a correct K_{CFC} is outside of that agent's non-managerial control to at least some extent.

This claim is extremely plausible, and it provides a close connection between luck and correctness-fixing control kinds. As defined earlier, part of what it is for x to be a member of a correctness-fixing control kind is for x to have some features that the agent can control (non-managerially). And it is plausible that correctness is typically one such feature (e.g. I can control whether or not I spell 'chiaroscuro' correctly). But, given principle (*), we can now see that this control is sharply limited to instances of correctness that are not the product of luck. Thus, it appears that the correctness of a K_{CFC} is within the realm of the agent's control, unless that correctness is the product of luck. It thereby seems plausible that what it is to be a K_{CFC} itself fixes R as a relevant respect.²¹

Since R is plausibly among the relevant K_{CFC} respects referred to within the Copied-Luck Claim, it is therefore quite plausible that the Copied-Luck Claim is true. Moreover, the Copied-Luck Claim combined with the Instrumental-Correctness Claim jointly entails our desired anti-luck conclusion: For any correctness-fixing control kind K_{CFC} and any x which is a K_{CFC} , and some particular degree d*,

General Anti-Luck Conclusion:

If x is good qua K_{CFC} to degree d*, then it's not a matter of luck that x is a correct K_{CFC} .

Basically, this claim states that non-lucky correctness is a standard that a K_{CFC} must meet if it's to be good qua K_{CFC} to the appropriate degree.

²¹ The claims in this paragraph do not amount to a deductive argument for the conclusion that R is a relevant respect. The aim is merely to draw out a connection between control and luck that makes it intuitively plausible that R is relevant.

The General Anti-Luck Conclusion follows from various principles advanced in this paper. But for those who are skeptical of the argument leading up to this point, I'd like to provide some additional explanatory support in favor of the General Anti-Luck Conclusion. Consider three luck cases involving three different correctness-fixing kinds—spelling ‘chiaroscuro,’ answering a question, and being a map of England.

Spelling Bee: Sam is asked to spell the word ‘chiaroscuro’ at a spelling bee. But, having no idea how to spell this word, he starts blurting out letters at random. By chance, however, he just so happens to blurt out the letters ‘c-h-i-a-r-o-s-c-u-r-o,’ which means that, to Sam’s surprise, his spelling of ‘chiaroscuro’ was correct.

Trivial Pursuit: In a game of trivial pursuit, Sue was asked the question “What’s the state bird of Ohio?” Unfortunately, she misheard the question, and mistakenly thought she was asked about the state bird of *Iowa*. However, since Sue recalls seeing a number of cardinals during her latest trip to Iowa, she answers the question by saying “the cardinal”. As a matter of fact, the cardinal is the state bird of Ohio (but not Iowa). So her answer to the question “What is the state bird of Ohio?” is correct, even though she thought she was answering a different question to which her answer would have been incorrect.

Rand McNally: Occasionally, the printing presses at Rand McNally jam up and cause misprints. This occurred during the printing of a particular map of England called ‘M’. As result, M portrayed a river as veering westward in a place where it actually veered eastward. Everything else on M was correct, except for this one misprint. However, shortly after the misprint occurred, that river just so happened to change course. It now veers westward at the exact spot of the misprint. Due to this turn of events, M now portrays the river accurately, and is indeed a completely correct map of England.

Each of these three cases describes how a given item becomes a correct K as a matter of luck. In Spelling Bee the relevant K is the event-kind spelling ‘chiaroscuro’; in Trivial Pursuit the relevant K is the event-kind answering the question “What’s the state bird of Ohio?”; and in Rand McNally the relevant K is the object-kind map of England.²² I would now like to focus on an intuitive difference between these three cases, and to suggest that it is best explained by the General Anti-Luck Conclusion.

²² In the Rand McNally case there is a relevant event-kind that I am ignoring—namely, the kind map-making. We are instead focusing on the kind map, because, in the case at hand, it is a matter of luck that the map is a correct map. By contrast, it is false to say that it is a matter of luck that the map-making is a correct map-making.

In particular, it seems intuitively clear that the relevant K within Spelling Bee and Trivial Pursuit would have been better qua K had it been correct without being the product of luck. For example, Sam's spelling of 'chiaroscuro' would have been better qua spelling had he gotten it correct without getting lucky (e.g. by sounding out the word rather than blurting out letters at random). And similarly, Sue's answering the question "What's the state bird of Ohio?" would have been better qua answering of that question had it been correct without being the product of luck (e.g. had she not meant to answer a different question). However, the Rand McNally case seems importantly different from the other two. It seems clearly false to say that M would be better qua map of England had it been correct without being the product of luck. Since M is one-hundred percent correct as it is, it just seems false to say that it would've been better qua map of England had its correctness come about in the normal way. To see this difference in another light, consider that we would feel comfortable saying that M is a perfectly good map of England. But we wouldn't feel comfortable saying that Sam's spelling was a perfectly good spelling of 'chiaroscuro', or that Sue's answering was a perfectly good answering of that question.

This intuitive difference is easily explained by the General Anti-Luck Conclusion. Since this conclusion generalizes over all correctness-fixing control-kinds, it thereby applies to Sam's spelling and Sue's answering. And since these items are luckily correct, the General Anti-Luck Conclusion accurately predicts that they would fail to achieve a certain degree of goodness-qua- K_{CFC} . And in fact that's exactly what seems true of those cases. Intuitively, they would be better qua K_{CFC} if they were correct without being the product of luck. However, the General Anti-Luck Conclusion does not apply to the Rand McNally case, since the kind map of England is not a correctness-fixing *control*-kind. So, the General Anti-Luck Conclusion is compatible with our intuitive judgment that M is a perfectly good map of England even though it was the product of luck. In short, the fact that

the General Anti-Luck Conclusion explains the difference between these cases is additional support in its favor.

There are other potential ways of explaining these intuitions, but such ways are inferior to the General Anti-Luck Conclusion. For example, it will be pointed out that Spelling Bee and Trivial Pursuit each involve an *event* that is produced by luck, whereas Rand McNally involves an *object* that is produced by luck.²³ And perhaps it might be claimed that correctness-fixing *event* kinds are resistant to lucky correctness, whereas correctness-fixing *object* kinds are not.²⁴ However, this explanation leaves it completely mysterious as to why correctness-fixing event kinds should resist lucky correctness. What difference exists between events and objects that would illuminate the difference with regard to lucky-correctness? The answer is completely unclear. To be sure, the event/object distinction provides a difference between the above cases. But it is not a difference that relates in any interesting way to the issue of lucky correctness. By contrast, the General Anti-Luck Conclusion involves a subclass of control kinds and, as I noted earlier, there is a close connection between luck and control kinds stemming from a previous principle:

(*) If x's being a correct K_{CFC} is lucky for the agent to whom it's attributed, then x's being a correct K_{CFC} is outside of that agent's non-managerial control to at least some extent.

