

---

## How We Know Pain Is Bad

---

Andrew Y. Lee | [andrewyuanlee@gmail.com](mailto:andrewyuanlee@gmail.com) | <http://andrewyuanlee.com/>  
University of Oslo, Department of Philosophy

---

WORD COUNT: 8286

KEYWORDS: badness of pain, epistemic gap, hedonic inversion, moral epistemology, pain, fact-value gap, is-ought gap, value of consciousness, well-being

---

### **Abstract:**

Consider two widely endorsed claims: (1) that no evaluative conclusion is derivable from purely descriptive premises, and (2) that one can know pain is bad simply on the basis of knowing how pain feels. These claims are in tension. The aim of this paper is to argue that the best way of resolving the tension is to restrict the scope of the first claim. In particular, I develop and defend the thesis that value facts about experiences (such as the fact that pain is bad) are epistemically entailed by phenomenal facts (such as the fact that pain feels a certain way). I argue that this thesis aligns with common sense, that denying it leads to counterintuitive consequences, that it explains doxastic asymmetries between different classes of ethical beliefs, that it accounts for the epistemic defects of hedonic inverts, and that it is defensible in the face of objections concerning rationality, analyticity, inference, explanation, debunking, and more. I also argue that we should embrace the surprising upshot that some value facts are genuinely derivable from purely descriptive facts.

---

---

## Introduction

Consider two widely endorsed claims that seem to be in tension with each other. The first is that no evaluative conclusion is derivable from purely descriptive premises. This principle—which I will call the *epistemic gap*—captures the idea that knowledge purely of how things are can never by itself suffice for knowledge of what is good or bad. The second is that one can know pain is bad simply on the basis of knowing how pain feels. In other words, anyone who knows what it is like to for an ordinary person to put their hand on a hot stove, to stub their toe, or to break a bone, is in a position to know that that kind of experience is bad.

This paper argues that the best way to resolve the tension is to restrict the scope of the epistemic gap. In particular, I will argue that value facts about experiences—a special subclass of ethical facts—are genuine counterexamples to the epistemic gap. Here is my core thesis:

### The Phenomenal Value Theory

Value facts about experiences are epistemically entailed by phenomenal facts.

By *value facts about experiences*, I mean facts purely about which kinds of experiences are good or bad. By *phenomenal facts*, I mean facts purely about what it is like to have certain experiences. I will later explain in more detail what I mean by *epistemic entailment*, but I will assume that if P epistemically entails Q, then P a priori entails Q, that Q is knowable on the basis of P, that P justifies Q, that it is rationally impermissible to believe that P and not Q, and that Q is a priori deducible from P.

The phenomenal value theory takes the aforementioned claim about how we know pain is bad at face value. According to the theory, anyone who knows how pain feels is thereby in a position to know that pain is bad, and someone who knows how pain feels yet infers that pain is good would be making an epistemic mistake. These claims may strike some as obvious, but they are puzzling when considered in light of the epistemic gap. In this respect, a goal of this paper is to follow Nagel's suggestion of getting "rid of the obstacles to the admission of the obvious."<sup>1</sup>

§1 explains and motivates the phenomenal value theory; §2 defends the theory from objections.

---

<sup>1</sup> Nagel [1980, p.109].

---

## §1 | The Phenomenal Value Theory

I will begin by characterizing the phenomenal value theory and the epistemic gap. Then I will present some motivations for the phenomenal value theory.

### Foundations

The focal example of this paper is the badness of pain. By *pain*, I mean the kind of nociceptive experience had by ordinary humans—in particular, a kind of experience that hurts. By *hurt*, I mean the unpleasant aspect of ordinary pain experiences. Putting it another way, I take ‘pain’ to by definition involve experiential suffering.<sup>2</sup> There are non-ordinary cases of nociceptive experiences that do not hurt, but such experiences lie outside the extension of ‘pain’ as I use the term.<sup>3</sup> Note that this is a verbal stipulation, rather than a stance on the nature of pain: for those skeptical that this definition of ‘pain’ sufficiently accords with either ordinary or scientific usage, we could simply replace all instances of the term ‘pain’ in this paper with another term (such as ‘unpleasant pain’). And for those skeptical that pain is bad, we could simply substitute in any other kind of experience that one takes to be either bad (e.g., torturous pain, nausea, or fear) or good (e.g., pleasure, beauty, or insight).

This paper is about the relationship between phenomenal facts and a certain class of value facts. By *phenomenal fact*, I mean a fact that ascribes only phenomenal properties (meaning properties that characterize only what it is like to have an experience). By *value fact*, I mean a fact that ascribes only value properties (meaning properties that characterize only how good or bad something is).<sup>4</sup> The fact that pain

---

<sup>2</sup> In ordinary language, we also predicate ‘hurt’ to bodily events (such as a pinprick) and body parts (such as one’s toe). But we can stipulate that ‘hurt’ is a theoretical term denoting a phenomenal property of the experiences associated with those bodily events or body parts.

<sup>3</sup> The most prominent cases involve pain asymbolia, a neurological disorder where subjects report feeling pain sensations that do not hurt. For discussion of pain asymbolia and its philosophical significance, see Grahek [2007], Bain [2014], and Klein [2015]. For an overview of theories of pain, see Aydede [2013].

<sup>4</sup> I focus on universal facts (where badness is attributed to properties, such as pain), but I take the arguments to generalize to particular facts as well (where badness is attributed to particulars, such as a particular pain).

---

is bad is a value fact (but not a phenomenal fact); the fact that pain hurts is a phenomenal fact (but not a value fact). Note that I am individuating facts by meanings, rather than by extensions: even if the fact that pain hurts necessitates the fact that pain is bad, that does not make the fact that pain hurts a value fact (since what it *means* for pain to hurt is for pain to feel a certain way). These remarks provide an initial gloss, but we will return to these issues about categorizing facts in §2.

Though I use the term ‘value facts about experiences’ without qualification, it is best to understand the phenomenal value theory as a restricted universal claim. Consider a pluralist about value who agrees that some but not all value facts about experiences are explained by phenomenal facts. Suppose, for example, that the pluralist thinks that pain is bad both because of how it feels and because it is disliked (where the badness due to each factor is non-redundant). Since nearly all the arguments in this paper support this sort of pluralist position (as well as monist positions), it is best to include such views within the scope of the phenomenal value theory. In light of this, the phenomenal value theory can be more strictly understood as the claim that there is a special class of value facts about experience that are explained by phenomenal facts (where that special class may or may not include all value facts about experiences). To simplify the prose, though, I will continue using the term ‘value facts about experience’ (with the quantifier restriction implicit).

