

# Metaethical Experientialism

ANDREW Y. LEE

Australian National University, Philosophy

## Abstract

I develop and defend *metaethical experientialism*, the thesis that phenomenal facts epistemically and metaphysically entail value facts. I argue, for example, that anyone who knows what it's like to feel extreme pain is in a position to know that that kind of experience is bad. I explain how metaethical experientialism yields genuine counterexamples to the principle that no ethical conclusion can be derived from purely descriptive premises. I also discuss the prospects for a pluralistic metaethics, whereby different classes of ethical facts are subject to different kinds of metaethical analyses.

## Introduction

Imagine what it's like for you to burn your hand, break your leg, experience nausea, or feel heartbreak. It's very natural to think that knowing what those experiences feel like puts you in a position to know that those experiences are bad. But that common-sensical claim is in tension with a widely endorsed metaethical principle:

THE EPISTEMIC GAP<sup>1</sup>

No ethical conclusion is derivable from purely descriptive premises.

---

<sup>1</sup> This principle more commonly goes by other names. But—at least in the present context—the label 'epistemic gap' is better than the alternative labels. Other common names include 'is-ought gap' (but my concern will be with value facts, rather than with deontic facts), 'fact-value gap' (but 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 reduction relations between properties, rather than entailment relations between facts).

---

The epistemic gap strikes many as compelling. How could premises merely about how things are ever suffice for conclusions about which things matter? Although most agree that ethical facts supervene on descriptive facts, few think that ethical facts are epistemically entailed by descriptive facts. And while there are some famous counterexamples to simple formulations of the epistemic gap, these are widely regarded as exposing technical problems with the formulation rather than as challenging the core idea.<sup>2</sup>

I think the epistemic gap is false. Maybe it holds for many kinds of ethical claims. But I think there's a special class of ethical claims that are genuinely derivable from purely descriptive premises. Here's my core thesis:

METAETHICAL EXPERIENTIALISM

Value facts about experiences are explained by phenomenal facts.

By *facts*, I mean true propositions. 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's like to have certain kinds of experiences. By F-facts *explain*<sup>3</sup> G-facts, I mean that it's both the case that (a) F-facts epistemically entail G-facts (where this might be understood as a priori entailment), and (b) F-facts metaphysically entail G-facts (where this might be understood as grounding).

My formulation of metaethical experientialism purposefully leaves open *which* phenomenal facts and *which* value facts are relevant. The most

---

<sup>2</sup> For some classic papers concerning these technical problems, see Prior [1960], Jackson [1974], and Pidgen [1989]. For more recent discussions, see Singer [2015], and Fine [2108]. For non-technical discussions of the epistemic gap in metaethics, see Sayre-McCord [2014], Maguire [2017], and Sparks [2021]. For general discussions of epistemic gap principles (including for other domains of facts), see Jackson [1998], Chalmers [2012], and Mehta [2019].

<sup>3</sup> I'm using the term 'explains' in a stipulative way (to mean epistemic and metaphysical entailment). Obviously, there are some forms of explanation (such as causal explanation) that don't require epistemic and metaphysical entailment.

---

modest version of metaethical experientialism makes a merely existential claim: *some* phenomenal facts explain *some* value facts about experiences. The version I favor is stronger: *every* phenomenal fact that *completely* characterizes what it's like to have an experience explains *every* value fact about that experience. Since even a single counterexample to the epistemic gap suffices to show that it's false, I'll focus mostly on the modest thesis. But I'll occasionally mention the version of metaethical experientialism that I favor, at least when doing so paints a more systematic picture.

Here's the structure of the paper. §1 presents metaethical experientialism; §2 defends metaethical experientialism from a variety of objections; §3 discusses some metaethical consequences.

## §1 Metaethical Experientialism

Nearly everyone agrees that some experiences—such as the experiences of burning your hand, breaking your bone, feeling extreme nausea, or experiencing heartbreak—are bad. Even if these experiences lead to good outcomes, and even these experiences have some good features, it's hard to deny that the experiences are bad to some degree. Nearly everyone also agrees that at least some of the badness of these experiences is due to how they feel. Maybe there are also other factors that contribute to their overall badness, and maybe the way those experiences feel is itself explainable in non-phenomenal terms, but it's hard to deny that one reason that these experiences are bad is because of how they feel.

If you have residual doubts, then I recommend the following exercise:

### Exercise

1. Put a skillet on your stovetop. Set the heat to medium. Wait three minutes for the skillet to sufficiently heat up.
2. Press your hand against the surface of the skillet for several seconds. As your hand burns, consider what it's like to have that experience.

- 
3. Ask yourself: Would anyone having exactly *that* kind of experience thereby be having an experience that is bad? Would anyone who knows what *that* experience feels like thereby be in a position to know that that experience is bad?

You might quibble that a set of instructions isn't an argument, or that it won't be persuasive to someone who doesn't already accept my view (though notice you can't say both!). I'd counter: following these instructions will put you in a better epistemic position for evaluating the entailment relations between the relevant facts, and a vivid reminder of how pain feels is a powerful method of persuasion. But I don't want to fight over the epistemic force of this exercise. If you aren't convinced even after directly confronting the phenomenal character of your own pain experiences, then there may be nothing I can say to sway you. But my aim isn't to convince the skeptic. Instead—following Nagel [1980: 109]—my aim is to get “rid of the obstacles to the admission of the obvious.” Think of metaethical experientialism as a starting point. The question I want to focus on is whether there are any compelling arguments that force one off that starting point.

