

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

Brent G. Kyle

## I. Introduction.

The Gettier problem has stymied epistemologists. After Edmund Gettier (1963) proved the insufficiency of the tripartite analysis of knowledge, his paper was followed by a barrage of reformulated analyses, most of which were returned with potential counterexamples.<sup>1</sup> A succession of reformulations and 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 exceptions<sup>2</sup>). But the abatement of this project is not due to any consensus over one particular analysis—quite the opposite. Despite the incredible effort 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 (e.g. Williamson 2000, 30).

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,<sup>3</sup> or else an instance of a general problem about conceptual analysis.<sup>4</sup> 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 evaluative concept, and a Gettier-like problem is just what we should expect from attempts at analyzing a concept of this sort.<sup>5</sup> These claims unfold as follows.

---

<sup>1</sup> See Shope (1983 and 2004) for a survey of responses.

<sup>2</sup> For example, see Greco (2003). Pritchard (2005) also seems to provide an analysis in the traditional sense.

<sup>3</sup> 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.

<sup>4</sup> For this type of suggestion, see Laurence and Margolis (1999, 15) and Williamson (2000, 30-3).

<sup>5</sup> 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 evaluative concepts? Very roughly, these are hybrid concepts that somehow combine both evaluative and non-evaluative content.<sup>6</sup> 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*.<sup>7</sup> 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 focus.

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 some of our prized ethical concepts will be subject to a problem similar to the Gettier problem. To the extent that the Gettier problem should trouble epistemologists, very similar problems should also

---

<sup>6</sup> 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.

<sup>7</sup> 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. Their view 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.

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, let me note that my desired conclusion that *knowledge* is thick conflicts directly with the dominant view about *knowledge*. A number of philosophers have recently claimed that both *knowledge* and *justification* are thin concepts (Battaly 2008, 435; Axtell and Carter 2008, 427; Elgin 2008, 371; Thomas 2008, 363; Väyrynen 2008, 392). 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.<sup>8</sup> *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.

---

<sup>8</sup> 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 7 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 (Gibbard 1992, 273; Sreenivasan 2001, 3; Payne 2005, 91). 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 good place to start. They state that thick terms have both evaluative and non-evaluative content. But what does it mean for a term to "have" these contents? In other words, how exactly are thick terms associated with these contents?

The issue of how thick terms are associated with their evaluative and non-evaluative contents is highly controversial, and rather than resolving it here I will try to remain neutral between the two most predominant ways of approaching this issue.<sup>9</sup> 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 the analyzability of thick concepts, there are also many who claim that thick concepts cannot be analyzed.<sup>10</sup> 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

---

<sup>9</sup> These two views might aptly be classified as semantic theories of the relation between thick terms and their evaluative contents. For a defense of a semantic view, see Kyle (forthcoming). Some people reject semantic theories and instead hold that the evaluative contents are only pragmatically associated with thick terms (e.g. see Väyrynen (2009)). However, although an attempt of including these pragmatic theories would be too cumbersome in this paper, I believe it is possible to adapt the main points of this paper to a pragmatic view. Doing so would require formulating (c1) so as to allow for a pragmatic connection between thick terms and their evaluative contents. This would in turn lead to other changes to certain claims in this paper, but none of these changes, to my knowledge, would obstruct the main theses advanced herein.

<sup>10</sup> 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.

positions. So, we should not interpret (c1) and (c2) as saying that  $Tk$  is analyzable in terms of evaluative and non-evaluative contents.

A slightly different option is to understand (c1) and (c2) as saying that  $Tk$  conceptually entails both evaluative and non-evaluative contents. The main worry about this option is that the most widely accepted forms of conceptual entailment are not restrictive enough for this interpretation to work. Consider epistemic theories, according to which  $P$  conceptually entails  $Q$  if it is necessarily true that whoever fully grasps the thought *that  $x$  is  $P$  only if  $x$  is  $Q$*  assents to it, or is justified in assenting to it.<sup>11</sup> On this theory of conceptual entailment, all thin concepts (e.g. *good*) conceptually entail a number of trivial non-evaluative concepts, like *self-identical*. But then, under the present interpretation of (c1) and (c2), all thin terms would satisfy both conditions. Moreover, if Judith Thomson (2010, 10) is correct in claiming that everything is good in some way or other, then it appears that all non-evaluative concepts (e.g. *red*) conceptually entail the evaluative concept *good in a way*. So, under the present interpretation, 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). These conditions are supposed to capture what's most distinctive about the thick. So, we obviously need a more restrictive way of interpreting these conditions.