The kinds of value facts that I have in mind are *intrinsic* (as opposed to instrumental), and *pro-tanto* (as opposed to all-things-considered). But I will be neutral on whether the relevant value facts are personal (concerning what makes one’s life better or worse) or impersonal (concerning what makes things in general better or worse). My own view is that the phenomenal value theory applies to value facts about experiences of both varieties. However, since the main aim of this paper is to argue that there is a special class of value facts about experiences that are epistemically entailed by phenomenal facts, those who are skeptical about the applicability of the thesis to one variety of value facts (e.g. impersonal) may simply read the paper as focused on the other variety (e.g. personal).

I take value facts about experiences to be a subclass of *ethical facts*, where ethical facts may concern not only the values of experiences but also the values of other kinds of things, and not only what is good but also what is right, or permissible, or obligatory, or fitting. Though I argue that value facts about experiences provide reasons for restricting the epistemic gap, I will take for granted that the

---

epistemic gap remains for other kinds of ethical facts. A consequence is that there is a source of knowledge for value facts about experiences that does not yield knowledge of other kinds of ethical facts.

### The Epistemic Gap

In the abstract, the epistemic gap<sup>5</sup> strikes many as intuitively plausible: how could facts about what matters be derived merely from facts purely about what is? The epistemic gap is most commonly characterized as the principle that no evaluative conclusion is derivable from purely descriptive premises.<sup>6</sup> But for our purposes, let us slightly reformulate it as the thesis that no value fact is epistemically entailed by any set of purely descriptive facts. This reformulation ensures that if the phenomenal theory is true, then the epistemic gap is false.

Let me say more about *epistemic entailment*. The term ‘epistemic entailment’ is intended to stand in for whichever epistemic relation one thinks best characterizes the epistemic gap. I will argue that on any reasonable interpretation of that relation, there are some violations of the epistemic gap. More specifically, here are some ways of precisifying what it means to say that P epistemically entails Q: (1) Q is knowable solely on the basis of P, (2) P a priori entails Q, (3) Q is scrutable from P, (4) P fully justifies Q, (5) P is conclusive evidence for Q, (6) P & ¬Q is inconceivable, and (7) it is ideally rationally impermissible to believe the conjunction P & ¬Q. For the purposes of this paper, I will take epistemic entailment to entail all of these conditions. Given

---

<sup>5</sup> The term ‘epistemic gap’ is not the most frequently used term in the metaethics literature. But it is arguably better in this context than the alternatives, which include ‘is-ought gap’ (my concern is primarily with value rather than oughts), ‘fact-value gap’ (there are facts about values), ‘Hume’s Law’ (which is liable to be confused with Hume’s Principle or Hume’s Dictum), and the ‘Open Question Argument’ (which arguably concerns reductive analysis, as I will discuss later in this section). For more on the epistemic gap in metaethics, see Sayre-McCord [2014], Maguire [2017], and Sparks [forthcoming]. For more general discussions of epistemic gaps, see Jackson [1998], Chalmers [2012], and Mehta [2019].

<sup>6</sup> There are well-known counterexamples to simple formulations of the epistemic gap, but these are widely regarded as exposing technical problems with its formulation rather than casting doubt on the core idea. For a few classic papers, see Prior [1960], Jackson [1974], and Pidgeon [1989]. For more recent discussions, see Singer [2015], and Fine [2108].

---

this, the phenomenal value theory postulates a particularly intimate epistemic connection between knowledge of how pain feels and knowledge that pain is bad.

It is important to distinguish epistemic entailment from reductive analysis. Whereas epistemic entailment is an epistemic relation between facts, reductive analysis is a metaphysical relation between properties. To say that a value property is reductively analyzable in terms of a phenomenal property is to say that what it is for that value property to be instantiated is for that phenomenal property to be instantiated. Though I think that value facts about experiences are epistemically entailed by phenomenal facts, I am skeptical that value properties are reductively analyzable in terms of phenomenal properties. For example, I do not think that what it is for badness to be instantiated is for hurting to be instantiated (since it is at least conceivable for badness to be instantiated in the absence of hurting). As an analogy, consider the common view that microphysical facts epistemically entail macrophysical facts even though there are no reductive analyses of macrophysical properties in terms of microphysical properties.<sup>7</sup>

There may be some who think the epistemic gap ought to be understood as a thesis about reductive analysis rather than epistemic entailment.<sup>8</sup> Though it is beyond the scope of this paper to fully address this issue, it is worth noting that it is common to characterize the epistemic gap in terms of epistemic entailment. For example, Sayre-McCord [2014, §4], in discussing Hume's original formulation of the epistemic gap, characterizes the idea as follows: "If...one infers from the fact that someone is feeling pain that something bad is happening, one is at least presupposing that pain is bad. And that presupposition, in turn, is not entailed by any claims concerned solely with plain matters of fact." My aim is to argue that the idea expressed by this sort of claim is false, at least for value facts about experiences.

### **The Phenomenal Value Theory**

The phenomenal value theory says that value facts about experiences are epistemically entailed by phenomenal facts. Let me briefly note a few issues that the

---

<sup>7</sup> See Fodor [1974] for a classic argument against reductive analyses of macroproperties. See also Moore [1903] on the open question argument, and Suikkanen [2016] and Ridge [2019] for more recent discussions of reductive analyses in metaethics.

<sup>8</sup> See Mehta [2019] for a recent example.

---

phenomenal value theory leaves open. First, the theory is neutral on whether all value facts (or only a privileged subset) are explained by phenomenal facts. Second, it is neutral on whether all phenomenal facts (or only a privileged subset) explain value facts. Third, it is neutral on how difficult it is to acquire knowledge of any given value fact on the basis of the relevant phenomenal facts. Fourth, it is neutral on which kinds of experiences are valuable (and, more generally, on issues at the level of first-order ethics).