So, there is an interesting connection between luck and control kinds, but no similar connection seems to exist between luck and event kinds.

The last item on the agenda for this section is to show that belief is a correctness-fixing control-kind. We've already established that it is a correctness-fixing kind; but is it also a control-kind? The primary issue here is whether agents can exhibit non-managerial control over their beliefs—and this kind of control need not be voluntary (as mentioned earlier). It is clear that agents

²³ Similar distinctions such as process/product or acts/objects could be posed instead, but my reply would be the same as what follows.

²⁴ It's worth pointing out that, even if this explanation were adequate, it would have little impact on my overall argument. Since belief is better classified as an event (i.e. a static event), rather than an object, my overall claim that goodness-qua-belief resists lucky correctness would be equally plausible.

can exhibit *managerial* control over their beliefs. Borrowing an example from Richard Feldman (2000, 671), suppose that I now believe that the lights are on in my office. I can obviously exhibit a form of control over this belief simply by turning the lights off—as I turn off the lights, I correspondingly believe that they are off. In this way, agents can exercise managerial control over their beliefs. But can they also exhibit a form of control that’s *non-managerial*?

Perhaps the best way to see that we can exercise non-managerial control over our own beliefs is to consider how we often change our beliefs simply by reflecting on whether they are justified. In particular, if a person reflects upon one of her beliefs and finds that it is unjustified, then that reflection by itself is oftentimes sufficient to eliminate the belief. For instance, suppose I believe that the trip will take only one hour, and I believe this because that’s how long it took me last weekend when there was no traffic. But suppose further that I know that the trip this time will take place during heavy rush-hour traffic. Once I reflect on this latter fact in connection with my belief, I realize that my belief about the trip taking one hour is completely unjustified—and immediately upon realizing this, I give up my belief. Thus, simply by reflecting on whether a belief is justified, the subject can often exhibit a sort of control over that belief. Of course, agents are not *always* in a position to exhibit this sort of control—it is possible for an agent to continue believing something that she also believes to be unjustified. But in that case, she is being irrational with respect to at least one of these beliefs. So, more precisely, insofar as a subject is rational with respect to a given belief, she can exhibit a sort of control over that belief simply by reflecting on whether it is justified. Henceforth, let us call this “reflective control.”²⁵

Does reflective control count as non-managerial? It seems so. First off, it is obvious that reflective control differs from the sort of managerial control exhibited over my belief that the lights are on. In the light-switch case, I am clearly managing my belief by flipping the light-switch on and

²⁵ For discussion on reflective control, see Hieronymi (forthcoming) and Owens (2000). The label comes from Hieronymi.

off—I know that my beliefs will behave according to certain regularities, and that’s precisely why I am able to exercise managerial control over that belief. But nothing similar is happening when I change my belief by reflecting on whether it’s justified. In the traffic case, for example, I do not rely on any knowledge about whether my belief will behave according to certain regularities.

But second, if reflective control were a form of managerial control, then we would expect that one person could exercise reflective control over another person’s beliefs. More generally, whatever I can managerially control about myself, I can also managerially control about other people in the right circumstances—for example, just as I can managerially control my left arm by lifting it with my right arm, so too can I managerially control *your* left arm by lifting it with my right arm. However, nothing similar is true about reflective control. To be sure, I can *managerially* control your belief, for example, by flipping the light switch on and off. But I cannot control your belief simply by reflecting on whether your belief is justified. Thus, it seems that reflective control is clearly not a form of managerial control.

In short, we can exhibit a form of non-managerial control over our beliefs. And that means that belief is a correctness-fixing control kind. If that’s true, then we can finally arrive at the basis of the anti-luck condition for my theory of knowledge: For any x and some particular degree d^* ,

Specific Anti-Luck Conclusion:

If x is good qua belief to degree d^* , then it’s not a matter of luck that x is a correct belief.

This Specific Anti-Luck Conclusion combines with the account of knowledge in terms of goodness-qua-belief to degree $d_T (=d^*)$ so as to yield that a subject knows a proposition only if it’s not a matter of luck that her belief is correct. Thus, the subject in Russell’s stopped clock example fails to have knowledge because his belief is not good qua belief to the appropriate degree.

III.3. A Thomsonian Objection. Before proceeding any further, my account of knowledge must face an objection regarding one of its basic assumptions. I have so far assumed that belief is a goodness-fixing kind. But is this assumption correct? It would be devastating to my account of knowledge if it turns out that belief is not a goodness-fixing kind. In that case, there would be no such property as being good qua belief. And therefore, on my account, it would be impossible to have knowledge—an absurd conclusion! Judith Thomson has in fact claimed that belief is not a goodness-fixing kind (2008, 112). It is hard to ascertain exactly what Thomson’s argument is.²⁶ So, instead of focusing on her negative argument, let me propose my positive argument for the claim that belief is a goodness-fixing kind.

There is good reason to think that belief is a goodness-fixing kind. In particular, I propose that this claim provides the best explanation for a truism about belief:

Belief Truism:

For any proposition P, there is a particular way in which a true belief regarding P is better than a false belief regarding P.

Note that this claim does not state that true belief is better *all things considered* than false belief, but only in *a particular way*. Moreover, the principle says nothing about suspension of judgment; more specifically, it does not entail the highly contentious claim that, for any true proposition P, believing

²⁶ Although I cannot understand her argument, I will note that there is a notable tension in Thomson’s claims about what counts, and what fails to count, as a goodness-fixing kind. In particular, she denies that belief is a goodness-fixing kind, but accepts that certain relevantly similar kinds are goodness-fixing. As noted earlier, Thomson explicitly claims that the kind map of England is a goodness-fixing kind. The question, then, is why she denies that the kind belief is goodness-fixing. Thomson does not explain this disparity. But one potential answer is that she thinks there are clear standards that are required for something to be good qua map of England, but not for being good qua belief. Indeed, she explicitly claims that “being close-enough-to-correct” is a standard for being good qua map of England (2008, 85). But if that’s so, then why couldn’t *exact* correctness be a standard required for being good qua belief? What’s the difference between the kind belief and the kind map of England such that correctness could be a standard associated with one but not the other? This question is especially pressing given that some philosophers understand beliefs to be very much like maps (e.g. Lewis 1994 and Brandon-Mitchell and Jackson 1996). But Thomson provides no difference between maps and beliefs, and it’s not at all clear what relevant difference she could provide. After all, both beliefs and maps have representational content; and both of them are correct when their representational contents, in some sense, correspond to reality. To be sure, there is a difference in that maps are objects while beliefs are states, but I cannot understand why this should make any difference with regard to whether something is a goodness-fixing kind. Thus, I submit that Thomson’s dismissal of belief as a goodness-fixing kind is highly suspicious, given that she concedes that maps are goodness-fixing kinds.