When ethicists propose (c1) and (c2) as conditions that capture what's most distinctive about thick terms, they seem to understand (c1) and (c2) as claiming that  $Tk$  is "accounted for" in terms of evaluative and non-evaluative content.<sup>12</sup> But what exactly does this accounting-for relation involve?

---

<sup>11</sup> See Boghossian (1997, 334), and see Williamson (2007, ch. 4) for a recent critical survey. In this paper, I assume broadly an epistemic view of conceptual entailment, although nothing of substance hangs on the assumption.

<sup>12</sup> For example, an accounting-for relation is employed by Susan Hurley in a similar context. On Hurley's view, in order to account for  $X$  in terms of  $Y$  we must provide "a conceptual 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  $X$ " (1989, 10). Although her formulation is rough, the view I go on to propose might be equivalent in extension to what Hurley has in mind. Also, on Hurley's view the accounting-for relation does not necessarily involve conceptual analysis, conceptual priority, or conceptual independence. I also intend to avoid these commitments for my version of the accounting-for relation.

As I understand it, the accounting-for relation provides a restrictive subclass of conceptual entailments, a subclass that includes only the entailments that are explanatory of the target concept, thereby leaving out trivial entailments (e.g. from *good* to *self-identical*). I define ‘accounting for’ as follows:

Concept *X* is accounted for in terms of *Y* if and only if

- (i) *X* conceptually entails *Y*, and
- (ii) this entailment plays a role in the best explanation of *X*.

On this view, an account consists only of conceptual entailments that help to best explain the target concept. The best explanation of a concept, as I construe it, is a bi-conditional that best explains why that concept applies when it does, and why it doesn’t apply when it doesn’t. There are two ways in which clause (ii) might go unsatisfied. First, the entailment may play no role in explaining the target concept. Trivial entailments fall in this category. Since everything falls under the concept *self-identical*, the fact that *good* entails *self-identical* does not explain any cases where *good* fails to apply. And second, the entailment may help to explain the target concept but play no role in the best explanation of that concept. For example, the concept *sister* conceptually entails *not male*. This entailment can explain why *sister* fails to apply in certain cases  $c_1 \dots c_n$ , but it wouldn’t play a role in the best explanation of *sister*, because there’s a stronger entailment—from *sister* to *female*—that would explain  $c_1 \dots c_n$  along with many other cases in which *sister* fails to apply (e.g. to nonsexual beings).

It is important to note that the above definition of ‘accounting for’ allows for the possibility of a circular account, where *X* is accounted for in terms of *Y* and *Y* is accounted for in terms of *X*. This situation is possible because the best explanation of each concept might, for all we know, be a circular bi-conditional. (The allowance of circularity is one way in which a full-blown account differs from a conceptual analysis). The reason we should leave open the possibility of a circular account is because some ethicists, like Jonathan Dancy, deal with individual thick concepts in

precisely this way.<sup>13</sup> 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” (1995, 276). Even though this account is circular, Dancy is still attempting to account for the concept *lewd*. And so, for the sake of including views like Dancy’s, our accounting-for relation should not automatically rule out circularity.<sup>14</sup>

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

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

(c2) ‘Tk’ expresses a concept that is 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)? A very simple argument supports the claim that ‘knowledge’ satisfies both (c1) and (c2). Traditionally it has been claimed that *knowledge* entails *justified true belief* and that the former is 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 partially accounted for in terms of an evaluative content. So, under

---

<sup>13</sup> Wiggins (1987) has a similar view.

<sup>14</sup> 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.”

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 so simple. Although I think this argument ultimately succeeds, it nonetheless makes a number of assumptions. Allow me to clarify and substantiate four of them in what follows.

First, the argument relies substantively on a few entailments that have been called into doubt—namely that *knowledge* entails *truth*, *belief*, and *justification*.<sup>15</sup> Space prohibits a 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 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.<sup>16</sup>

Second, 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 it is quite plausible that *justification* is thin, 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’

---

<sup>15</sup> 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 Williamson (2007, 111-12).

<sup>16</sup> The argument also assumes the traditional view that *knowledge* is accounted for in terms of *belief*. But Timothy Williamson has proposed that we should understand *belief* in terms of *knowledge*. On his view, “to believe *p* is to treat *p* as if one knew *p*” (2000, 46). But this is not a problem for my assumption of the traditional view. First, Williamson does not think that *belief* is to be accounted for in terms of *knowledge* in my sense of “accounting for,” because he does not think that *belief* entails *knowledge*. And second, Williamson’s view is not incompatible with my assumption that *knowledge* is accounted for in terms of *belief*. Williamson is concerned to explain the entailment from *knowledge* to *belief*, and he does so by elucidating the concept of belief. But it is nonetheless compatible with his view that the entailment from *knowledge* to *belief* plays a role in the best explanation of *knowledge*. And that’s all I’m assuming.