### **'Pain' and 'Bad'**

Pain is the canonical example of a bad experience. For the purposes of this paper, I'll use it as my primary example.

By *pain*, I mean the ordinary nociceptive experiences had by ordinary humans. Pains *hurt*, meaning that they have an unpleasant phenomenal character. There may be non-ordinary cases of nociceptive experiences that don't hurt, but I'm using 'pain' in a stipulative way, whereby those experiences won't count as instances of pain.<sup>4</sup> Keep in mind that pain is merely one

---

<sup>4</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].

---

example that we can focus on: if desired, we could substitute in a different kind of experience, such as torturous pain, or nausea, or fear. As noted earlier, even a single counterexample to the epistemic gap is enough to show that the principle is false. If you think pain isn't the best or cleanest example, then substitute in whatever you think is.

By *bad*, I'll always mean final, pro-tanto badness (as opposed to instrumental or all-things-considered badness). I'll stay neutral on whether value facts entail deontic facts or fittingness facts (for example, you might think the fact that  $x$  is bad entails that one ought not bring about  $x$  or that it's unfitting to bring about  $x$ ). If you think there are such entailments, then accepting metaethical experientialism will commit you to phenomenal facts entailing those sorts of facts as well. Otherwise, metaethical experientialism may be understood as a thesis solely about value, leaving open questions about facts associated with these other families of ethical concepts.

The term *bad* may be interpreted as concerning either what makes an individual worse off or what makes the world in general worse off. I favor the view that what it is for a world to be worse off is for its individuals to be worse off. Given this, I think instances of badness in the first sense just are instances of badness in the second sense. But those who think that these two notions come apart are free to focus on one over the other. For brevity, though, I'll continue to use the unmodified term 'bad'.

You might ask which theory of badness I'm operating with. But that would be asking the wrong question. My claim isn't that phenomenal facts explain value facts *given* some particular theory of badness. Instead, my claim is that anyone who knows the relevant phenomenal facts and who possess the concept BAD is in a position to see that the phenomenal facts explain the target value facts. There are theories of badness that are in tension with metaethical experientialism. But that simply means that those theories are rival views; it doesn't mean that metaethical experientialism is false.

Throughout the paper, I'll take value facts about experiences to be a subclass of *ethical facts*, where ethical facts include not only facts about value,

---

but also facts about rightness, permissibility, fittingness, and anything else within the subject-matter of ethics. Although I'll argue that value facts about experiences are counterexamples to the epistemic gap, I'll assume that the epistemic gap remains for other kinds of ethical facts. If that assumption is correct, then there's a source of knowledge for value facts about experiences that doesn't generalize to other kinds of ethical facts.

### **Metaethical Experientialism vs. Metaethical Hedonism**

To further characterize metaethical experientialism, I'll contrast it with the class of views that are closest to metaethical experientialism.<sup>5</sup> This is a class of views that I'll call *metaethical hedonism*.<sup>6</sup>

Metaethical hedonists endorse *hedonism*, the claim that only pleasures/pains are good/bad. But metaethical hedonists argue for hedonism by appealing to a metaethical asymmetry: (a) our ethical beliefs about the values of hedonic experiences are immune to skeptical worries, but (b) our other ethical beliefs are susceptible to skeptical worries. The idea is that there's a source of knowledge about the values of hedonic experiences that doesn't yield knowledge about the values of other kinds of things.

There's no consensus amongst metaethical hedonists about how exactly we come to acquire knowledge of the values of hedonic experiences. However, the epistemological stories endorsed by metaethical hedonists tend to differ from the epistemological story developed in this paper. Some think that it's analytic that pain is bad and pleasure is good, but I'll argue in

---

<sup>5</sup> Metaethical experientialism might also be compared to other theories that take phenomenal facts to explain facts of another domain. Consider: *phenomenal conservatives* hold that certain epistemic facts are epistemically entailed by phenomenal facts (see CITATIONS); *phenomenal intentionalists* hold that certain intentional facts are epistemically entailed by phenomenal facts (see CITATIONS); and *phenomenal powers theorists*<sup>5</sup> hold that certain causal powers facts are epistemically entailed by phenomenal facts (see CITATIONS). Similarly, *metaethical experientialists* hold that certain value facts are epistemically entailed by phenomenal facts.

<sup>6</sup> See Wright [1963], Mendola [1990], Pearce [1995], Hewitt [2008], Massin [2016], and Sinhababu [2022]. I take the term 'metaethical hedonism' from Carlsmith [2022].

---

§2 that such claims aren't analytic. Others think that we come to know such facts through introspection, but I'll explain later in this section why metaethical experientialism isn't committed to any claims about introspective knowledge.

If you're sympathetic metaethical hedonism, then you'll probably be sympathetic to metaethical experientialism.<sup>7</sup> But even if you're skeptical of metaethical hedonism, you might still be sympathetic to metaethical experientialism. The reason is that that metaethical experientialism sidesteps many of the objections to metaethical hedonism. In what follows, I'll note some contrasts between metaethical hedonism and metaethical experientialism, explain how these contrasts make metaethical experientialism more defensible, and clarify some of the commitments of metaethical experientialism.

A preliminary remark: I'm going to use metaethical hedonism as a foil for characterizing metaethical experientialism. Because of this, I'll focus on the contrasts between the views. But it's also worth emphasizing the commonality: both views hold that we have a distinctive way of knowing value facts about experiences that renders those beliefs immune to skeptical challenges. Given this, it's fair to think of the views as allies, and the commonalities as greater than the differences.