P is better in a particular way than suspending judgment with regard to P.²⁷ And lastly, I would like to add that Belief Truism holds true even when we are comparing an unjustified true belief with a justified false belief. In other words, I think it's correct to say that there is a particular way in which an unjustified true belief is better than a justified false belief. This becomes apparent once we realize that believing P just is taking P *to be true* in the right sort of way.²⁸ Once we recognize this point about the nature of belief, it becomes obvious that there's a particular way in which an unjustified true belief would be better than a justified false belief.

But what explains Belief Truism? More specifically, what is that particular way mentioned within Belief Truism? It cannot be that true belief is more useful or beneficial for the believer than false belief. After all, it is sometimes more useful and beneficial for the believer to have a false belief about P than to have a true belief about P. For example, it may be more useful and beneficial to have a false belief about your bleak chances of survival than to have a true belief about those chances.

The explanation also cannot come from the fact that Belief Truism partly compares a true proposition with a false proposition. In other words, we cannot explain Belief Truism by claiming that true belief is better than false belief in that the former always involves a true proposition. After all, there are obviously false claims that are analogous to Belief Truism in respect of comparing true propositions to false propositions. Consider the following:

Doubt Claim:

For any proposition P, there is a particular way in which doubting P while P is true is better than doubting P while P is false.

This analogous claim about doubt seems clearly mistaken. Doubting false propositions is typically the better alternative. And there seems to be no particular way in which it would be better to doubt

²⁷ For example, having a true belief about how many names are in the NYC phonebook might not be better in any way than suspending judgment about that fact.

²⁸ Supposing and conjecturing might be ways of taking P to be true, but not in the same way as believing.

true propositions. Thus, the explanation of Belief Truism cannot come from the fact that it partly compares true propositions to false proposition.

How then can Belief Truism be explained? I think it is clear that the particular way referred to within Belief Truism is that of being better *as a belief*, or better *qua belief*. In particular, we can explain Belief Truism by claiming that, for any proposition P, a true belief regarding P is better *qua belief* than a false belief regarding P. This claim entails, and explains, Belief Truism. And if we accept it, then we must accept that there is such a property of being better qua belief and, relatedly, of being good qua belief. In other words, we must accept that belief is a goodness-fixing kind. Therefore, the assumption I have made so far in this paper seems to be very plausible.

IV. The Gettier and Value Problems.

The goal up to this point has been to argue for the following account of knowledge:

Exemplary Belief Analysis

S knows P iff S believes P and this belief is good qua belief to at least degree d_T .

Degree d_T is the degree of goodness-qua-belief that has correctness as a standard (i.e. a standard that must be met for a belief to qualify as good qua belief to that degree). It was also argued that, since belief is a correctness-fixing control kind, this same degree of goodness-qua-belief sets non-lucky correctness as an additional standard. We might plausibly suppose that these are fairly high standards, and, correspondingly, that degree d_T is a fairly high degree of goodness-qua-belief. Nothing that I say in this paper requires this, or relies on it, but it will be useful for the discussions below.

The aim of this final section is to explain how this account of knowledge can be applied to both the Gettier and Value problems. I shall discuss these problems in turn and consider objections as they arise.

IV.1 The Gettier Problem. As noted at the outset, I take the Gettier Problem to be the task of stating informative, necessary, and sufficient conditions for knowledge, ones that explain why knowledge excludes the kind of belief described within Gettier-style cases. The main question I would like to focus on is this: Does my account deal adequately with Gettier-style counterexamples, and if so, how?

The way in which my account deals with Gettier-style counterexamples is simple: My account entails that, for a subject to have knowledge, it must not be a matter of luck that her belief is true. And Gettier-style counterexamples are cases in which it *is* a matter of luck that the subject's belief is true. So, the Gettier subject lacks knowledge. Thus, on my account, the diagnosis of Gettier-style cases will be plain and simple across the board.

Of course, one might initially find this approach to the Gettier Problem unsatisfactory. Compare my analysis with the following:

No-Luck Analysis

S knows that P iff S has a true belief that P and it is not a matter of luck that this belief is true. In the absence of a further account of luck (i.e. of epistemic luck), this No Luck analysis would be unsatisfactory. However, it is not inadequate on account of being untrue, but on account of the fact that it's not sufficiently informative. Nearly everyone accepts the platitude that knowledge excludes luck (e.g. see Dancy 1985, 134 and Pritchard 2005, 1-6). And the No-Luck Analysis simply analyzes knowledge in terms of that platitude. But epistemologists typically want an analysis of knowledge that can explain why that platitude is true in the first place. We want to know why knowledge excludes luck. All by itself, the No-Luck Analysis does not tell us this. If we were to supply a detailed account of epistemic luck (e.g. as in Pritchard 2005), and substitute it into the No-Luck Analysis, then this problem could be avoided. But in the absence of an account of epistemic luck, the No-Luck analysis is not sufficiently informative.

Does my account of knowledge fall prey to this same problem? No. Although I do not provide an account of epistemic luck, my account explains the anti-luck platitude in a different way. As we've already seen, my analysis of knowledge excludes luck by virtue of the General Anti-Luck Conclusion. Let us recall that principle: For any correctness-fixing control kind K_{CFC} and any x which is a K_{CFC} , and some particular degree d^* ,

General Anti-Luck Conclusion:

If x is good qua K_{CFC} to degree d^* , then it's not a matter of luck that x is a correct K_{CFC} .

My account explains the anti-luck platitude in the following way: Knowledge is analyzed in terms of goodness-qua-belief to degree $d_T (=d^*)$, and this property is a member of a general class of properties that exclude lucky correctness (i.e. the class containing all properties of the form goodness-qua- K_{CFC} to d^*). Recall that other properties, such as goodness-qua-spelling and goodness-qua-answering, also exclude lucky correctness by virtue of the relevant applications of the General Anti-Luck Conclusion. The way in which goodness-qua-belief excludes luck is simply another application of this general principle. Thus, my account allows us to explain the anti-luck platitude as a particular instance of a general kind of anti-luck.

So far, I've dealt with the Gettier Problem on a very general level. But my account will also have to face specific objections arising from particular Gettier-style cases. Let's consider a variation on Carl Ginet's barn-façade example.²⁹ In this variation, we are asked to compare two different cases in which a subject forms a true belief that there is a barn.

In the first case, the subject is in barn-façade country (unknownst to him) and he forms the belief that there is a barn (call it 'B1') on the basis of what visually appears to him to be a barn. His justification for B1 is then increased by other factors, for example, by his gathering extra evidence for B1 or by checking for defeaters (e.g. he may have talked to a farmer who claims that

²⁹ For the print version of this example, see Goldman (1976, 772-73).

the barn is real). But, as it turns out, B1 is luckily correct despite its high degree of justification—since he is in barn façade country, he had encountered many fake barns along the way, and this one just so happens to be the only real one in that area.