‘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.<sup>17</sup> Of course, the above ways of understanding 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.<sup>18</sup> 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.<sup>19</sup>

Third, 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

---

<sup>17</sup> 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 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 8, 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.

<sup>18</sup> Goldman (1986, 22-7) believes that justification supervenes on facts about reliability. See Conee and Feldman (2001) for the view on which justification supervenes on internalist factors.

<sup>19</sup> Richard Fumerton (2001) argues, by process of elimination, that there isn’t any interesting sense in which justification is evaluative (or, rather, normative). Without going into detail, let me note that Fumerton (2001, 52) seems to think that a certain form of noncognitivism about value judgments could uphold the view that justification is normative. The problem, as he correctly notes, is that few epistemologists would happily accept the view that claims about justification are not truth-apt. However, Fumerton fails to recognize a possibility that has recently attracted many meta-ethicists, known as “hybrid expressivists” (e.g. see Copp 2001). Very roughly, these theorists accept that value claims are truth-apt but they appropriate certain features of traditional expressivism, such as its explanation of moral motivation and response to the Open Question argument. If one was a hybrid expressivist and explained the normativity of justification in terms of hybrid expressivism, then such a view would avoid the above problem, since it can allow claims about justification to be truth-apt. I am not claiming that hybrid expressivism is correct (see Schroeder 2009 for critical discussion), but only that Fumerton’s argument overlooks this viable possibility.

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.<sup>20</sup>

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 must be "significant" in a certain sense.<sup>21</sup> If we accept this requirement, 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'?

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).<sup>22</sup> 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 accounted for in terms of some non-evaluative content (namely *true belief*), *knowledge* does not have enough non-evaluative content for us to correctly

---

<sup>20</sup> Of course, my claim here is compatible with the widely accepted view that true beliefs are highly valuable. Consider an analogy: the concept *diamond* is non-evaluative, even though diamonds are highly valuable.

<sup>21</sup> For the most straightforward use of a 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 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.

<sup>22</sup> See Väyrynen (2008, 391). Williams (1995, 234) and Brewer (2009, 185-86) appear to have a similar type of view in mind. The basic idea behind the continuum view most likely originates in Scheffler (1987, 417).

place it on the thick end of the continuum. In reply, 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.<sup>23</sup> But many will agree that *deliberate killing* is the non-evaluative content associated with *murder* (e.g. Suter 1973, 361; Anscombe 1979, 74). 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.

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 terms "ad hoc conjunctions." As I shall understand this notion, an ad hoc conjunction

---

<sup>23</sup> 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. But I don't see why we should accept this claim. 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.

expresses the mere conjunction of a thin concept with a wholly non-evaluative concept. Ad hoc conjunctions satisfy both (c1) and (c2), but they are not thick.<sup>24</sup> Thick concepts cannot be manufactured simply by conjoining a thin concept with a wholly non-evaluative concept. Catherine Elgin (2008, 372) has explicitly excluded ad hoc conjunctions from the thick. 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 she still wouldn’t be courageous. Metaphorically speaking, the evaluative and non-evaluative contents of a paradigmatic thick term seem more entwined than those of an ad hoc conjunction.

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.

---

<sup>24</sup> Note that terms expressing mere conjunctions of thick concepts with wholly non-evaluative concepts are not ad hoc conjunctions in the relevant sense. For example, if ‘monrageous’ expresses the concept *courageous and born on a Monday*, this expression would not count as an ad hoc conjunction.

I believe that terms like ‘monrageous’ should count as expressing thick concepts. Intuitively speaking, a thick concept is supposed to involve an evaluative content and a non-evaluative content that are appropriately related to one another. “Appropriately related” is admittedly vague, but we need not specify this before we can see that ‘monrageous’ will meet this intuitive requirement. Since ‘monrageous’ has a thick concept built into it, it will satisfy this intuitive requirement by virtue of its “in built” thick concept—that is, ‘monrageous’ has an evaluative and a non-evaluative content that are appropriately related by virtue of containing the thick concept *courage*. The mere fact that ‘monrageous’ also has some extra non-evaluative content does not, to my mind, disqualify it from counting as thick. None of this can be said for ad hoc conjunctions, since they do not have an “in built” thick concept. The condition I go on to propose will allow for the inclusion of ‘monrageous’ with the thick.