- a. **Monism vs. Pluralism:** Metaethical hedonism is committed to hedonism. By contrast, metaethical experientialism leaves open what the basic goods/bads are.<sup>8</sup>

---

<sup>7</sup> In principle, you could endorse metaethical hedonism without endorsing metaethical experientialism. Suppose, for example, that you think both that (a) we know value facts about experiences via introspection (and that this renders those beliefs immune to skeptical worries), and that (b) phenomenal facts don't explain value facts. By endorsing (a), you endorse metaethical hedonism; by endorsing (b), you deny metaethical experientialism.

<sup>8</sup> As examples, Mendola [1990: 713] suggests that "hedonism is supported by normative realism," Rawlette [2008: 244] says that "we ought to do what will produce the greatest possible balance of positive over negative phenomenal experience," a view she calls "hedonistic utilitarianism."

---

Metaethical experientialism not only leaves open the possibility that some non-hedonic experiences are good/bad, but also the possibility that some non-experiential things are good/bad. If some non-experiential things are good/bad, then facts about those things being good/bad won't be explained by phenomenal facts. But that's compatible with thinking that value facts about experiences are explained by phenomenal facts. You could, for example, be an objective list theorist (perhaps pleasure, friendship, and knowledge are all basic goods) and a metaethical experientialist (if you accept that the way pleasure feels explains why it's good). Given this, metaethical experientialism is an ecumenical thesis that's compatible with a wide range of views about which things are good/bad.

You could even endorse metaethical experientialism while holding that some facts about the values of experiences aren't explained by phenomenal facts. Consider a view on which painful experiences are bad both because of how they feel and because they're disliked (where the badness due to each factor is non-redundant, and where the badness due to how they feel is explained by the way they feel). This view holds that only some (rather than all) facts about the values of experiences are explained by phenomenal facts. But it's still a version of metaethical experientialism.

- b. **Reduction vs. Explanation:** Many metaethical hedonists take goodness/badness to be reducible to pleasurableness/painfulness.<sup>9</sup> By contrast, metaethical experientialism is merely a claim about entailments between facts, rather than reductions between properties.

To say that value property F is reducible to phenomenal property G is to say that what it is for something to be F is for it to be G. To say that phenomenal

---

<sup>9</sup> As examples, Mendola [1990: 706] says that "value-disvalue is a distinct range of phenomenal properties," and Rawlette [2018: 72] says that "intrinsic goodness and badness are phenomenal qualities of experience."



---

fact P epistemically (or metaphysically) entails value fact Q is to say that if one knows that P (or if P obtains) then one is thereby in a position to know that Q (or that Q thereby obtains). Even if the fact that *x* is painful entails the fact that *x* is bad, it doesn't follow that what it is for something to be bad is for it to be painful. As an analogy, consider the common view that microphysical facts explain macrophysical facts even though macrophysical properties are irreducible to microphysical properties.<sup>10</sup>

You might wonder whether the epistemic gap is better understood as a thesis about reduction (between properties) instead of entailment (between facts). Well, there are many canonical formulations of the epistemic gap that appeal to entailments between facts. Here's one salient example: Sayre-McCord [2014: §4] discusses Hume's original formulation of the epistemic gap, and characterizes the principle 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." The aim of this paper is to argue that the view expressed by that claim is false, at least when we focus on value facts about experiences.

- c. **Hedonic Phenomenology:** Most metaethical hedonists hold that all hedonic experiences have some phenomenal property in common.<sup>11</sup> By contrast, metaethical experientialism leaves open questions about the nature of hedonic experiences.

---

<sup>10</sup> See Fodor [1974] for a classic argument against reductive analyses of the macroproperties studied by the special sciences. See Suikkanen [2016], Sinhababu [2018], and Ridge [2019] on reductive analyses in metaethics. See Sayre-McCord [1988] and Brink [1989] for some proponents of non-reductive metaethical naturalism.

<sup>11</sup> Mendola [1990: 704] says that pleasure is "like a sensation or feeling" and that there is a "phenomenal hedonic tone present in all cases of pleasure."

---

There are two prominent views on the question of what makes an experience pleasurable.<sup>12</sup> According to phenomenological theories, pleasurable experiences are those that instantiate a certain phenomenal property. According to attitudinal theories, pleasurable experiences are those that are liked by their subjects. The debate centers on the heterogeneity problem, or the challenge of finding something in common between the diverse range of experiences that we call ‘pleasurable’. While it may be tempting to think of metaethical experientialism as committed to a phenomenological theory, metaethical experientialism is neutral on this debate.

Suppose an attitudinal theory is correct. Then there are many different kinds of experiences that can be pleasurable, each with its own distinctive phenomenology. But even so, it remains plausible that for any particular pleasurable experience, knowing how that experience feels enables you to know how good/bad that experience is. As examples, consider the experiences of receiving a massage, eating a delicious snack, having sex, and conversing with friends. You might think that there’s nothing in common between the ways that all these experiences feel. But that doesn’t undercut the fact that for each of those experiences, knowing how it feels would enable you to know something about how good/bad it is.

- d. **Introspection vs. Inference:** Some metaethical hedonists argue that we know value facts through *introspection*.<sup>13</sup> By contrast, metaethical experientialism says merely that knowing what it’s like to have an experience enables one to know how good/bad that experience is.