In the second case, the subject is *not* in barn façade country. He forms the belief that there is a barn (call it ‘B2’) purely on the basis of the visual appearance of a barn, just as we normally form perceptual beliefs. But in this case he did not check for defeaters or gather additional evidence. Importantly, however, in this second case B2 is true without being the product of luck—there were no fake barns anywhere in the area.

The important point about these two cases is that it seems very plausible to say that B2 counts as knowledge even though it is not as highly justified as B1. Moreover, given our natural reactions to cases like the first one, I think some people will be initially inclined to say that B1 does not count as knowledge. After all, B1 was deemed to be luckily true. However, a problem may arise for my account once we notice that B1 seems to be at least as good qua belief as B2, due to the fact that B1 has such a high level of justification. These intuitions are a problem for my account in the following way: Since B1 is at least as good qua belief as B2, and B2 counts as knowledge, my account entails that B1 is good qua belief to at least degree d_T . But if that’s so, then being good qua belief to degree d_T does not exclude lucky correctness, since B1 is good to that degree but is luckily correct. So, my account would not be able to explain why Gettier beliefs like B1 fail to count as knowledge.

In reply, let me point out that the first case, involving B1, is deceptive. If we grant that B1 has greater justification than B2, this is presumably because the subject who holds B1 has either checked for defeaters or acquired additional evidence while the subject who holds B2 does neither. But, if this is so, then it’s not at all clear why we should think either that B1 fails to count as knowledge or that it’s a matter of luck that B1 is true. For example, if we imagine that the subject

who holds B1 asks a farmer about the status of the barn³⁰ or looks at the barn from different angles, then it seems very plausible that B1 does count as knowledge and that it's not a matter of luck that B1 is true. The fact that there were fake barns in the area is completely irrelevant, given that the subject has performed these additional checks about *this particular barn*. On the other hand, if we alter the examples so that the justification for B1 is equal to that of B2, then it no longer seems true that B1 is as good qua belief as B2. In short, I do not think this general way of describing cases (as modeled above) will be successful in supplying counterexamples to my account.

The above claims lead one to wonder how exactly justification relates to my account of knowledge. I have not been concerned with the supposed entailment from knowledge to justification, but one way of explaining that putative entailment would be to claim that justification is a property entailed by goodness-qua-belief *to some degree or other*. More precisely, we could claim that justification is a standard that a belief must meet if it's to be good qua belief to any degree whatsoever. However, in order to maintain this position I must insist that justification is not sufficient for goodness-qua-belief to degree d_T . Otherwise, justification would entail truth.³¹ Thus, for present purposes, I shall take justification to be a property entailed by goodness-qua-belief to some degree or other, though it does not entail goodness-qua-belief to degree d_T . Nothing I say in this paper will rely on this view of justification, but it will be useful when discussing the Value Problem to have a rough conception of its relationship to goodness-qua-belief.

IV.2. The Value Problem. At the outset I explained the Value Problem in its most basic form. That way of stating the problem is typically called “The Primary Value Problem,” which is

³⁰ Of course, we could imagine that all farmers in that region would lie about the status of their barn facades, but in that case I don't see why asking farmers in this region about their barns would increase one's justification.

³¹ No doubt, there are other ways in which my account could potentially explain the entailment from knowledge to justification. The view proposed here is merely aimed at giving us a general idea of how justification might fit into the picture. I am open to the possibility that further examination of the notion of goodness-qua-belief might yield a different account of this relation.

Primary Value Problem: The task of explaining why knowledge is more valuable than mere true belief.

But as some epistemologists have argued (e.g. Pritchard 2010), there are two other forms of the Value Problem, both of which I shall discuss:

Secondary Value Problem: The task of explaining why knowledge is more valuable than that which is necessary but insufficient for knowledge.

Tertiary Value Problem: The task of explaining why knowledge has a different *kind* of value from that which is necessary but insufficient for knowledge.

These value claims seem to compare the value of properties, such as the property of having knowledge versus the property of having mere true belief. But what does it mean to say that one property is more valuable than another? Although this is rarely clarified, I believe that we are to understand these claims in terms of the value that each property adds to its possessor. Thus, for example, the Primary Value Problem is basically claiming that knowledge adds more value to its possessor than the property of having mere true belief.

In what follows, I explain how my account deals with these three forms of the Value Problem. But since the Primary Value Problem is a specific instance of the Secondary Value Problem, any solution to the latter will also solve the former. Thus, I shall skip straightaway to the Secondary Value Problem.

IV.2.a. Secondary Value Problem. In order to address this, we need to compare knowledge with some properties that, according to my account, are plausibly necessary but insufficient for knowledge. Consider the following two: (1) being good qua belief to some degree or other, and (2) being a justified true belief. These are the relevant properties that I would need to show have lesser value than knowledge.

Can my account support the greater value of knowledge over (1) and (2)? Yes it can. On this account, knowledge guarantees a greater degree of value from what is guaranteed by (1). Notice that anything that has (1) could have a very low degree of goodness-qua-belief—indeed, the lowest

degree possible. But my account states that knowledge requires the belief in question to be good qua belief to at least degree d_T , and this guarantees a higher degree of goodness-qua-belief than what is guaranteed by (1).³² Therefore, knowledge requires a greater degree of goodness-qua-belief from that involved in (1).

By the same token, my account reveals that knowledge has a kind of value that is absent from (2). This is because knowledge requires goodness-qua-belief to degree d_T , whereas (2) does not. Thus, since knowledge includes (2) but also has an additional kind of value, we can thereby cite that additional value in explaining the greater value of knowledge.

One might challenge this explanation by claiming that, for some K's, being good qua K is worse overall than not being good qua K. For instance, being good qua hit-man is worse overall than not being good qua hit-man. However, while this may be correct, the objection itself mistakenly presupposes that knowledge is of greater *overall* value than that of (2). But in fact this is not true. There are possible worlds in which knowledge is not better overall than merely having a justified true belief. Suppose there's an evil demon bent on torturing all and only those who possess knowledge of some proposition P—he will torture you if you know P but not if you merely have a justified true belief that P. Clearly, in this possible world, having knowledge that P is not better overall than having a justified true belief that P.