As far as I can tell, there has been no explicit defense of the phenomenal value theory. However, there are a few philosophical literatures that are relevant to the theory, and a few other theories that are closely related or structurally similar. An obvious connection is to theories of ethical knowledge. Most such theories compare ethical knowledge to some other kind of knowledge: for example, to mathematical knowledge (where we can know certain ethical truths just by intuition), perceptual knowledge (where we have an ethical sense, akin to perception), scientific knowledge (where we apply reasoning to our ethical intuitions to reach reflective equilibrium), analytic knowledge (where we acquire ethical knowledge just by understanding the meanings of ethical terms), or propositional attitude knowledge (where we acquire ethical knowledge by understanding the contents of our evaluative attitudes).<sup>9</sup> By contrast, this paper argues that we can acquire knowledge of the values of our experiences solely on the basis of knowledge of the phenomenal characters of our experiences. In this respect, this paper favors a *phenomenal model* of ethical knowledge (at least for value facts about experiences).<sup>10</sup>

The phenomenal value theory is largely independent of most contemporary debates at the intersection of philosophy of consciousness and value theory. These include the literature on hedonism and utilitarianism, some parts of the pleasure

---

<sup>9</sup> See Sayre-McCord [2012] for an overview of these models of ethical knowledge. Notably, there is no mention of a phenomenal model of ethical knowledge, and the only discussion of ethical knowledge from experience concerns the perceptual model.

<sup>10</sup> This is a somewhat delicate point. To be precise, I am arguing that value facts about experience can be known through *inference* (from phenomenal facts). However, the more general picture I favor is that value facts about experience can be known via just the same methods by which phenomenal facts can be known: namely, through (1) inference from more fundamental phenomenal facts, (2) introspection, and (3) imagination. This paper defends only the first claim, but mentioning these other claims provides context.

---

and pain literature, and some recent discussions on the value of consciousness itself.<sup>11</sup> Whereas the phenomenal value theory is a metaethical thesis, nearly all of the work within the debates above falls within first-order ethics. Though philosophers working on these issues sometimes make claims suggestive of the phenomenal value theory, there has been little work explicitly developing the view.

A few philosophers have explicitly argued for positions similar to the phenomenal value theory. Most notably, some have argued that phenomenal facts analytically entail value facts about experiences: for example, that the concept BAD is built into the concept PAIN.<sup>12</sup> However, I will later explain why I think we can always conceptually distinguish descriptive properties from value properties, and I will later argue that the right way to capture the relation between phenomenal facts and value facts about experiences is in terms of epistemic entailment rather than in terms of analyticity. Though this difference may strike some as subtle, it has important ramifications: most notably, only the phenomenal value theory entails that there are genuine counterexamples to the epistemic gap.

Finally, there are structural parallels between the phenomenal value theory and other theories connecting phenomenal facts to facts of another domain. Just as (1) phenomenal conservatives hold that certain epistemic facts are explained by phenomenal facts, (2) phenomenal intentionalists hold that certain intentional facts are explained by phenomenal facts, and (3) phenomenal powers theorists hold that certain causal powers facts are explained by phenomenal facts,<sup>13</sup> so too (4) phenomenal value theorists hold that certain value facts are explained by phenomenal facts. Each of these positions is independent of the others, but they share a common structure. Seeing the parallels can help in understanding the phenomenal value theory.

---

<sup>11</sup> See Bramble [2016] and Deijl [forthcoming] for defenses of hedonism and experientialism, Kahane [2009, 2010] and Bradford [2020] on the badness of pain, and Siewert [1998], Lee [2018], and Kriegel [2019] on the relationship between consciousness and value.

<sup>12</sup> See Von Wright [1963], Mendola [1990], Hewitt [2008], and Massin [forthcoming].

<sup>13</sup> For phenomenal conservatism, see Huemer [2001]. For phenomenal intentionalism, see Kriegel [2013]. For the phenomenal powers view, see Mørch [2018].



---

## Motivations

An initial motivation for the phenomenal value theory is that it aligns with common sense. If you are asked how you know pain is bad, the natural response is to cite how pain feels. Perhaps you might appeal to the fact that pain is disliked—a view I will discuss in the next section—but at first pass, it is natural to think that we dislike pain because of how it feels and that the way pain feels explains both why we dislike it and how we know that it is bad. And perhaps you might note that pain is distracting and associated with bodily damage, but these seem to point to instrumental badness rather than intrinsic badness. If we are concerned with how we know that pain is intrinsically bad, there seems nothing more relevant than the fact that pain hurts. If someone is skeptical of that answer, then the most convincing method of persuasion may be to remind them of how pain feels.

A second motivation is that denying the phenomenal value theory leads to counterintuitive consequences. To deny the theory, then one must hold that a subject could know that pain hurts yet not be in a position to know that pain is bad. But it is implausible that such a situation is possible, at least if we stipulate that the subject possesses the relevant concepts and the right kinds of reasoning capacities and that their pain feels the way that ordinary pains do. And as before, if you are unmoved by the example of pain, then we could instead substitute in a different kind of experience, such as torturous pain, sickly nausea, or horrific disgust. Even if all one knows about those experiences is what it is like to have them, it seems that that is already enough to know that those experiences are bad.

There may be temptation to claim that knowledge of phenomenal facts necessarily entails whatever factor enables knowledge of the corresponding value facts. Suppose, for example, that one favors intuitionism, according to which all ethical knowledge is acquired on the basis of intuition. Then one might accept that anyone who knows that pain hurts is in a position to know that pain is bad but contend that this is because necessarily anyone who knows that pain hurts has the intuition that pain is bad. However, this picture is more complex, since it posits the intuition as an intermediary state (between knowledge of how pain feels and knowledge that pain is bad). Moreover, a theorist who favors this picture would have to explain why there should be a necessary connection between knowledge of how pain feels and the intuition that pain is bad. In fact, the connection is particularly unobvious, since

---

those two mental states differ not only with respect to their contents (that pain hurts vs. that pain is bad) but also with respect to the attitudes (knowledge vs. intuition). The simpler hypothesis is that we know pain is bad just on the basis of knowing how pain feels.