Questions about introspection concern one of the processes by which we acquire knowledge of facts about our experiences. But metaethical

---

<sup>12</sup> I formulate the discussion here in terms of pleasure because most of the relevant literature on this issue focuses on pleasure. But it’s straightforward to translate the discussion into pain. See Broi [forthcoming] for discussion of the debate on the heterogeneity of pleasure.

<sup>13</sup> The most notable example is Sinhababu [2022].

---

experientialism makes a claim about what follows from phenomenal facts, regardless of how we come to know them. You could accept metaethical experientialism while thinking that introspection yields knowledge of only phenomenal facts (and that value facts aren't introspectable), or while thinking that introspective beliefs are no more epistemically secure than beliefs formed on the basis of other epistemic processes. Whatever your view on introspection, it remains plausible that knowing how an experience feels puts you in a position to know how good/bad that experience is.

You might object that all knowledge of phenomenal facts is acquired through introspection. I think there's good reason to doubt that claim.<sup>14</sup> But even if the claim were true, it would be irrelevant. Identifying how we come to know a class of facts is one thing. Identifying what those facts explain is another. As an analogy, two people could agree that microphysical facts explain macrophysical facts while disagreeing on how we acquire knowledge of microphysical facts in the first place. Similarly, two people could agree that phenomenal facts explain value facts about experiences while disagreeing about how we acquire knowledge of phenomenal facts in the first place.

## §2 Objections

I'll now address a variety of objections.

1. **Good Pains:** Some pains are good. Consider the experiences you have when eating spicy food or receiving a massage.

---

<sup>14</sup> Let Swamperson be an internal duplicate of you who spontaneously materialized a few seconds ago. Swamperson knows what it's like to feel pain, but they have never had any pains. Since one can introspect only experiences that one has actually had, Swamperson has never introspected any pains. If you're tempted to object that thinking about pain entails actually experiencing that pain, then I invite you to consider whether you'd rather think about being tortured or actually undergo torture.

---

Let's call these kinds of experiences *good pains*. Suppose we ask: What makes it the case that good pains are good? Any plausible answer will retain the idea good pains are good because of how they feel.<sup>15</sup> You could appeal to good pains to challenge the claim that all instances of pain are bad, or to question exactly which phenomenal properties are good-making or bad-making properties, or to argue that merely knowing that an experience is painful doesn't yet suffice for knowing whether it's bad to some degree. But it remains plausible that the fact that good pains feel the way they do explains why they are good.

A more general point is that metaethical experientialism is neutral on which kinds of phenomenal properties make an experience good/bad. In fact, you could even be a *holist* who thinks that to know the value of any given experience, one must evaluate the experience as a whole (as opposed to the parts or properties of the experience in isolation). The holist might deny that the fact that an experience is painful explains the fact that the experience is bad, since merely specifying that an experience is painful leaves out the rest of what it's like to have that experience. But even the holist could still accept that phenomenal facts that characterize whole experiences explain facts about the values of those experiences.

It may be useful to draw a distinction between *complete phenomenal facts*, which completely characterize what it's like to have an experience, and *partial phenomenal facts*, which only partially characterize what it's like to have an experience. You might be skeptical that partial phenomenal facts (such as the fact that an experience is painful) explain value facts about experiences. But even so, it remains plausible that complete phenomenal facts

---

<sup>15</sup> Bradford [2020] argues that good pains are good partly *because* they're unpleasant. While I think that good pains are good because of how they feel overall, I'm skeptical that good pains are good because of their unpleasantness. But those who instead agree with Bradford still ought to accept metaethical experientialism, since taking good pains to be good because of their unpleasantness is a way of taking good pains to be good because of how they feel.

---

epistemically entail value facts. For the rest of the paper, I'll take it for granted that the phenomenal facts under consideration are complete.

2. **Appearance vs. Reality:** Just because something is *experienced as* bad doesn't mean that the *experience itself* is bad.

The following claim is certainly true: just because something is experienced as F doesn't mean that the experience itself is F. But the truth of that claim doesn't undermine metaethical experientialism. To say that *a* is experienced as F is to say that one's experience represents *a* as F. But metaethical experientialism is neutral on questions about what experiences represent. You could, for example, think that pains have evaluative contents ('*a* is bad'), imperatival contents ('stop *a* from persisting'), or even no contents at all. But whatever you think, it remains plausible that painful experiences are themselves bad.

If pain merely represented badness, then it should be possible for pain to be instantiated without anything bad being instantiated. Compare: if phenomenally red experiences merely represent redness (instead of themselves being red), then it should be possible for a phenomenally red experience to be instantiated without anything red being instantiated. But imagine finding out that you've been the victim of a Cartesian demon your whole life: all of your experiences have been hallucinations induced by the demon. You might then conclude that nearly all of your experiences have been illusory: even though it appears to you that there's a red tomato in front of you, there isn't in fact any red tomato in front of you. But even in this skeptical scenario, where you might be ready to give up nearly all inferences from 'appears F' to 'is F', there's little temptation to think that nothing bad was actually instantiated whenever you had painful experiences. The badness associated with painful experiences isn't merely a matter of what's represented by those experiences. Instead, it's a matter of how those experiences feel.

3. **Validity:** To derive 'x is bad' from 'x feels painful', we need not only the minor premise ('x feels painful') but also a conditional premise ('if x feels painful, then x is bad'). But the conditional premise contains a value term ('bad'), so we don't have a genuine counterexample to the epistemic gap.