How, then, are we to interpret the phrase “is more valuable than” within the Secondary Value Problem? For reasons just mentioned, I do not think we can interpret it as meaning “has more *overall* value than.” I think it is more plausible to interpret the phrase in one of two ways: “is better *in a particular way* than,” or “is better than... *other things being equal*.” Either of these interpretations would stave off the evil demon counterexample. Even in the evil demon world, it is plausible that there is a particular way in which knowledge is better than mere justified true belief; it

³² Note that, if degree d_T were the lowest degree of goodness-qua-belief, then (1) would not be insufficient for knowledge. And this means that (1) would be irrelevant to the Secondary Value Problem.

is also plausible that knowledge is better than mere justified true belief other things being equal. But, on either interpretation, my account seems able to explain the greater value of knowledge. After all, on my view, knowledge is better *qua belief* than mere justified true belief. This is guaranteed by the fact that knowledge requires goodness-qua-belief to at least degree d_T , whereas mere justified true belief does not. And being better *qua belief* seems to constitute a way of being better. It also seems to make a belief better other things being equal. Thus, the account in question shows great promise in resolving the Secondary (as well as the Primary) Value Problem.

IV.2.b. Tertiary Value Problem. My approach to the Tertiary Value Problem is different from my approach to the Primary and Secondary Value Problems. In particular, I believe that the Tertiary Value Problem is not a genuine problem at all. But, as I'll explain, my account of knowledge can easily undermine the reasons for thinking it is a genuine problem.

First of all, let me explain why my account could not solve the Tertiary Value Problem (if indeed it were a genuine problem). My account does not predict that knowledge should have a different *kind* of value from all properties that are necessary but insufficient for knowledge. Recall property (1)—being good *qua belief* to some degree or other. This property is necessary but insufficient for knowledge. But the kind of value attributed by this property seems to be the same kind of value attributed within my account of knowledge. Thus, my account does not appear to have the resources to solve the Tertiary Value Problem.

Fortunately, there is no good reason to think that the Tertiary Value Problem is a genuine problem. Why should we think that knowledge has a different kind of value from *every* property that's necessary but insufficient for knowledge? Duncan Pritchard, the main proponent of the Tertiary Problem, motivates it in two different ways, both of which I reject.

Pritchard's first argument for the Tertiary Value Problem is as follows:

...[I]f one regards knowledge as being more valuable than that which falls short of knowledge merely as a matter of degree rather than kind, then this has the effect of

putting knowledge on a kind of continuum of value with regard to the epistemic, albeit further up the continuum than anything that falls short of knowledge. The problem with this 'continuum' account of the value of knowledge, however, is that it fails to explain why the long history of epistemological discussion has focused specifically on the stage in this continuum of value that knowledge marks rather than some other stage (such as a stage just before the one marked out by knowledge, or just after) (Pritchard 2010, 7-8).

Before addressing this argument, let's first consider the kind of continuum view that Pritchard has in mind. He seems to envision a value continuum consisting of items like (i) belief, (ii) true belief, (iii) justified true belief, and (iv) knowledge. And he thinks that the proponent of this 'continuum' view is simply ranking knowledge at a higher place above other epistemic values like (i)-(iii), without explaining why epistemologists have been so concerned with that particular place on the continuum.

But my account does not have the effect of placing knowledge on a value continuum along with notions like (i)-(iii). Indeed, my account predicts that knowledge has a different *kind* of value from (i)-(iii), because knowledge requires goodness-qua-belief to at least degree d_T , whereas (i)-(iii) do not. However, my account *does* place knowledge on a different sort of value continuum, one that consists of all the varying degrees of goodness-qua-belief. Knowledge is placed at degree d_T or higher (which, as I have supposed, is a fairly high degree of goodness-qua-belief). But Pritchard's objection does not succeed against this sort of value continuum. In particular, he objects that the continuum view "fails to explain why the long history of epistemological discussion has focused specifically on the stage in this continuum of value that knowledge marks...." But my account does not fail to explain this. Degree d_T is the lowest degree of goodness-qua-belief that guarantees truth. And epistemological discussion has focused on this particular degree (and higher) because epistemologists are interested in how we can "get at the truth." Focusing on some lower degree would not cater to this interest, because a lower degree would not guarantee truth. Thus, the needed explanation called for by Pritchard falls right out of my account of knowledge.

Pritchard's second argument can be summed up with two premises. He thinks (P1) that any property that is necessary but insufficient for knowledge will be merely instrumentally valuable. But

he also thinks (P2) that knowledge has some non-instrumental value. It follows from (P1) and (P2) that knowledge has a different kind of value from that which is necessary but insufficient for knowledge (Pritchard 2010, 8).

However, Pritchard provides very little reason for accepting (P1) in full generality, and my account actually provides reason to think it is false. Recall a principle that I advanced in section III.3: For any proposition P, a true belief regarding P is better *qua belief* than a false belief regarding P. This claim holds for *any* proposition, and it is important to notice that there are many trivial propositions for which it is of no instrumental value for a subject to truly believe—it is of no instrumental value for me to truly believe that there are exactly *n* blades of grass on the neighbor's lawn. Therefore, if this principle is correct, then a true belief can be better *qua belief* than its false counterpart, even if that true belief is no more instrumentally valuable than its false counterpart. It follows that there can be increases in value-*qua-belief* even if there is no increase in instrumental value. For example, suppose I hold a false belief B in a trivial proposition, like the one mentioned above. Merely by changing the truth-value of B from false to true (and without changing anything else about B) we can increase B's value-*qua-belief* without thereby adding any instrumental value to B. Now let's focus on the precise quantity of value-*qua-belief* that has increased as a result of this change—call this quantity of value 'V'. On my account, it is necessary but insufficient for knowledge that the relevant belief has V. (After all, true belief is necessary but insufficient for knowledge, and V is entailed by any true belief in accordance with the above-mentioned principle.) But, as already mentioned, V is not merely instrumentally valuable, because it could be possessed by a trivial belief, like B, that has no instrumental value whatsoever. Therefore, on my account of knowledge, V is necessary but insufficient for knowledge even though it is not merely instrumentally valuable. My account entails that Pritchard's (P1) is false.

In short, there appears to be no good reason to think the Tertiary Value Problem is a genuine problem. Pritchard provides two arguments for the claim that knowledge has a different kind of value from anything that falls short of it. But my account undermines the appeal of both of his arguments. So, the fact that my account cannot solve this pseudo-problem is no mark against it.

V. Conclusion.

The theory of knowledge advanced in this paper is very simple: To know a proposition just is to believe it and for this belief to be good qua belief to the appropriate degree. Although this account may initially seem *all too* simple, I have argued that it entails both a factivity condition and an anti-luck condition. Factivity and anti-luck can be seen as standards that a belief must meet if it's to be good qua belief to the appropriate degree. In the final section of this paper, I argued that this simple account of knowledge shows promise in resolving both the Gettier and Value problems.

Conclusion

I have argued in this dissertation that *knowledge* is a thick concept, and that, because thick concepts exhibit certain general features, *knowledge* must also possess these features. The result is a new approach to the Gettier and Value problems. In closing, let me review two general features of thick concepts, and then tie together several points from each individual chapter.