A third motivation concerns doxastic asymmetries. Though there is little consensus amongst philosophers over which value claims are true, it is nearly universally believed that pain is bad. There are exceptions, of course, but the fact that there are outliers does not negate the genuine pattern. Even philosophers moved by arguments for global skepticism about ethical knowledge tend to be reluctant to deny that they know pain is bad. And while objective list theorists, desire-satisfactionists, and experientialists have many disagreements about what kinds of things are valuable, there is consensus across these theories that pain is bad (and pleasure is good). Although these doxastic asymmetries need not be taken to be epistemically probative, they do call out for explanation. A natural explanation is that the asymmetries exist because there is a source of knowledge for the fact that pain is bad that does not apply to other kinds of ethical facts.

A fourth motivation concerns hedonic inversion. A *hedonic invert* is a subject who has the same kinds of hedonic experiences as ordinary humans (their pains hurt just as much as yours and mine), but who forms the opposite value beliefs on the basis of their hedonic experiences. Whereas any ordinary subject would form the belief that pain is bad on the basis of how pain feels, the hedonic invert instead forms the belief that pain is good on that very same basis. Intuitively, there is something deeply epistemically defective about the hedonic invert. The phenomenal value theory provides a diagnosis of the epistemic defect. The exact nature of the defect depends on how one chooses to precisify the notion of epistemic entailment. But the basic idea is that the hedonic invert is in a position to know that pain is bad (because they know how pain feels), yet they instead form the opposite belief. Putting it in other terms, the hedonic invert forms a belief that contradicts what they have justification for believing, that is counter to what their evidence supports, that results in believing a mutually inconceivable set of propositions, and that is not only rationally impermissible but is such the negated belief is rationally obligatory. These all strike me as plausible analyses of the hedonic invert's epistemic condition.

Suppose we accept this diagnosis of hedonic inversion. Do other cases of subjects with unusual ethical beliefs have the same epistemic structure? There is

---

reason to think the answer is ‘no’. At least, a number of philosophers have argued for the possibility of ideally rational *ethical eccentrics*, or subjects with ethical beliefs very different from our own (e.g. an amoralist who believes that torturing people for fun is permissible).<sup>14</sup> To my knowledge, however, nobody has defended the possibility of ideally rational hedonic inverts. Suppose, then, that hedonic inverts are necessarily irrational but that other kinds of ethical eccentrics (that do not possess unusual value beliefs about experiences) are possibly rational. Then the phenomenal value theory explains this asymmetry: even though phenomenal facts epistemically entail value facts about experiences, they do not epistemically entail other kinds of ethical facts. Consequently, hedonic inverts are rationally compelled to deny that pain is good since they know how pain feels, but other kinds of ethical eccentrics may not be rationally compelled to deny their own unusual ethical beliefs.

These motivations elicit a positive case for the phenomenal value theory. Nevertheless, I think the most interesting philosophical action occurs when we consider whether the theory is defensible in the face of objections. Those sympathetic to the epistemic gap may concede that the phenomenal value theory is intuitive yet still contend that there *must* be some objection that shows the theory is untenable.

## §2 | Objections

The rest of the paper defends the phenomenal value theory from objections. Along the way, I will also address some further issues about how to best characterize the epistemic gap.

### The Self-Evidence Objection

The self-evidence objection is that it is self-evident that pain is bad. If that is right, then it seems mistaken to think that the fact that pain hurts epistemically entails the fact that pain is bad; instead, the conclusion is knowable even without any base premises whatsoever. The standard way of characterizing self-evident propositions is as propositions such that anyone who understands them is in a position to

---

<sup>14</sup> For discussion of ideally coherent ethical eccentrics, see Street [2009]. One of the cases Street discusses is Future Tuesday Indifference, which concerns beliefs about pain. However, that case is structurally different from hedonic inversion, and it is possible to accept both my claims about hedonic inversion and Street’s claims about Future Tuesday Indifference.

---

know they are true.<sup>15</sup> There are two relevant precisifications of this idea, which generate two different versions of the self-evidence objection.

The first version claims that the sentence ‘pain is bad’ is analytic, meaning that the truth of ‘pain is bad’ follows from the definitions of the terms.<sup>16</sup> Or, framing the idea in terms of concepts, this version claims that the concept BAD is contained in the concept PAIN. However, the standard tests for analyticity support the conclusion that ‘pain is bad’ is not analytic. First, the denial of an analytic fact should seem nonsensical, but those who deny that pain is bad are not speaking nonsensically—instead, they are speaking falsely. Second, the assertion of an analytic fact should seem cognitively insignificant, but those who claim that pain is bad seem to be making a substantive (even if obvious) claim. Third, analytic truths can be translated into logical truths through the substitution of synonyms, but there seems no way to translate ‘pain is bad’ into a logical truth through such a method. Finally, ethical nihilists are competent users of the term ‘pain’ but do not deny that pains exist.

The second version of the self-evidence objection claims that it is merely a priori (rather than analytic) that pain is bad, meaning that anyone competent with the concepts PAIN and BAD is in a position to know that pain is bad. Now, suppose that we grant that it is a priori that pain is bad.<sup>17</sup> To acquire a piece of knowledge a priori, one must possess the relevant concepts. Consequently, we can ask what kind of concept of pain is needed in order to know a priori that pain is bad. Suppose that one possesses a purely functional concept of pain that yields no knowledge of how pain feels (but which still enables one to think thoughts that refer to pain). Even if it

---

<sup>15</sup> See DePaul & Hicks [2016] for discussion of self-evident propositions in metaethics.

<sup>16</sup> The term ‘analytic’ is ambiguous in contemporary philosophy. In some contexts, ‘analytic’ is used narrowly, to mean (roughly) truth in virtue of the definition of the terms. In other contexts, ‘analytic’ is used broadly, to mean (roughly) truth in virtue of meaning as opposed to form. I will reserve the term ‘analytic’ for the narrow sense and use the term ‘a priori’ to cover the broader sense. For more detailed discussion of these notions (and some reason to favor the usage of terms adopted here), see Russell [2017] and Rey [2018].

<sup>17</sup> Is knowledge that pain is bad indeed a priori? This question is tricky, but my own view is that the answer is ‘no’. In brief, I think non-testimonial justification for believing that pain is bad requires introspective or imaginative knowledge of how pain feels, and that any knowledge of phenomenal character acquired on the basis of introspection or imagination must be a posteriori.