Suppose, *per reductio*, that whenever it's the case that (a) no set of purely F-premises ever logically entails a G-conclusion, it follows that (b) there's an epistemic gap between F-facts and G-facts. Well, an inference from the premise 'Fx' to the conclusion 'Gx' is never logically valid. No matter what 'F' and 'G' are, deriving the conclusion 'Gx' will require a conditional premise 'Fx → Gx'. This means we can substitute in any predicates whatsoever for 'F' and 'G' to yield an epistemic gap. But those who accept an epistemic gap between the descriptive and the ethical aren't merely appealing to the general observation that there's a logical gap whenever 'F' and 'G' are distinct predicates. Instead, there are some pairs of subject-matters (descriptive/axiological, physical/phenomenal, abstract/concrete) where one seems epistemically (rather than merely logically) isolated from the other. Therefore, we should reject the supposition that lack of logical entailment suffices for an epistemic gap. This means that the present objection misidentifies the nature of epistemic gaps. Instead of characterizing epistemic gaps in terms of logic entailment, we ought to instead appeal to epistemic entailment.

4. **Dislike:** Pain is bad because we dislike it.

To evaluate this objection, we need to ask: What is the relationship between the way pain feels and the fact that it's disliked?<sup>16</sup> There are three options:

---

<sup>16</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.

---

OPTION A: *Pain is disliked because of how it feels.* This is the most natural option: intuitively, pain is disliked because it hurts. But if (a) the fact that pain hurts explains the fact that it's disliked, and (b) the fact that pain is disliked explains the fact that it's bad, then (c) the fact that pain hurts explains the fact that it's bad. In other words, anyone who has an experience that hurts would thereby be having an experience that is bad, and anyone who knows what such an experience feels like would thereby be in a position to know that that experience is bad. This means that this first option is compatible with metaethical experientialism.

OPTION B: *Pain feels the way it does because we dislike it.* This option yields an explanation for why pain feels the way it does. But that explanation is compatible with thinking that pain is bad because of how it feels. If you think that fact that pain is disliked explains the fact that pain is bad (the dislike objection) *and* that pain feels the way it does because it's disliked (OPTION B), then you can still accept that the fact that pain feels the way it does explains the fact that it's bad (metaethical experientialism). Finding an explanation for a fact doesn't undercut the explanatory power of that fact.

OPTION C: *The way pain feels is independent from the fact that it's disliked.* This is the only option where the dislike objection has force. More precisely, one would have to hold both that (a) pain and dislike can come apart, and (b) the fact that pain hurts is independent of facts about its badness. However, those who endorse this view must face some radical and implausible consequences. Think about the most painful experience you've ever had (or imagine a torturous experience even more painful than any you've had) and imagine a phenomenal duplicate who's undergoing that exact same total experience but who doesn't dislike that experience.<sup>17</sup> It's very hard to believe that that experience isn't bad for that subject. The idea that pain is bad because

---

<sup>17</sup> I find it very hard to genuinely imagine such a scenario. If I imagine a subject undergoing the exact same total experience I have when I have a painful experience, then it seems to me that that subject thereby dislikes that experience. In other words, it's natural to think that the disliking of the experience is constitutively connected to how it feels to have the experience.

---

it's disliked is credible only if one retains the claim that pain is bad because of how it feels.

5. **Uncertainty:** One can know what it's like to have an experience yet be rationally uncertain about how good/bad that experience is.

If you know what it's like to have an experience, then you're in a position to know at least approximately how good/bad that experience is. It's hard to think of a counterexample to the inference from 'S knows what it's like to  $\phi$ ' to 'S is in a position to know approximately how good/bad the experience of  $\phi$ 'ing is'. Consider, as random examples, what it's like for you to drink coffee, go swimming, do pushups, or read this paper. I suspect you will easily know, to at least some degree of precision, how good/bad those experiences are. There may, of course, be cases where one is unsure whether an experience is overall good or overall bad. But even in those cases, one can usually still know that the experience is close to value neutral.

6. **Analyticity:** It's analytic that pain is bad. Therefore, the fact that pain is bad isn't a purely descriptive fact.

To say that a sentence is analytic is to say that the truth of the sentence follows from the definitions of its terms. There are a number of standard tests for evaluating whether a sentence is analytic: (1) denials of analytic sentences should sound nonsensical, (2) assertions of analytic sentences should sound cognitively insignificant, and (3) analytic sentences should be translatable into logically true sentences through the substitution of synonyms. And these tests all indicate that the sentence 'Pain is bad' isn't analytic. Consider: (1) those who deny that pain is bad aren't speaking nonsensically (instead, they're speaking falsely), (2) those who assert that pain is bad seem to be making a substantive (even if obvious) claim, and (3) there's no clear way of translating 'Pain is bad' into a logical truth through substitution of



---

synonyms. Furthermore, if it were analytically true that pain is bad, then ethical nihilists ought to deny that pains exist (since they deny that anything bad exists). But it's obvious that ethical nihilists can accept that there are pains, which is evidence that it's not built into the meaning of 'pain' that it's bad.

7. **A Prioricity:** It's a priori that pain is bad. Therefore, it's trivially true that the fact that pain is bad is explained by phenomenal facts, since that fact is explained by any fact whatsoever.

It's certainly not obvious that this fact is a priori. On the one hand, anyone competent with the concepts PAIN and BAD seems in a position to know that pain is bad. But on the other hand, competency with the concept PAIN seems to require a posteriori knowledge of how pain feels. How we think about this issue turns on delicate issues about what it is for experience to play a justifying role in acquiring knowledge of a fact. But let's suppose, for argument, that it's indeed a priori that pain is bad.