One general feature exhibited by positive thick concepts is their connection to the concept *good in a way*. It was argued in chapter 1—“How Are Thick Terms Evaluative?”—that positive thick concepts conceptually entail *good in a way*. Although this was never made explicit, I believe that many of the considerations brought in favor of that conclusion can apply similarly in the case of knowledge. For example, it was pointed out that the following sentences sound odd:

- # Nancy is generous and she’s not good in any way.
- ? Mark is generous {and/but} he’s good in a way.

The view that positive thick concepts conceptually entail *good in a way* was seen as the best way of explaining this kind of linguistic data. But it is worth pointing out that similar data is available with regard to the verb ‘knows’.

- # Jon knows that p, and his position with regard to this fact is not good in any way.
- ? Sue knows that p {and/but} her relation to this fact is good in a way.¹

The infelicity of these sentences supports the view that *knowledge* conceptually entails *good in a way*.

And this is just what we should expect if *knowledge* is a thick concept.

In chapter 3, “How to Analyze a Thick Concept,” the connection between thick concepts and *good in a way* (or *bad in a way*) was taken as something that cries out for explanation. There, it was argued that the Qua Schema can provide a particularly elegant explanation of this connection.

According to the Qua Schema, thick concepts can be analyzed in terms *good qua* N_{SUB} and *bad qua*

¹ Of course, there are contexts in which these sentences would sound fine. But, as argued in chapter 1, this can be explained by the fact that ‘good in a way’ is context sensitive.

N_{SUB} . And since N_{SUB} must be a goodness-fixing kind (or a badness-fixing kind), the Qua Schema thereby guarantees that every thick concept will conceptually entail *good in a way* (or *bad in a way*).

The Qua Schema also explains a second general feature of thick concepts—namely, the fact that thick concepts “hold together” an evaluation plus a description. Many philosophers believe these two aspects of a thick concept are so closely entwined that they cannot be analyzed into distinct types of contents. The Qua Schema diverges from this view by allowing for the analyzability of thick concepts, but it respects the intuitive sense in which evaluation and description are “entwined”. According to the Qua Schema, a thick concept’s evaluative and non-evaluative contents are combined by virtue of the fact that N_{SUB} is conceptually entailed by the totality of the thick concept’s non-evaluative content.

In many ways, chapter 2 of this dissertation (i.e. “Knowledge as a Thick Concept”) was also concerned with the way in which thick concepts combine evaluation and description. There, I approached this issue by asking why ad hoc conjunctions (e.g. ‘gred’ and ‘wronday’) do not express thick concepts. These terms express concepts that are mere conjunctions of evaluative and non-evaluative content. But, intuitively, mere conjunction is not the way in which thick concepts combine such contents. As noted, the Qua Schema is my preferred way of characterizing the way in which these contents are combined in a thick concept.

We have so far reviewed two general facts about thick concepts—(i) the fact that thick concepts combine evaluation and description, and (ii) the fact that they conceptually entail *good/bad in a way*. And I argued in chapter 3 that these two facts are best explained by the hypothesis that each thick concept can in principle be analyzed according to the Qua Schema. If I am correct in advancing this hypothesis, then it can be employed as the first premise of a very simple argument:

P1: For any concept *C*, if *C* is a thick concept, then *C* can in principle be analyzed according to the Qua Schema.

P2: *Knowledge* is a thick concept.

C: Therefore, *knowledge* can in principle be analyzed according to the Qua Schema.

In many ways, this argument is the backbone of the dissertation. Each step was defended in at least one of the four chapters. P1 was defended in chapter 3 (and indirectly in chapter 1). P2 was defended in chapter 2. And the conclusion was defended in chapter 4.

Although chapter 2 was partly concerned to defend P2, this chapter was also aimed at explaining why the Gettier Problem arises to begin with. I argued that this can be explained by the fact that *knowledge* is a thick concept. More precisely, I argued that there is a general problem about thick conceptual analysis, and that the Gettier Problem is a particular instance of this general problem. This naturally suggests that a new direction should be taken with regard to research into the Gettier problem, a direction that inquires into the nature of thick concepts more generally. Chapters 3 and 4, together, are my contribution to this general approach to the Gettier Problem. Chapter 3 establishes my view about the nature of thick concepts more generally (as portrayed by the Qua Schema), and chapter 4 applies this general view to the particular case of knowledge.

The simple argument laid out above would be unsound if *knowledge* could not be analyzed in accord with the Qua Schema. But chapter 4, “Knowledge as Exemplary Belief,” aimed at showing that *knowledge* can be analyzed according to this schema. There, I argued (on grounds independent of P1 and P2) that knowledge consists in having a belief that is good qua belief to the appropriate degree. Generally speaking, an item must meet certain standards if it is to qualify as being good qua K. And I argued that non-lucky correctness is a standard that a belief must meet if it’s to be good qua belief (to the appropriate degree). If this is right, then my account of knowledge can resolve the

Gettier Problem, since Gettier-style cases are ones in which it is a matter of luck that the subject acquires a correct belief.

The fourth chapter of this dissertation aimed not merely at resolving the Gettier Problem, but also the Value Problem—i.e. the problem of explaining the greater value of knowledge over lesser qualities, like mere true belief. It is highly important that this dissertation tackle the Value problem, in addition to the Gettier Problem. Attempted solutions to the Gettier Problem are often criticized by the supposed fact that they render the value of knowledge mysterious. Moreover, it is natural to expect that the value of knowledge could be illuminated by an inquiry into knowledge as a thick *evaluative* concept. And this was the final conclusion of this dissertation. The appropriate degree of goodness-qua-belief, not only excludes the kind of belief described in Gettier-style cases, but it also adds additional value to lesser qualities, like mere true belief.

This dissertation has aimed to show that *knowledge* is importantly thick. The inclusion of *knowledge* among thick concepts, as well as a better understanding of what thick concepts are, has led to new insights into the nature and value of knowledge.