---

is a priori that pain is bad, it is plausible that such a purely functional concept of pain does not enable one to acquire that a priori knowledge.<sup>18</sup> What more is needed? It is natural to think that acquiring a priori knowledge that pain is bad requires a concept of pain that yields knowledge of how pain feels. But this means that even if it is a priori that pain is bad, it remains the case that knowing that pain is bad requires knowing how pain feels. Consequently, the phenomenal value theory still captures the epistemic structure of how we know pain is bad.

### The Validity Objection

The validity objection is that in order to know that pain is bad, one must know not only the minor premise (that pain hurts) but also the conditional premise (that if pain hurts, then it is bad). However, the conditional premise contains a value term (namely, 'bad'), so we do not have a genuine counterexample to the epistemic gap. More generally, one might observe that any unconditional evaluative conclusion will require a conditional premise containing a value term.

Suppose, per reductio, that an epistemic gap occurs between F-facts and G-facts whenever no set of purely F-premises logically entails a G-conclusion. Then epistemic gaps would be ubiquitous. Since an inference from premise  $Fx$  to conclusion  $Gx$  is never logically valid, deriving the conclusion  $Gx$  will always require the conditional premise  $Fx \supset Gx$ . But since F and G are arbitrary predicates, we can substitute in any terms whatsoever for F and G to yield an epistemic gap. However, epistemic gaps are not ubiquitous, so we should reject the initial supposition that lack of logical entailment suffices for an epistemic gap.

What about a version of the validity objection that appeals to the disjunction of logical or analytic validity (rather than merely logical validity) as the criterion for bridging an epistemic gap? The problem with such an approach is that analytic entailments are rare whereas a priori entailments are rife. Consider again the standard tests for analyticity mentioned previously: negations of analytic facts seem nonsensical, assertions of analytic facts seem cognitively insignificant, and substitution of synonyms can translate analytic facts into logical facts. These tests concern facts rather than inferences, but we can apply the tests to inferences by examining the

---

<sup>18</sup> See Kahane [2010] for argument that someone who does not know how pain feels is not in a position to know that pain is bad.

---

relevant conditionals: for example, we can assess whether an inference from  $Fx$  to  $Gx$  is analytic by assessing whether the fact  $Fx \supset Gx$  is analytic. When we do so, it is easy to find cases of epistemic entailment that do not satisfy these tests. As an example, let  $P$  be a microphysical fact specifying the total microphysical state of the universe and let  $Q$  be the fact that at least one person exists: it is plausible that  $P$  epistemically entails  $Q$  but implausible that  $P$  analytically entails  $Q$ .<sup>19</sup>

The general lesson is that epistemic gaps ought to be characterized in epistemic terms, rather than in logical or semantic terms. Because of this, epistemic entailment (rather than rather logical or analytic validity) is the relevant relation when investigating epistemic gaps.

### The Closure Objection

The closure objection is that phenomenal facts epistemically entail only other phenomenal facts. To support this objection, one might appeal to the common idea that there are different domains of facts—normative, descriptive, phenomenal, physical, etc.—that are conceptually isolated from one another, and contend that epistemic entailment cannot bridge conceptually isolated domains.

As a reminder, a value fact is a fact that ascribes only value properties, a phenomenal fact is a fact that ascribes only phenomenal properties, and in general an  $x$ -fact is a fact that ascribes only  $x$ -properties. This categorization scheme preserves the idea that some domains are conceptually isolated from other domains: for example, it remains plausible that physical facts do not epistemically entail phenomenal facts. Nevertheless, there are good reasons for thinking that phenomenal facts epistemically entail facts of other kinds.

I have already mentioned some theories that contradict the closure objection. In particular, phenomenal conservatives hold that phenomenal facts epistemically entail certain epistemic facts, phenomenal intentionalists hold that phenomenal facts epistemically entail certain intentional facts, and phenomenal powers theorists hold that phenomenal facts epistemically entail certain causal power facts. In order for the closure objection to work, all of these views must be false. This means that anyone who has some credence in some of these views should temper their credence

---

<sup>19</sup> This conclusion aligns with standard understandings of a prioricity, analyticity, and logical validity. For further discussion, See De Paul & Hicks [2016], Russell [2017], and Rey [2018].

---

in the closure objection. Nevertheless, since each of the aforementioned views is controversial, I do not wish to lean too much on them.

Consider instead *structural facts*, or facts ascribing only purely structural properties (such as parthood, quantity, or dimensionality properties) and that are directly captured by formal representations (such as mathematical models). Since structural facts do not ascribe phenomenal properties, they are not phenomenal facts. But structural facts about experiences are epistemically entailed by phenomenal facts. Consider how color experiences can be represented via points in a bounded three-dimensional space (with dimensions representing hue, saturation, and brightness) such that color experiences that are more similar to each other correspond to points that are closer in the space.<sup>20</sup> These formal models capture various structural facts about color experiences, such as facts about their similarity structure. And these structural facts are epistemically entailed by phenomenal facts: by knowing facts about what it is like to see various colors, one is in a position to know that those color experiences stand in the structural relations captured by the model.

In response, the objector might contend that structural facts about experiences are just further phenomenal facts. But consider how we can ask whether formal models of color experiences also capture the structural relations amongst perceptible colors, whether spatial experience and physical space have the same geometrical properties, or more generally whether experiences instantiate the same structural properties as other kinds of things. Since these questions are intelligible, it is conceivable that the structural properties instantiated by experiences are also instantiated by other kinds of things. That means that those structural properties are not phenomenal properties (since, by definition, phenomenal properties can be instantiated only by experiences).

The objector might point out that the relevant structural facts are still facts about experiences. But that is irrelevant in the present context, since the fact that pain is bad is also a fact about experiences. If we were to instead categorize facts by their referents, then the phenomenal value theory would no longer be in tension with the epistemic gap, since the fact that pain hurts and the fact that pain is bad have the same referent. To draw out the tension between the phenomenal value theory and the epistemic gap, we need to categorize facts by appeal to the properties

---

<sup>20</sup> See Lee [forthcoming] for more on the formal structure of quality-space models.