To know a fact a priori, one must possess the concepts needed to understand that fact. What kind of concept of pain one must possess in order to understand the fact that pain is bad? Suppose that you possess a purely functional concept of pain that yields no knowledge of how pain feels (but that still enables you to have thoughts that refer to pain). Even if it's a priori that pain is bad, a purely functional concept of pain would arguably be insufficient for acquiring a priori knowledge of that fact.<sup>18</sup> What more is needed? Well, it's plausible that in order for one to know a priori that pain is bad, one must grasp how pain feels. But this means that even if it's a priori that pain is bad, it remains the case that in order to know that pain is bad, one must know how pain feels. Consequently, metaethical experientialism still captures the epistemic structure of how we know pain is bad.

---

<sup>18</sup> See Kahane [2010] for a more sustained argument that anyone who doesn't know how pain feels isn't in a position to know that pain is bad.

- 
8. **Rationality:** For any descriptive fact P and ethical fact Q, it's possible for a rational agent to know P while denying Q.

For many ethical beliefs, there are philosophers who have defended the possibility ideally rational subjects with ethical beliefs that are inversions of our own ethical beliefs. Consider, for example, someone who believes that killing people for fun is permissible.<sup>19</sup> To my knowledge, however, nobody has defended the possibility of ideally rational subjects who believe that pain is good and pleasure is bad. Let's call these sorts of subjects *hedonic inverts*.

Hedonic inverts have the same kinds of experiences as ordinary humans (their pains hurt just as much as yours and mine) in the same kinds of circumstances (bodily damage causes them to feel pain). But they form the opposite value beliefs on the basis of those experiences. Both the ordinary subject and the hedonic invert know what it's like to burn one's hand on a hot skillet. But whereas the ordinary subject forms the belief that that kind of experience is bad (on the basis of their experience), the hedonic invert instead forms the belief that that experience is good (on the basis of their experience). Even those who think that there could be ideally rational ethical eccentrics still tend to think that there's something epistemically defective about hedonic inverts. This generates a puzzle: Why do hedonic inverts seem epistemically worse than other kinds of ethical eccentrics?

Metaethical experientialism provides an answer. A hedonic invert is in a position to know that their experience is bad (because they know how the experience feels), yet they instead form the opposite belief (that the experience is good). Note the peculiarity of the situation: the hedonic invert has knowledge (of how pain feels) that justifies a certain belief (the belief that

---

<sup>19</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's possible to accept both my claims about hedonic inverts and Street's claims about Future Tuesday Indifference.

---

pain is bad), yet on the basis of that very knowledge forms the opposite belief (that pain is good). This is epistemically worse than more mundane cases where a subject merely believes  $\neg P$  when they ought to believe that  $P$ , or where a subject believes  $P$  on the basis of  $Q$  when  $Q$  itself leaves open whether or not  $P$ . By contrast, supposing that phenomenal facts don't explain other kinds of ethical facts, other kinds of ethical eccentrics aren't manifesting the same kind of epistemic incoherence.

9. **Debunking:** We believe painful experiences are bad because it was evolutionarily advantageous, rather than because it's true.

To debunk a belief, one needs more than merely the premise that there is an evolutionary explanation for why we have the belief. If that were the only requirement, then nearly every belief would be debunkable. While there's dispute over what more is required for a debunking argument to succeed, it's generally accepted that it must at least be conceivable for the relevant facts to have been different.<sup>20</sup> The role of the conceivability premise is to show that in scenarios where the relevant facts were different, we would still have had the same beliefs that we actually have (because our beliefs track what is evolutionarily advantageous, rather than what is true).

If metaethical experientialism is true, then it's inconceivable that the phenomenal facts obtain without the value facts explained by those phenomenal facts obtaining. This follows from more general principles about epistemic gaps: if there's no epistemic gap from  $F$ -facts to  $G$ -facts, then it should be inconceivable for the  $F$ -facts to obtain without the corresponding  $G$ -facts obtaining. Compare: it's inconceivable for the microphysical facts to obtain without the corresponding macrophysical facts obtaining (at least assuming there's no epistemic gap from microphysical facts to macrophysical facts).

---

<sup>20</sup> See Street [2006] for a classic example of a debunking argument against ethical beliefs. See Kahane [2011], Clarke-Doane [2012], and Vavova [2015] for some recent discussions of evolutionary debunking arguments.

---

Hence, if you know what it's like to burn your hand on the skillet, then (according to metaethical experientialism) it's inconceivable for you to be having an experience with exactly that phenomenal character yet for that experience to not be bad. This means that there's no non-question begging debunking argument against metaethical experientialism. Instead, any debunking argument will have to appeal to premises that the metaethical experientialist already contests.

10. **Closure:** If  $P$  is explained by some set of phenomenal facts, then  $P$  is itself a phenomenal fact. Therefore, phenomenal facts cannot explain value facts.

Let's call the principle expressed above *phenomenal closure*. Recall that 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  $F$ -fact is a fact that ascribes only  $F$ -properties.<sup>21</sup> I'll argue that there are compelling reasons to think that phenomenal facts sometimes explain other kinds of facts.