Others believe that 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.<sup>25</sup> But it is obvious that ad hoc conjunctions do not admit of this style of analysis. My car is good 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 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.<sup>26</sup> 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

---

<sup>25</sup> 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.

<sup>26</sup> 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 option 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 Kyle (2012, ch. 3).

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 agree that thick concepts are 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.<sup>27</sup> On this view, the relevant thin concept would be *merit*. Henceforth, let  $T_n$  stand for the thin concept that's used within the account of a given concept.<sup>28</sup> So the claim here is that, if  $T_k$  is a thick concept, then there is some thin concept  $T_n$  such that  $T_k$  is to be accounted for in terms of  $T_n$ .

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 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  $T_n$  and the concept *red* as its corresponding  $N$ . Consider some provisional examples:<sup>29</sup>

---

<sup>27</sup> Here I'm conjecturing based on how he accounts for *lewd*. See Dancy (1995, 275-76).

<sup>28</sup> If there is more than one such thin concept, then let  $T_n$  stand for their conjunction.

<sup>29</sup> 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  $T_n$  to a concept like *merit* or *appropriate*, but that wouldn't make any important difference to my main claims.

<u>Target Concept</u>	<u>Tn</u>	<u>N</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 *C* is (i) logically stronger than *N*, (ii) compatible with *Tn*, but (iii) incompatible with *Tk*.<sup>30</sup>

(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 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

---

<sup>30</sup> The additional concept *C* must also be internally consistent. Otherwise ad hoc conjunctions would satisfy (c3) (e.g. ‘gred’ could have *good*, *red*, and *not red* as its concept *C*, which would satisfy clauses (i)-(iii)). There is good reason to require that *C* be internally consistent. After all, the main reason for focusing on concept *C* is that *C* applies within certain possible cases where *N* and *Tn* apply without *Tk*. But if *C* is internally inconsistent then there will be no such possible case.

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* fulfills all of (i), (ii), and (iii).<sup>31</sup> Ad hoc conjunctions do not satisfy the condition in question.

---

<sup>31</sup> Here is a proof showing that ad hoc conjunctions don’t satisfy (c3). (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*. Thus, we can assume the following:

(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,

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

From these claims, we get a contradiction. Given (2), and assuming *C* is internally consistent, we’re entitled to assume that we have a possible 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)]



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 (e.g. *bad* or *wrong*). 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 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*.

---

(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).

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 such an act would not count as cruel. And if one kills an aggressor in order to save one's own life, then such an act would not count 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.<sup>32</sup>

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'.<sup>33</sup> So, we might wonder whether these questionable cases present problems for my claim that (c1)-(c3) are sufficient.

---

<sup>32</sup> 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 kind of maneuver. First, even if the expanded content is the appropriate *N*, this 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. But second, this 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*.

<sup>33</sup> 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 by claiming that its evaluative content does not obtain in virtue of the obtaining non-evaluative content:

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 (e.g. a deliberate killing not done out of self-defense, etc).

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's unclear how the requirement won't end up excluding many paradigmatic thick

I am ultimately unsure whether the above 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 unlikely that these conditions would 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 clear that there is such a concept, one that has recently come under the name 'veritic luck.'<sup>34</sup> 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, 29):

---

terms. We should be open to the possibility that, in at least some contexts of utterance, 'good human being' does express a thick concept.

<sup>34</sup> The name comes from Mylan Engel (1992, 59) and has been adopted recently by Pritchard (2005, 145).

[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.

Or consider three more cases from Bertrand Russell (1948, 154-55):

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.

These represent a larger body of cases that I'll call cases of veritic luck. They typically 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 true.<sup>35</sup> 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.

---

<sup>35</sup> 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.

Given this notion of veritic luck, we can now ask whether the concept *veritically lucky belief* satisfies clause (i) of (c3). 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* also satisfies clause (iii). In other words, the following is very plausible:

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

(iii') is a form of the anti-luck platitude advanced by many epistemologists—"knowledge excludes luck".<sup>36</sup> 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 *veritically lucky belief* 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*.

---

<sup>36</sup> 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.

Indeed, (ii) is most certainly true. 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 Smith," Brown is in Barcelona (1963, 232-33). 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. In general, I take Gettier-style cases to be cases where the subject has a veritically lucky belief. And I take the main insight of Gettier's paper 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.<sup>37</sup> In sum, Gettier cases show us that (ii) is true.