---

they ascribe. Since facts ascribing only phenomenal properties can epistemically entail facts ascribing other kinds of properties, the closure objection does not work.<sup>21</sup>

### **The Rationality Objection**

The rationality objection is that the phenomenal value theory cannot account for the fact that some philosophers rationally deny that pain is bad. Since everyone knows how pain feels, and since the inference from ‘pain hurts’ to ‘pain is bad’ is easy to make, the phenomenal value theory entails that everyone is in a position to know that pain is bad. Yet some philosophers nevertheless deny that pain is bad.

In response: philosophy sometimes leads us to radical conclusions that seem philosophically compelling but turn out to be false. Consider how illusionists about consciousness deny that consciousness exists, how color eliminativists deny that colors exist, how mereological nihilists deny that macroscopic objects exist, how skeptics deny that we have any knowledge of the external world, how folk-psychology eliminativists deny that we have beliefs and desires, or how ontic structuralists deny that things exist.<sup>22</sup> I am not suggesting that any of these views is in fact false (except illusionism), but I know of nobody who thinks that all such views are true. No matter your philosophical commitments, you almost certainly believe that in some cases theoretical considerations lead philosophers to deny facts that they are in a position to easily know.<sup>23</sup>

---

<sup>21</sup> What if the objector concedes that phenomenal facts can explain facts of other kinds but still maintains that phenomenal facts cannot explain ethical facts? This version of the objection is close to the just-too-different objection against ethical naturalism advanced by Enoch [2011]. By itself, this claim is close to simply denying the phenomenal value theory, so independent support is needed in order for the objection to not be question-begging. I have already argued against the most natural way of supporting the objection. Moreover, I suspect that many will find the intuition behind the just-too-different objection to be less compelling when considering value facts about experiences.

<sup>22</sup> See Frankish [2016] on illusionism, Holman [2002] on color eliminativism, van Inwagen [1990] on mereological nihilism, Churchland [1981] on folk-psychology eliminativism, Stroud [1984] on skepticism, and French & Ladyman [2011] on ontic structuralism.

<sup>23</sup> Note that being in a position to easily know that P need not mean that it is easy to justify the claim that one knows that P. Consider how many philosophers think that we are in a



---

Philosophers who endorse these sorts of views typically concede that their position is counterintuitive but argue that the balance of theoretical considerations favors their view nonetheless. And of course, it is uncontested that theoretical considerations should sometimes outweigh common sense. But it is also possible to be mistaken about how strong the countervailing theoretical considerations are, and as a result it is unsurprising that philosophers sometimes deny claims that they are in a position to easily know. In my view, this is what is going on when philosophers deny that pain is bad. Just as there are clever arguments to the false conclusion that consciousness does not exist, so too there are clever arguments to the false conclusion that pain is not bad.

The rationality objection is also undercut by distinguishing ideal rationality from non-ideal rationality. On the one hand, the phenomenal value theory entails that philosophers who deny that pain is bad are not ideally rational. But ideal rationality is an impossible standard to achieve, and it is plausible that no philosopher has an ideally rational set of philosophical beliefs. On the other hand, the phenomenal value theory allows that philosophers arguing for counterintuitive positions still satisfy the non-ideal standards of rationality that are more pertinent to ordinary human life. This means we can reconcile the fact that everyone is in a position to easily know that pain is bad with the fact that philosophers who deny that pain is bad are not irrational by ordinary standards.

### **The Dislike Theory Objection**

The dislike theory objection is that we know pain is bad by knowing that we dislike it. There are two versions of this objection that must be addressed separately. The first version contends that pain hurts because it is disliked (and by consequence pain is bad because it is disliked). The second version contends that pain is bad because it is disliked (but not because hurting is a matter of being disliked).<sup>24</sup>

---

position to easily know that there are external objects even though it is difficult to justify the claim that we have such knowledge.

<sup>24</sup> See Street [2006] for an instance of the dislike theory objection. See Heathwood [2007], Bramble [2013], and Lin [2018] on the relationship between pleasure and desire, and Kahane [2009] Bain [2017], and Bradford [2020] on the relationship between the badness of pain, the phenomenal character of pain, and the disliking of pain.

---

Suppose first that what it is for an experience to hurt is for it to be disliked. Then we have a reductive analysis of the nature of hurting. But that reductive analysis is compatible with thinking that the fact that pain hurts epistemically entails the fact that pain is bad. In particular, one can accept the analytic claim that what it means for an experience to hurt is for it to feel a certain way while also endorsing the substantive claim that what it is for an experience to hurt is for that experience to be disliked. In light of this, those who contend that pain is bad because it is disliked can accept that pain is bad because it hurts, so long as they also think that pain hurts because it is disliked.

There may be a temptation to respond by contending that there is no distinctive phenomenology of hurting. This harkens to the debate in the pleasure and pain literature about whether unpleasantness is a phenomenal property or an attitudinal property. However, that debate concerns whether all unpleasant experiences have some phenomenal property in common, an issue that is irrelevant in the present context. Suppose there are many different ways for an experience to feel unpleasant and that what makes any given experience count as unpleasant is that the experience is disliked (rather than the instantiation of a common phenomenal property). Then there is no single phenomenal property instantiated by all unpleasant experiences. But one could still think that for any given unpleasant experience, the fact that it is bad is epistemically entailed by the fact that it feels the way that it does.

The other version of the dislike theory objection denies the claim that hurting is a matter of being disliked, as well as the claim that pain is bad because it hurts, yet nevertheless claims that pain is bad because it is disliked. Unlike the previous version, this version of the dislike theory is incompatible with the phenomenal value theory. But while the previous version was able to accommodate the datum that pain is bad because of how it feels, the current version is forced to deny it. This strikes me as a decisive cost, since it is extremely implausible that the way pain feels has no essential connection to its badness. Moreover, taking dislike to be independent of hurting commits one to the dubious consequence that it is possible for a subject to know how pain feels yet not be in a position to know that pain is bad.

### **The Debunking Objection**

The debunking objection is that we believe pain is bad because it was evolutionarily advantageous, rather than because it is true. If we can explain why we

---

believe pain is bad without appeal to any facts about the badness of pain, then that seems to undercut the motivation for the phenomenal value theory.<sup>25</sup>

In order to construct a debunking argument, one needs more than merely the premise that there is an evolutionary explanation for why we have the beliefs under consideration. If that were the only requirement, then nearly any set of beliefs whatsoever would be debunkable. So, what else is required for a debunking argument to succeed? Though there is disagreement about the details, one condition that is generally accepted is that it must be conceivable for the relevant facts to have been different.<sup>26</sup> The role of the conceivability premise is to show that even if the relevant facts had been different, we would still have had the same beliefs we actually do (because our beliefs are determined by what is evolutionarily advantageous rather than what is true).