Let a *structural fact* be a fact that ascribes only purely structural properties (such as the kinds of properties captured by mathematical models).<sup>22</sup> Structural facts don't ascribe any phenomenal properties, so they aren't phenomenal facts. But some structural facts are nevertheless explained by phenomenal facts. Consider, for example, how the similarity relations between color experiences can be represented via a geometrical space, where color experiences that are more similar to each other correspond to regions that are

---

<sup>21</sup> This categorization scheme still leaves some classes of facts epistemically isolated from others: for example, it remains plausible that phenomenal facts are epistemically isolated from physical facts and that physical facts are epistemically isolated from mathematical facts.

<sup>22</sup> I'll focus on structural facts because I think it's the most compelling case. But it's also worth noting that phenomenal conservativists, phenomenal intentionalists, and phenomenal powers theorists will likewise deny phenomenal closure, since each of those positions takes phenomenal facts to explain facts of another domain.

---

closer in the space.<sup>23</sup> If we consider a purely formal model of color experiences, then the model will specify only structural facts about those experiences.<sup>24</sup> But the structural facts specified by that model will themselves be explained by phenomenal facts about the target experiences. Given this, we have reason to reject phenomenal closure.

You might point out that the aforementioned structural facts are still facts *about* experiences. But that's irrelevant in the present context, since facts about the values of experiences are likewise facts about experiences. If we were to categorize facts by their referents (instead of by the properties of the they ascribe), then metaethical experientialism would no longer be in tension with the epistemic gap, since the value facts that figure into the conclusion have the same referents as the phenomenal facts that figure into the premises. To draw out the tension between metaethical experientialism and the epistemic gap, we need to categorize facts by the properties they ascribe.

11. **Explanation:** This paper hasn't yet explained *why* phenomenal facts explain value facts about experiences. To justify metaethical experientialism, we need an explanation.

There are many cases where one can be justified in believing that F-facts explain G-facts even when if one doesn't know why that explanatory relation holds. Consider again the common belief that microphysical facts explain macrophysical facts. There's disagreement over how to explain that explanation. The disagreements concern which kind of epistemic entailment is relevant (analytic vs. synthetic? a priori vs. a posteriori?), which metaphysical

---

<sup>23</sup> See Lee [2021] on the formal structure of quality-space models.

<sup>24</sup> You might object that these structural facts are themselves phenomenal facts. But note that the very same formal models could be used to represent non-experiential things. Since phenomenal facts are facts that ascribe phenomenal properties, and since phenomenal properties are instantiable only by experiences, the structural facts specified by formal models of experiences aren't phenomenal facts. See Lee [2022] for a more sustained version of this argument.

---

relations back the relevant epistemic relations (grounding? identity? realization? something else?), and why those metaphysical relations hold (essences? conceptual truths? something else?). These are live philosophical debates. But you don't need to know the answer to *why* microphysical facts explain macrophysical facts in order to be justified in believing *that* the explanations hold. By similar lights, you don't need to know *why* phenomenal facts explain value facts about experiences in order to be justified in believing *that* those explanations hold. To be justified in accepting an explanation, one doesn't necessarily need an explanation for that explanation.

### §3 Metaethical Consequences

How does metaethical experientialism relate to existing metaethical theories?

#### The Nature of Ethical Facts

To start, metaethical experientialism is incompatible with theories that deny the existence of any (non-trivial) ethical facts. This includes *non-cognitivism*, according to which ethical claims don't express truth-evaluable propositions, and *nihilism*, according to which ethical claims are systematically false. The incompatibility with these views is obvious: metaethical experientialism holds that some claims about the values of experiences are true.

A more subtle relationship concerns *subjectivism*, according to which ethical facts are true in virtue of the attitudes of valuers. You might think that metaethical experientialism entails that value facts about experiences are subjective, since they're explained by phenomenal facts and since phenomenal facts are themselves subjective. But we should distinguish 'subjective' in the phenomenal sense (where subjective facts are facts about subjective experiences) from 'subjective' in the alethic sense (where subjective facts are facts that are true in virtue of the attitudes of agents). While value facts about experiences are subjective in the phenomenal sense, they probably aren't

---

subjective in the alethic sense.<sup>25</sup> Since it's the alethic sense of 'subjective' that's standardly taken to demarcate realism from anti-realism, metaethical experientialism is naturally understood as a version of ethical realism.

Realist theories may be partitioned into *non-naturalist* theories, according to which ethical facts are non-natural facts, and *naturalist* theories, according to which ethical facts are natural facts. Metaethical experientialism is incompatible with non-naturalism. At least, this holds if we accept the following two claims: (1) phenomenal facts are natural facts, and (2) if value facts about experiences are explained by natural facts, then value facts about experiences are themselves natural facts. Both these claims strike me as plausible. Even if you're a non-physicalist, phenomenal facts arguably ought to count as natural facts in the present context (where we're contrasting descriptive facts from ethical facts, as opposed to, say, physical facts from phenomenal facts). And it's hard to see how there could be epistemic and metaphysical entailments from natural facts to value facts if those value facts are themselves non-natural facts. Although a non-naturalist version of metaethical experientialism may be coherent, it strikes me as unmotivated.

Given this, I'll henceforth understand metaethical experientialism as a version of ethical naturalism. Yet metaethical experientialism differs from the most prominent versions of ethical naturalism. One version is *analytic naturalism*, according to which ethical terms are synonymous with natural terms. Another version is *a posteriori naturalism*, according to which ethical properties can be a posteriori identified with natural properties. Both of these views are committed to a certain picture of the epistemology of ethical facts. But neither of those pictures aligns with the picture painted by metaethical experientialism. According to metaethical experientialism, our source of

---

<sup>25</sup> Well, you could have a version of metaethical experientialism where the phenomenal facts that explain value facts are themselves explained by attitudinal facts. How we classify this view depends on delicate issues about how to distinguish realism from anti-realism. And it may turn out that these views are best thought of as borderline between realism and anti-realism. For the purposes of this section, I'll set these views aside and assume that metaethical experientialism is committed to ethical realism.