---

<sup>37</sup> 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

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 (2007, 283).

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

If *safe belief* were the relevant *N*, then 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. If he's right, then *veritically lucky safe belief* is not a coherent concept. However, Avram Hiller and Ram Neta (2007, 307-8) have together pointed out that safe beliefs can be veritically lucky. Following Hiller and Neta, we can simply revise Gettier's original counterexample by describing Brown in such a way that he is in Barcelona in all nearby possible worlds—for example, suppose that "Brown is so constituted (psychologically, financially, and otherwise) that he would only leave Barcelona in the strangest of circumstances." So, by Pritchard's definition of safety,

---

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."

Smith has a safe belief that either Jones owns a Ford or Brown is in Barcelona. Moreover, we can continue to maintain Smith's ignorance of Brown's whereabouts by stipulating that Smith got to know Brown "through an internet chat room" and that Brown never disclosed information about his location (2007, 308). So, it's a matter of luck that Smith's belief is true. Therefore, in this revised case Smith has a veritically lucky safe belief, which means that the concept *veritically lucky safe belief* is coherent after all. And since *veritically lucky safe belief* is logically stronger than *safe belief*, (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:

(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\*). Hiller and Neta's revised Gettier case also shows that *veritically lucky safe belief* is compatible with *justification*. After all, we can easily imagine that Smith's belief is justified via the same mechanisms as the original Gettier case. Thus, (ii\*) also appears to be true. Now consider (iii\*). We've already seen that *veritically lucky belief* is incompatible with *knowledge*. If that's right, then *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.

Let us now turn to the possibility in which *reliably formed true belief* is the relevant N. For present purposes, 'reliable' can be defined as a process that produces either mostly true beliefs or at least a high ratio of true beliefs to false beliefs. A reliably formed true belief is therefore a true belief that is produced by a reliable process. Suppose then that *reliably formed true belief* is the relevant N associated with *knowledge*. In this case, what I said above with regard to *safe belief* can be said similarly with regard to the concept *reliably formed true belief*. In particular, we can take the following to be the



relevant *C*—*veritically lucky reliably formed true belief*. This concept simply refers to reliably formed true beliefs that are also veritically lucky. And the coherence of this concept can be demonstrated, not by a variation of Gettier’s original case, but by Carl Ginet’s barn façade example, first published in Goldman (1976, 772-73):

Henry is driving in the country-side with his son. For the boy’s edification Henry identifies various objects on the landscape as they come into view. “That’s a cow,” says Henry, “That’s a tractor,” “That’s a silo,” “That’s a barn,” etc. Henry has no doubt about the identity of these objects; in particular, he has no doubt that the last-mentioned object is a barn, which indeed it is. Each of the identified objects has features characteristic of its type. Moreover, each object is fully in view, Henry has excellent eyesight, and he has enough time to look at them reasonably carefully, since there is little traffic to distract him. [...] Suppose we are told that, unknown to Henry, the district he has just entered is full of papier-mâché facsimiles of barns. The facsimiles look from the road exactly like barns, but are really just façades, without back walls or interiors, quite incapable of being used as barns. They are so cleverly constructed that travelers invariably mistake them for barns. Having just entered the district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. But if the object on that site were a facsimile, Henry would mistake it for a barn.

We can easily suppose that in this case Henry has a true belief that is reliably formed—it is formed by a perceptual process that produces mostly true beliefs, or at least a high ratio of true beliefs to false beliefs. But it also appears to be veritically lucky—it is a matter of luck that Henry has a true belief because he could have easily been mistaken. So, the concept *veritically lucky reliably formed true belief* is coherent, and we can therefore take it to be the relevant *C* in the event that *reliably formed true belief* is the relevant *N*. Claims (i’)-(iii’) can be revised accordingly.<sup>38</sup>

Next, let’s consider whether *sensitive belief* could be the relevant *N* associated with *knowledge*. Following Robert Nozick (1981, 172), we can say that a belief that *P* is sensitive only if, were *P* false, the subject would not believe *P*. The concept *sensitive belief* cannot be dealt with in the same way that I dealt with previous candidates for *N*. This is because it’s not obvious that sensitive beliefs can be

---

<sup>38</sup> Alvin Goldman (1976) originally used the barn façade example to motivate his relevant alternatives account of knowledge, which he understands to involve a form of reliability (although different from how I defined ‘reliability’). I do not address this view here because the distinction between relevant and irrelevant alternatives is widely seen as unclear and problematic. See Vogel (1999, 163-68) for discussion.