Bibliography

- Alston, W. 1989. "Justification and Knowledge," in *Epistemic Justification: Essays in the Theory of Knowledge*, Ithaca, NY: Cornell University Press, 172-82.
- Anscombe, E. 1958. "Modern Moral Philosophy," *Philosophy* 33, 1-19.
- 1979. "Prolegomenon to a Pursuit of the Definition of Murder," *Dialectics and Humanism* 4, 73-77.
- Atlas, J. and S. Levinson. 1981. "It-Clefts, Informativeness, and Logical Form: Radical Pragmatics (Revised Standard Version)," in P. Cole, ed., *Radical Pragmatics*, New York: Academic Press, 1-61.
- Audi, R. 1995. "Memorial Justification," *Philosophical Topics* 23(1), 31-45.
- 2001. "An Internalist Theory of Normative Grounds," *Philosophical Topics* 29, 19-46.
- 2007. "Justifying Grounds, Justified Beliefs, and Rational Acceptance," in M. Timmons, J. Greco, and A. Mele, eds., *Rationality and the Good*, New York: Oxford University Press, 222-47.
- Ayer, A.J. 1956. *The Problem of Knowledge*, London: Macmillan and Company Limited.
- Axtell, G. and A. Carter. 2008. "Just the Right Thickness: A Defense of Second-Wave Virtue Epistemology," *Philosophical Papers* 37, 413-34.
- Bach, K. 1999. "The Myth of Conventional Implicature," *Linguistics and Philosophy* 22, 327-66.
- Battaly, H. 2008. "Metaethics Meets Virtue Epistemology: Salvaging Disagreement about the Epistemically Thick," *Philosophical Papers* 37, 435-54.
- Blackburn, S. 1984. *Spreading the Word*. Oxford: Clarendon Press.
- 1992. "Through Thick and Thin," *Proceedings of the Aristotelian Society*, suppl. vol. 66: 285-99.
- 1998. *Ruling Passions*. Oxford: Oxford University Press.
- Blutner, R. 2004. "Pragmatics and the Lexicon," in L. Horn and G. Ward, eds., *The Handbook of Pragmatics*, Oxford: Blackwell, 488-514.
- Brewer, T. 2009. *The Retrieval of Ethics*, Oxford: Oxford University Press.
- Brower, B. 1988. "Virtue Concepts and Ethical Realism," *The Journal of Philosophy* 85, 675-93.
- Burton, S. 1992. "'Thick' Concepts Revisited," *Analysis* 52, 28-32.

- Conee, E. and R. Feldman. 2001. "Internalism Defended," *American Philosophical Quarterly* 38, 1-18.
- Chierchia, G. and S. McConnell-Ginet. 1990. *Meaning and Grammar: An Introduction to Semantics*. Cambridge, Mass: MIT Press.
- Dancy, J. 1985. *Introduction to Contemporary Epistemology*. Oxford : Blackwell.
- . 1995. "In Defense of Thick Concepts," *Midwest Studies in Philosophy* 20, 263-79.
- Davis, W. 2010. "Implicature," in E. Zalta, ed., *The Stanford Encyclopedia of Philosophy (Winter 2010 Edition)*, URL = < <http://plato.stanford.edu/archives/win2010/entries/implicature/>>.
- DePaul, M. 2009. "Ugly Analyses and Value," in Haddock, Millar, and Pritchard, eds., *Epistemic Value*, Oxford: Oxford University Press, 112-38.
- Elgin, C. 2005. "Williams on Truthfulness," *The Philosophical Quarterly* 55, 343-52.
- . 2008. "Trustworthiness," *Philosophical Papers* 37, 371-87.
- Eklund, M. 2011. "What Are Thick Concepts?" *The Canadian Journal of Philosophy* 41, 25-50.
- Engel, M. 1992. "Is Epistemic Luck Compatible with Knowledge," *The Southern Journal of Philosophy* 30, 59-75.
- Feldman, R. 2000. "The Ethics of Belief," *Philosophy and Phenomenological Research* LX, 667-695.
- Foot, P. 1958. "Moral Arguments," *Mind* 67, 502-13.
- Frankfurt, H. 1988. "The Problem of Action," in *The Importance of What We Care about*, Cambridge, UK: Cambridge University Press, 69-79.
- Gazdar, G. 1979. *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.
- Geach, P. 1956. "Good and Evil," *Analysis* 17, 32-42.
- Gettier, E. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23, 121-23.
- Geurts, B. 1998. "Mechanisms of Denial," *Language* 74, 274-307.
- Gibbard, A. 1992. "Thick Concepts and Warrant for Feelings," *Proceedings of the Aristotelian Society*, suppl. vol. 66, 285-99.
- . 2003. "Reasons Thin and Thick," *The Journal of Philosophy* 100, 288-304.
- Ginet, C. 1988. "The Fourth Condition," in D. Austin, ed., *Philosophical Analysis*, Kluwer Academic Publishers, 105-17.
- Goldman, A. 1976. "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73, 771-91.

- 1986. *Epistemology and Cognition*, Cambridge, Mass.: Harvard University Press.
- Greco, J. 1995. "A Second Paradox Concerning Responsibility and Luck," *Metaphilosophy* 26, 81-96.
- 2003. "Knowledge as Credit for True Belief," in DePaul and Zagzebski, eds., *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford: Oxford University Press.
- Greenberg, J. 1966. *Language Universals*. The Hague: Mouton and Co.
- Grice, H.P. 1989. *Studies in the Way of Words*. Cambridge, Mass: Harvard University Press.
- Hare, R.M. 1963. *Freedom and Reason*. Oxford: Oxford University Press.
- 1967. "Geach: Good and Evil," in Foot, ed., *Theories of Ethics*, Oxford: Oxford University Press.
- 1995. "Inhaltsreiche ethische Begriffe und die Unterscheidung zwischen Tatsachen und Werten: Replik auf Millgram," in C. Fehige and G. Meggle, eds., *Zum moralischen Denken*, Frankfurt a. M.: Suhrkamp, 354-88.
- 1997. *Sorting Out Ethics*. Oxford: Clarendon Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hazlett, A. Forthcoming. "The Myth of Factive Verbs," *Philosophy and Phenomenological Research*.
- Hieronymi, P. 2006. "Controlling Attitudes," *Pacific Philosophical Quarterly* 87, 45-74.
- Forthcoming. "Two Kinds of Agency," in L. O'Brien, ed., *Mental Agency*, Oxford: Oxford University Press.
- Horn, L. 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.
- Hurka, T. and D. Elstein. 2009. "From Thick to Thin: Two Moral Reduction Plans," *The Canadian Journal of Philosophy* 39, 551-36.
- Hurley, S. 1989. *Natural Reasons*, Oxford: Oxford University Press.
- Hursthouse, R. 1995. "Normative Virtue Ethics," in Roger Crisp, ed., *How Should One Live?* Oxford: Oxford University Press, 19-36.
- Jackson, F. and D. Brandon-Mitchell. 1996. *Philosophy of Mind and Cognition*, Oxford: Blackwell.

- Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Kadmon, N. 2001. *Formal Pragmatics: Semantics, Pragmatics, Presupposition, and Focus*, Oxford: Blackwell Publishers.
- Kornblith, H. 2009. "Knowledge Needs No Justification," in Q. Smith, ed., *Epistemology: New Essays*, Oxford: Oxford University Press, 5-23.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*, Cambridge, UK: Cambridge University Press.
- Laurence, S. and E. Margolis. 1999. "Concepts and Cognitive Science," in Laurence and Margolis, eds., *Concepts: Core Readings*, Cambridge, Mass.: MIT Press, 3-81.
- Levinson, S. 1983. *Pragmatics*. Cambridge, UK: Cambridge University Press.
- 2000. *Presumptive Meanings*, Cambridge, Mass: MIT Press.
- Lewis, D. 1994. "Lewis, David: Reduction of Mind," in S. Guttenplan, ed., *A Companion to the Philosophy of Mind*, Oxford: Blackwell, 412-431.
- 1996. "Elusive Knowledge," *Australasian Journal of Philosophy* 74, 549-67.
- 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 S. Holtzman and C. Leich, eds., *Wittgenstein: To Follow a Rule*, London: Routledge.
- Murdoch, I. 1971. *The Sovereignty of Good*, New York: Schocken.
- Nagel, T. 1969. "Sexual Perversion," *The Journal of Philosophy* 66, 5-17.
- Owens, D. 2000. *Reason without Freedom: The Problem of Epistemic Normativity*, London: Routledge.
- Payne, A. 2005. "A New Account of Thick Concepts," *The Journal of Value Inquiry* 39, 89-103.
- Plantinga, A. 1993. *Warrant and Proper Function*, Oxford: Oxford University Press.
- Plato. 1961. "Meno," in Hamilton and Cairns, eds., Guthrie, trans., *The Collected Dialogues of Plato*, Princeton, NJ: Princeton University Press.
- Platts, M. 1979. *Ways of Meaning*, Cambridge, Mass.: MIT Press.
- Potts, C. 2007. "Into the Conventional-Implicature Dimension," *Philosophy Compass* 2, 665-79.

- Pritchard, D. 2005. *Epistemic Luck*, Oxford: Oxford University Press.
- 2007. "Anti-Luck Epistemology," *Synthese*, 277-97.
- Priest, G. 1997. "Sexual Perversion," *Australasian Journal of Philosophy* 75, 360-72.
- Primoratz, I. 1997. "Sexual Perversion," *American Philosophical Quarterly* 34, 245-58.
- Richard, M. 2008. *When Truth Gives Out*, Oxford: Oxford University Press.
- Riggs, W. 2009. "Luck, Knowledge, and Control," in Haddock, Millar, and Pritchard, eds., *Epistemic Value*, Oxford: Oxford University Press, 204-21.
- Rønnow-Rasmussen, T. 2002. "Instrumental Values—Strong and Weak," *Ethical Theory and Moral Practice* 5, 23-43.
- Rosen, G. 2001. "Brandom on Modality, Normativity and Intentionality," *Philosophy and Phenomenological Research* 63.
- Russell, B. 1948. *Human Knowledge: Its Scope and Its Limits*. New York: Simon and Schuster.
- Sadock, J. 1978. "On Testing for Conversational Implicature," in P. Cole, ed., *Syntax and Semantics, vol. 9: Pragmatics*, New York: Academic Press, 281-97.
- Shah, N. and Velleman, J. 2005. "Doxastic Deliberation," *Philosophical Review* 114, 497-534.
- Scheffler, S. 1987. "Morality through Thick and Thin: A Critical Notice of *Ethics and the Limits of Philosophy*," *Philosophical Review* 96, 411-34.
- Shope, R. 1983. *The Analysis of Knowing*, Princeton, NJ: Princeton University Press.
- 2004. "The Analysis of Knowing" in Niiniluoto et.al., eds., *Handbook of Epistemology*, Kluwer, 283-329.
- Simmons, L. 1997. "On Not Destroying the Health of One's Patients," in Oderberg and Laing, eds., *Human Lives: Critical Essays on Consequentialist Bioethics*, St. Martin's Press: New York, 15-60.
- Slote, M. 1975. "Inapplicable Concepts," *Philosophical Studies* 28, 265-71.
- Sreenivasan, G. 2001. "Understanding Alien Morals," *Philosophy and Phenomenological Research* 62: 1-32.
- Stalnaker, R. 1970. "Pragmatics," *Synthese* 22, 272-89.
- Stanley, J. and Z. Szabó. 2000. "On Quantifier Domain Restriction," *Mind and Language* 15, 219-61.
- Statman, D. "Moral and Epistemic Luck," *Ratio* 4 (new series), 146-56.

- Steup, M. 2008. "The Analysis of Knowledge," *The Stanford Encyclopedia of Philosophy*, E. Zalta, ed., URL = <<http://plato.stanford.edu/archives/fall2008/entries/knowledge-analysis/>>.
- Sturgeon, S. 1993. "The Gettier Problem," *Analysis* 53, 156-64.
- Suter, R. 1973. "Moore's Defense of the Rule "Do No Murder"," *The Personalist* 4, 361-75.
- Tappolet, C. 2004. "Through Thick and Thin: *Good* and Its Determinates," *Dialectica* 58, 207-21.
- Thomas, A. 2008. "The Genealogy of Epistemic Virtue Concepts," *Philosophical Papers* 37, 345-69.
- Thomson, J. 2008. *Normativity*, Chicago: Open Court.
- Unger, P. 1968. "An Analysis of Factual Knowledge," *The Journal of Philosophy* 65, 157-70.
- Väyrynen, P. 2008. "Slim Epistemology with a Thick Skin," *Philosophical Papers* 37, 389-412.
- 2009. "Objectionable Thick Concepts in Denials," *Philosophical Perspectives: Ethics* 23: 439-69.
- 2012. Forthcoming. "Thick Concepts: Where's Evaluation?" in *Oxford Studies in Metaethics* 7, Oxford: Oxford University Press.
- Ward, G. 1988. *The Semantics and Pragmatics of Preposing*, New York: Garland Publishing.
- Wiggins, D. 1987. "A Sensible Subjectivism?" in *Needs, Values, Truth*, Oxford: Basil Blackwell, 185-214.
- 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.
- 1993. "Who Needs Ethical Knowledge?" in Griffiths, ed., *Royal Institute of Philosophy Supplement: 35 Ethics*, 213-222.
- 1995. "Truth in Ethics," *Ratio* 8(3), 227-42.
- Williams, C. 2003. "Perverted Attractions," *The Monist* 86, 115-40.
- Williamson, F. 2004. "Sex, Disorder and Perversion," *Philosophical Papers* 33, 203-29.
- Williamson, T. 2000. *Knowledge and its Limits*, New York: Oxford University Press.
- 2007. "On Being Justified in One's Head," in M. Timmons, J. Greco, and A. Mele, eds., *Rationality and the Good*, New York: Oxford University Press, 106-22.

-----, 2007. *The Philosophy of Philosophy*, Oxford: Blackwell.

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

Zagzebski, L. 2003. "The Search for the Source of Epistemic Good," *Metaphilosophy* 34, 12-28.

Zimmerman, M. 1993. "Luck and Moral Responsibility," in Statman, ed., *Moral Luck*, Albany: State University of New York Press.