The phenomenal value theory challenges this background premise of the debunking objection. At least, on my preferred way of understanding epistemic entailment, if  $P$  epistemically entails  $Q$  then it is inconceivable that  $P \ \& \ \neg Q$ . This means that if the phenomenal value theory is true, then it is inconceivable for the phenomenal facts to hold without the relevant value facts also holding. Note that, as discussed previously, accepting that it is inconceivable that pain is not bad (given knowledge of how pain feels) does not entail that those who deny that pain is bad are making nonsensical claims or that they are irrational by the ordinary standards of

---

<sup>25</sup> See Street [2006] for a classic example of a debunking argument against ethical beliefs.

<sup>26</sup> See Kahane [2011], Clarke-Doane [2012], and Vavova [2015] for some recent discussions of evolutionary debunking arguments. Notably, Clarke-Doane argues that peer disagreement amongst conceptually competent users over  $P$  is evidence that both  $P$  and  $\neg P$  are intelligible, which might seem to be evidence against my claim that it is inconceivable that pain is not bad. However, it seems to me that Clarke-Doane's argument deploys two distinct senses of 'intelligibility'. On the one hand, it is plausible that peer disagreement over  $P$  is good evidence for the analytic intelligibility of both  $P$  and  $\neg P$  (i.e. neither  $P$  nor  $\neg P$  is analytically false). However, that is irrelevant in the present context since I do not claim that 'pain is not bad' is analytically false. On the other hand, it is implausible that peer disagreement over  $P$  is good evidence for the conceivability of both  $P$  and  $\neg P$  (i.e. neither  $P$  nor  $\neg P$  can be ruled out a priori). After all, there are plenty of philosophical debates involving peer disagreement over truths that are presumably knowable a priori. See also my responses to the self-evidence objection (p.10) and the rationality objection (p.15).

---

rationality. Now, the debunker could respond by denying that it is inconceivable that pain is not bad, given knowledge of how pain feels. But this simply moves the dialectic back to the question of whether the phenomenal value theory is plausible in the first place. The upshot is that there is no debunking argument against the phenomenal value theory that does not presuppose its falsity.

### **The Generalization Objection**

The generalization objection is that the phenomenal value theory fails to properly generalize beyond the core example of the badness of pain. We have all sorts of phenomenal knowledge about our experiences, but we are often unsure of what to think about the values of those experiences. Perhaps focusing on how we know pain is bad stacks the deck in favor of the phenomenal value theory.

I think the generalization objection has a false starting point. For the vast majority of our experiences, we are able to know at least approximately how good or bad that experience is. As examples, consider what it was like for you to eat the last meal you ate, or what it was like for you to last exercise, or what it was like for you to last watch a movie or listen to music or play a game. I suspect you will easily know approximately how good or bad those experiences were. In fact, it is hard to think of any experiences where we know what it is like to have the experience yet where we do not have even approximate knowledge of how good or bad that experience is. Even when an experience is close to value neutral, we can usually know that it is close to value neutral (even when we do not know whether it has net positive value or net negative value).

These observations acknowledge that we do not always have exact knowledge of the values of our experiences. However, it is likewise plausible that we do not always have exact knowledge of the phenomenal characters of our experiences. Even those most optimistic about our introspective capacities tend to accept that our phenomenal judgments are less precise and reliable when targeting aspects of experiences that are detailed, subtle, or complex.<sup>27</sup> Consider, for example, questions about the precise character of your emotional experience, or exactly how many dimensions of variation your olfactory experiences have, or whether your temporal

---

<sup>27</sup> See Block [1995], Schwitzgebel [2006], Williamson [2000], and Lee [2019] for various arguments that our knowledge of our own experiences is limited.

---

experience is continuous or discrete. Given that our phenomenal knowledge of phenomenal character is often approximate, it is unsurprising that the corresponding evaluative knowledge is also often approximate.

## Conclusion

This paper may strike some as having developed an extended defense of something that was already obvious. To some extent, I agree with that sentiment. However, the epistemic gap also strikes many as obvious: in fact, there is arguably more philosophical literature on how to best formulate the epistemic gap than on whether the thesis is true in the first place. Yet these apparently obvious theses—about how we know pain is bad and about the epistemic gap—are in tension with each other. Although I have framed this paper as motivating and defending the phenomenal value theory, I think it is more perspicuous to think of the paper as identifying an apparent dilemma and explaining why we need not fear one of the horns.

This paper has adopted a *bottom-up* methodology, whereby we examine a particular domain and identify the metaethical principles that are most plausible for that domain. By contrast, most work in metaethics adopts a *top-down* methodology, whereby we search for principles and theories that apply universally to all domains. Though I think the top-down methodology is valid and fruitful, maintaining an exclusively top-down perspective risks overlooking details that distinguish one domain from another. Our panoramic view from the top of a mountain is different from our immediate view when standing on that very terrain.

In light of this, the phenomenal value theory yields a picture where the metaethics of consciousness is special. Even if we accept the epistemic gap for many kinds of ethical facts, there is a special class of ethical facts—namely, value facts about experiences—where the epistemic gap is surmountable. Though this inference may be surprising when considered in the abstract, it is compelling and defensible upon close examination. As a result, the best answer to how we know pain is bad turns out to be also the simplest and most obvious: by knowing how it feels.

---

**REFERENCES**

Aydede, Murat, "Pain", *The Stanford Encyclopedia of Philosophy* (Spring 2013 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2013/entries/pain/>>.

Bain, David (2014). Pains that Don't Hurt. *Australasian Journal of Philosophy* 92 (2):1-16.

Bain, David (2017). Why Take Painkillers? *Noûs* 2017 (2):462-490.

Block, Ned (1995). On a confusion about a function of consciousness. *Brain and Behavioral Sciences* 18 (2):227–247.

Bradford, Gwen (2020). The badness of pain. *Utilitas* 32 (2):236-252.