veritically lucky. Still, there is good reason to deny that *sensitive belief* could be the relevant *N* associated with *knowledge*. For it to be the relevant *N*, sensitive belief must be necessary for knowledge. But it's extremely doubtful that sensitivity is necessary for knowledge. Without going into detail, there are many difficulties that allegedly arise for a sensitivity requirement on knowledge. As Nozick himself points out (1981, 204-5), the sensitivity requirement is in conflict with the widely-held principle that knowledge is closed under known entailment. And others have argued that the sensitivity requirement is incompatible with inductive knowledge (Vogel 1987; Sosa 1999, 145), with higher-level knowledge (Vogel 2000), and with the view that knowledge requires margins for error (Williamson 2000, ch. 7). If even one of these problems proves fatal, then the *sensitive belief* cannot be the relevant *N* associated with *knowledge*. I will therefore proceed as if *sensitive belief* is not the requisite *N*.

Although I cannot discuss all possible substitutes for *N*, I have so far discussed four of the main candidates—*true belief*, *safe belief*, *reliably formed true belief*, and *sensitive belief*.<sup>39</sup> As we've seen, the first three concepts are compatible with veritic luck, whereas the fourth concept is not likely a necessary condition for knowledge. There is, however, one last possibility that must be treated differently.

Could *true belief that's not veritically lucky* be the relevant *N*? I shall assume not. Certainly it's true that *knowledge* entails *true belief that's not veritically lucky*, but it doesn't seem that *knowledge* is to be accounted for in terms of this concept. Recall the definition of 'accounting for' (p. 6). What I am assuming here is that this entailment would not play a role in the best explanation of *knowledge*, that is, in the best explanation of why *knowledge* applies when it does and why it doesn't apply when it doesn't. To be sure, many epistemologists are keen on providing an account of *knowledge* that

---

<sup>39</sup> 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.

explains why *knowledge* does not apply in veritic luck cases. But the standard approach is to account for *knowledge* in terms of some additional element (e.g., safety), and then to claim that this additional element provides the basis for excluding luck. A recent alternative to the standard approach is to reject the possibility of a non-circular account of *knowledge* (Williamson 2000, 2-5). Even this alternative approach must assume that *knowledge* is not accounted for in terms of *true belief that's not veritically lucky*; otherwise there obviously would be a non-circular account available. But no one, to my knowledge, has claimed that *knowledge* is to be accounted for in terms of *true belief that's not veritically lucky*.<sup>40</sup> And why is this? Presumably, it's because there's a pervasive sense among epistemologists that the best non-circular explanation of *knowledge* would explain the anti-luck platitude that knowledge excludes luck. But if we accounted for *knowledge* directly in terms of *being not veritically lucky*, then our account would simply assert the anti-luck platitude without explaining it. For this reason, I will assume, along with virtually all other epistemologists, that *knowledge* is not to be accounted for 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*). Therefore, '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.

As noted at the outset, many theorists have classified *knowledge* as a thin concept. We've seen that they are mistaken. But their 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.

---

<sup>40</sup> 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.

### 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  $C$  such that  $C$  is (i) 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 us now examine the issue of thick conceptual analysis in 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 task 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 pain<sup>41</sup> 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).

Is this claim 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

---

<sup>41</sup> 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.

very surgery. But then it looks like conditions (i) and (ii) both obtain, which means they are not sufficient for cruelty.<sup>42</sup>

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). The TCA problem is the task of stating non-circular, necessary and sufficient 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 is an adequate response remains to be seen.

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*. 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

---

<sup>42</sup> 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. Still, these treatments of one another are often cruel without the intention to induce pain.

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 Gettier-style cases.<sup>43</sup> As described, the Gettier problem looks very much 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 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 (2005, 148-9) and strongly suggested by Mylan Engel (1992, 64-72). 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).<sup>44</sup> For these reasons, I take it to be plausible that the Gettier problem is an instance of the TCA problem.

---

<sup>43</sup> 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.

<sup>44</sup> 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.

If it is an instance of the TCA problem, then we would have an answer to our initial question: Why does the Gettier problem arise in the first place? On my view, it 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 philosophers believe that most concepts are unanalyzable on the grounds that few of them (if any) admit of uncontroversial analyses.<sup>45</sup> 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 (c3) that explains this difference, a

---

<sup>45</sup> For a discussion on this, see Laurence and Margolis (1999, 14-16).



condition that connected our discussion to the Gettier problem. Gettier's counterexamples are at the heart of the fact that 'knowledge' satisfies (c3). The fact that 'knowledge' satisfies (c3) secured my conclusion that *knowledge* is thick, and 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 the TCA 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.

### Bibliography

- Alston, William (1989) "Justification and Knowledge," in *Epistemic Justification: Essays in the Theory of Knowledge*, 172-82, Ithaca, NY: Cornell University Press.
- Anscombe, G.E.M. (1979) "Prolegomenon to a Pursuit of the Definition of Murder," *Dialectics and Humanism* 4, 73-77.
- Audi, Robert (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 *Rationality and the Good*, eds. M. Timmons, J. Greco, and A. Mele, 222-47. New York: Oxford University Press.
- Ayer, A.J. (1956) *The Problem of Knowledge*. London: Macmillan and Company Limited.
- Axtell, Guy and Adam Carter (2008) "Just the Right Thickness: A Defense of Second-Wave Virtue Epistemology," *Philosophical Papers* 37, 413-34.
- Battaly, Heather (2008) "Metaethics Meets Virtue Epistemology: Salvaging Disagreement about the Epistemically Thick," *Philosophical Papers* 37, 435-54.
- Blackburn, Simon (1984) *Spreading the Word*. Oxford: Clarendon Press.
- Boghossian, Paul (1997) "Analyticity," in *A Companion to the Philosophy of Language*, eds. B. Hale and C. Wright, 331-68. Oxford: Blackwell Publishers.

- Brewer, Talbot (2009) *The Retrieval of Ethics*, Oxford: Oxford University Press.
- Burton, Stephan L. (1992) "'Thick' Concepts Revisited," *Analysis* 52, 28-32.
- Conee, Earl and Richard Feldman (2001) "Internalism Defended," *American Philosophical Quarterly* 38, 1-18.
- Copp, David (2001) "Realist-Expressivism: A Neglected Option for Moral Realism," *Social Philosophy and Policy* 18: 1-43.
- Dancy, Jonathan (1995) "In Defense of Thick Concepts," *Midwest Studies in Philosophy* 20, 263-79.
- Elgin, Catherine (2005) "Williams on Truthfulness," *The Philosophical Quarterly* 55, 343-52.
- (2008) "Trustworthiness," *Philosophical Papers* 37, 371-87.
- Engel, Mylan Jr. (1992) "Is Epistemic Luck Compatible with Knowledge," *The Southern Journal of Philosophy* 30, 59-75.
- Fumerton, Richard (2001) "Epistemic Justification and Normativity," in *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. M. Steup, 49-60. Oxford University Press.
- Gettier, Edmund (1963) "Is Justified True Belief Knowledge?" *Analysis* 23, 121-23.
- Gibbard, Allan (1992) "Thick Concepts and Warrant for Feelings," *Proceedings of the Aristotelian Society*, suppl. vol. 66: 267-83.
- (2003) "Reasons Thin and Thick," *The Journal of Philosophy* C(6), 288-304.
- Ginet, Carl (1988) "The Fourth Condition," in *Philosophical Analysis*, ed. D.F. Austin, Kluwer Academic Publishers, 105-17.
- Goldman, Alvin (1976) "Discrimination and Perceptual Knowledge," *The Journal of Philosophy* 73, 771-91.
- (1986) *Epistemology and Cognition*, Cambridge, Mass.: Harvard University Press.
- Greco, John (2003) "Knowledge as Credit for True Belief," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. DePaul and Zagzebski, Oxford: Oxford University Press.
- Hare, R.M. (1963) *Freedom and Reason*. Oxford: Oxford University Press.
- Hawthorne, John (2004) *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hazlett, Allan (forthcoming) "The Myth of Factive Verbs" *Philosophy and Phenomenological Research*.
- Hiller, Avram and Ram Neta (2007) "Safety and Epistemic Luck," *Synthese* 158, 303-313.