Bramble, Ben (2013). The distinctive feeling theory of pleasure. *Philosophical Studies* 162 (2):201-217.

Chalmers, D. (2012). *Constructing the World*. New York: Oxford University Press.

Churchland, Paul M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy* 78 (February):67-90.

Clarke-Doane, Justin (2012). Morality and Mathematics: The Evolutionary Challenge. *Ethics* 122 (2):313-340.

DePaul, Michael and Hicks, Amelia, "A Priorism in Moral Epistemology", *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2016/entries/moral-epistemology-a-priori/>>.

Enoch, D. 2011. *Taking Morality Seriously*, Oxford: Oxford University Press.

Fine, Kit. (2018). Truthmaking and the is–ought gap. *Synthese*. 10.1007/s11229-018-01996-8.

Fodor, J. A. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese* 28 (2):97-115.

- 
- Frankish, K. (2016). Illusionism as a Theory of Consciousness. *Journal of Consciousness Studies* 23 (11-12):11-39.
- French & Ladyman, 2011. "In Defence of Ontic Structural Realism," in A. Bokulich and P. Bokulich (eds.) 2011, pp. 25–42.
- Grahek, Nikola (2007). *Feeling Pain and Being in Pain* (2nd Edition), MIT Press.
- Heathwood, Chris (2007). The reduction of sensory pleasure to desire. *Philosophical Studies* 133 (1):23-44.
- Hewitt, S. (2008) *Normative Qualia and a Robust Moral Realism*. PhD Dissertation, New York University.
- Holman, Emmett L. (2002). Color eliminativism and color experience. *Pacific Philosophical Quarterly* 83 (1):38-56.
- Jackson, Frank (1974). Defining the autonomy of ethics. *Philosophical Review* 83 (1):88-96.
- Jackson, Frank, 1998, *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: Clarendon Press.
- Kahane, Guy (2009). Pain, dislike and experience. *Utilitas* 21 (3):327-336.
- Kahane, Guy (2010). Feeling pain for the very first time: The normative knowledge argument. *Philosophy and Phenomenological Research* 80 (1):20-49.
- Kahane, Guy (2011). Evolutionary Debunking Arguments. *Noûs* 45 (1):103-125.
- Klein, Colin (2015). What Pain Asymbolia Really Shows. *Mind* 124 (494):493-516.
- Kriegel, Uriah (2013). The Phenomenal Intentionality Research Program. In U. Kriegel (ed.), *Phenomenal Intentionality*. Oxford University Press.
- Lee, Andrew Y. (2019). The Microstructure of Experience. *Journal of the American Philosophical Association*:1-20.

- 
- Lee, Andrew (forthcoming). Modeling Mental Qualities. *Philosophical Review*.
- Lin, Eden (2018). Attitudinal and Phenomenological Theories of Pleasure. *Philosophy and Phenomenological Research*.
- Maguire, Barry (2017). The Autonomy of Ethics. In Tristram McPherson & David Plunkett (eds.), *The Routledge Handbook of Metaethics*. Routledge. pp. 431-442.
- Massin, Olivier (2016). Bad by Nature, An Axiological Theory of Pain. In Jennifer Corns (ed.), *The Routledge Handbook of Philosophy of Pain*. Routledge.
- Mehta, Neil (2019). Phenomenal, Normative, and Other Epistemic gaps: A General Diagnosis. *Philosophy and Phenomenological Research* 98 (3):567-591.
- Mendola, Joseph (1990). Objective value and subjective states. *Philosophy and Phenomenological Research* 50 (4):695-713.
- Moore, G.E., 1903. *Principia Ethica*, New York: Cambridge University Press.
- Mørch, Hedda Hassel (2018). The evolutionary argument for phenomenal powers. *Philosophical Perspectives* 31 (1):293-316.
- Nagel, Thomas. (1980). The limits of objectivity. *The Tanner Lectures on Human Values*, 1, 75-139.
- Nagel, Thomas (1986). *The View From Nowhere*. Oxford University Press.
- Peacocke, Christopher (2004). Moral Rationalism. *Journal of Philosophy* 101 (10):499 - 526.
- Pigden, Charles R. (1989). Logic and the autonomy of ethics. *Australasian Journal of Philosophy* 67 (2):127 – 151.
- Prior, A. N. (1960). The autonomy of ethics. *Australasian Journal of Philosophy* 38 (3):199 – 206.
- Rachels, Stuart (2000). Is unpleasantness intrinsic to unpleasant experiences? *Philosophical Studies* 99 (2):187-210.



---

Rey, Georges, "The Analytic/Synthetic Distinction", *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2018/entries/analytic-synthetic/>>.

Russell, Bruce, "A Priori Justification and Knowledge", *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2017/entries/apriori/>>.

Sayre-McCord, "Moral Epistemology," in *The International Encyclopedia of Ethics*, edited by Hugh La Follette (Blackwell, 2012), pp. 1674-1688.

Sayre-McCord, Geoff, "Metaethics", *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2014/entries/metaethics/>>.

Schroeder, Mark, "Value Theory", *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2016/entries/value-theory/>>.

Schwitzgebel, Eric (2006). The unreliability of naive introspection. *Philosophical Review* 117 (2):245-273.

Singer, Daniel J. (2015). Mind the Is-Ought Gap. *Journal of Philosophy* 112 (4):193-210.

Sparks, Jacob (2018). Is, Ought, and the Regress Argument. *Australasian Journal of Philosophy* 97 (3):528-543.

Street, Sharon (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies* 127 (1):109-166.

Street, Sharon (2009). In defense of future Tuesday indifference: Ideally coherent eccentrics and the contingency of what matters. *Philosophical Issues* 19 (1):273-298.

Stroud, B., 1984, *The Significance of Philosophical Scepticism*, Oxford: Clarendon Press.

Suikkanen, Jussi (2016). Naturalism in Metaethics. In Kelly James Clark (ed.), *Blackwell Companion to Naturalism*. Wiley-Blackwell. pp. 351-368.

van Inwagen, P (1990). *Material Beings*, Ithaca (NY): Cornell University Press.

Vavova, Katia (2015). Evolutionary Debunking of Moral Realism. *Philosophy Compass* 10 (2):104-116.

Von Wright, G. (1963) *The Varieties of Goodness*. London: Routledge & Kegan Paul.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.