- Hurka, Thomas and Daniel Elstein (2009) "From Thick to Thin: Two Moral Reduction Plans" *The Canadian Journal of Philosophy* 39, 551-36.
- Hurley, Susan (1989) *Natural Reasons*, Oxford: Oxford University Press.
- Kornblith, Hilary (2009) "Knowledge Needs No Justification," in *Epistemology: New Essays*, ed. Quentin Smith, 5-23.
- Kyle, Brent (forthcoming) "How Are Thick Terms Evaluative?" *Philosophers' Imprint*.
- (2012) *Knowledge as a Thick Concept: New Light on the Gettier and Value Problems*. Doctoral Dissertation.
- Laurence, Stephen and Eric Margolis (1999) "Concepts and Cognitive Science," in *Concepts: Core Readings*, eds. Laurence and Margolis, 3-81.
- Lewis, David (1996) "Elusive Knowledge," *Australasian Journal of Philosophy* 74, 549-67.
- Lovibond, Sabina (1983) *Realism and Imagination in Ethics*. Minneapolis: University of Minnesota Press.
- Nozick, Robert (1981) *Philosophical Explanations*. Cambridge, Mass.: Belnap Press of Harvard University Press.
- Payne, Andrew (2005) "A New Account of Thick Concepts," *The Journal of Value Inquiry* 39, 89-103.
- Plantinga, Alvin (1993) *Warrant and Proper Function*, Oxford: Oxford University Press.
- Platts, Mark (1979) *Ways of Meaning*, Cambridge, Mass: MIT Press.
- Pritchard, Duncan (2005) *Epistemic Luck*, Oxford: Oxford University Press.
- (2007) "Anti-Luck Epistemology," *Synthese*, 277-97.
- Russell, Bertrand (1948) *Human Knowledge: Its Scope and Its Limits*. New York: Simon and Schuster.
- Schroeder, Mark (2009) "Hybrid Expressivism: Virtues and Vices," *Ethics* 119: 257-309.
- Scheffler, Samuel (1987) "Morality through Thick and Thin: A Critical Notice of *Ethics and the Limits of Philosophy*," *Philosophical Review* 96, 411-34.
- Shope, Robert. (1983) *The Analysis of Knowing*. Princeton, NJ: Princeton University Press.
- (2004) "The Analysis of Knowing" in *Handbook of Epistemology*, eds. Niiniluoto et.al., Kluwer, 283-329.
- Simmons, Lance (1997) *Human Lives: Critical Essays on Consequentialist Bioethics*. eds. Oderberg and

- Laing. St. Martin's Press: New York, 15-60.
- Sosa, Ernest (1999) "How to Defeat Opposition to Moore," *Philosophical Perspectives 13 Epistemology*, 144-53.
- Sreenivasan, Gopal (2001) "Understanding Alien Morals" *Philosophy and Phenomenological Research* 62: 1-32.
- Steup, Matthias (2008) "The Analysis of Knowledge," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, URL = <<http://plato.stanford.edu/archives/fall2008/entries/knowledge-analysis/>>.
- Sturgeon, Scott (1993) "The Gettier Problem," *Analysis* 53, 156-64.
- Suter, Ronald (1973) "Moore's Defense of the Rule "Do No Murder","" *The Personalist* 4, 361-75.
- Tappolet, Christine (2004) "Through Thick and Thin: *Good* and Its Determinates" *Dialectica* 58, 207-21.
- Thomas, Alan (2008) "The Genealogy of Epistemic Virtue Concepts," *Philosophical Papers* 37, 345-69.
- Unger, Peter (1968) "An Analysis of Factual Knowledge," *The Journal of Philosophy* 65, 157-70.
- Väyrynen, Pekka (2008) "Slim Epistemology with a Thick Skin," *Philosophical Papers* 37, 389-412.
- (2009) "Objectionable Thick Concepts in Denials," *Philosophical Perspectives 23 Ethics*, 439-68.
- Vogel, Jonathan (1987) "Tracking, Closure, and Inductive Knowledge," in *The Possibility of Knowledge: Nozick and His Critics*, ed. Luper-Foy, Oxford: Rowman and Littlefield.
- (1999) "The New Relevant Alternatives Theory," *Philosophical Perspectives 13 Epistemology*, 156-80.
- (2000) "Reliabilism Leveled," *The Journal of Philosophy* 97, 602-23.
- Wiggins, David. (1987) "A Sensible Subjectivism?" in *Needs, Values, Truth*, 185-214, Oxford: Basil Blackwell.
- Williams, Bernard (1985) *Ethics and the Limits of Philosophy*, Cambridge, Mass: Harvard University Press.
- (1995) "Truth in Ethics" *Ratio* 8(3), 227-42.
- Williamson, Timothy (2000) *Knowledge and its Limits*. New York: Oxford University Press.
- (2007) "On Being Justified in One's Head," in *Rationality and the Good*, eds. M. Timmons, J. Greco, and A. Mele, 106-22, New York: Oxford University Press.
- (2007) *The Philosophy of Philosophy*. Oxford: Blackwell Publishing.