---

knowledge for value facts about experiences is neither conceptual analysis nor scientific investigation. Instead, metaethical experientialism is most naturally thought of as a version of *a priori naturalism*, whereby some ethical facts are a priori derivable from natural facts.

A priori naturalism is an uncommon view in metaethics.<sup>26</sup> The principal reason may be because of the dominance of the epistemic gap: if it's indeed true that no ethical conclusion can be derived from purely descriptive premises, then a priori naturalism is a non-starter. But a core point of this paper has been to argue that the epistemic gap isn't compelling when we focus on value facts about experiences. Nevertheless, my goal has also been to develop a view *only* about value facts about experiences. This means that metaethical experientialism leaves open how to think about the nature of other kinds of ethical facts. In fact—as I'll discuss next—one could develop a pluralist metaethics that combines metaethical experientialism with some other metaethical theory.

### **Metaethical Pluralism**

Nearly all discussions in metaethics assume *metaethical monism*, the view that the same metaethical analysis will apply to all ethical facts. By contrast, metaethical experientialism is compatible with *metaethical pluralism*, the view that different metaethical analyses hold for different classes of ethical facts.

If you accept metaethical experientialism, then you accept that (at least some) value facts about experiences are explained by phenomenal facts. But this leaves open the possibility that there are other kinds of ethical facts (say, deontic or fittingness facts about non-experiential things) that aren't explained by phenomenal facts. In fact, metaethical experientialism is a very modest thesis, at least in terms of the class of facts that it targets. It's combinable with just about any theory about the nature of ethical facts and ethical knowledge, so long as that theory is restricted to other classes of ethical facts (rather than value facts about experiences).

---

<sup>26</sup> But see Sparks [2018] for a recent endorsement of this sort of view.



---

There's been little exploration of metaethical pluralism. All of the metaethical hedonists discussed earlier favor metaethical monism. The reason is that they are all moved by the central challenge in ethical epistemology: if there are ethical facts, then how can we acquire knowledge of them? By postulating a special source of knowledge for facts about pleasures and pains, metaethical hedonists have an answer for one class of ethical facts. But since this solution doesn't generalize to other ethical subject-matters, metaethical hedonists tend to assume nihilism about ethical propositions that can't be explained by phenomenal facts. And other metaethical theories likewise tend to take metaethical monism for granted: both in discussions of the metaphysics of ethical facts (nihilism, subjectivism, non-naturalism, naturalism, etc.) and the epistemology of ethical facts (intuitionism, analytic naturalism, a posteriori naturalism, etc.), it's standardly assumed that a single metaethical analysis will apply universally.

Perhaps metaethical monism is a reasonable default assumption: unless we have reason to think otherwise, we may as well aim for our theories to be universally applicable. But there's no obvious reason to think that monism *has* to be the case. If you think about the variety of facts that fall under the subject-matter of ethics—facts about value vs. reasons vs. fittingness, about welfare vs. morality, about intrinsic properties vs. extrinsic properties, about experiences vs. non-experiential things—then you might even find it reasonable to apportion more credence to pluralism than to monism. And if you find metaethical experientialism plausible yet also

A view that strikes me as especially attractive is objectivism about value facts about experiences and subjectivism about other kinds of ethical facts. A central challenge in metaethics is reconciling ethical realism with ethical knowledge. If there are objective ethical facts, then how is it possible for us to acquire knowledge of them? Metaethical experientialism, in my view, offers a solution: we can come to know facts about the values of experiences by knowing what it's like to undergo those experiences. But this solution works only for a limited class of ethical facts. Granted, it's the class of ethical

---

facts that arguably has the greatest claim for being objective: if any ethical facts are objective, then the fact that pain is bad is plausibly one of them. But what about the metaethical status of other kinds of ethical facts?

Suppose the epistemological challenge persists for other classes of ethical facts. That is, suppose no other version of ethical realism can be combined with a satisfying ethical epistemology. Then we are faced with the choice of either nihilism or subjectivism. In my view, subjectivism is the more attractive option. Instead of taking other ethical beliefs to be systematically false, we could instead take them to have a different metaethical ground. There are big questions, of course, about what kind of subjectivist theory to favor, and what the relevant class of subjective ethical facts would look like. I won't speculate on that here; I merely wanted to sketch the kind of metaethical pluralism that strikes me as most plausible.

If the epistemic gap remains for other kinds of ethical facts, then the metaethics of conscious experiences is special. Even if the epistemic gap isolates many classes of ethical facts from the domain of descriptive facts, there's a class of ethical facts—value facts about experiences—where the epistemic gap is surmountable. This inference may feel surprising when it's considered abstractly. But I think it's compelling and defensible upon close examination. The best answer to how we know pain is bad turns out—in my view—to also be 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.
- Bennett, Karen (2017). *Making Things Up*. Oxford University Press.
- Bliss, Ricki and Kelly Trogdon, "Metaphysical Grounding", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/grounding/>>.
- 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.
- Brink, David O. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press, 1989.
- Broi, Antonin (forthcoming). From the Heterogeneity Problem to a Natural-Kind Approach to Pleasure. *Pacific Philosophical Quarterly*.
- Chalmers, D. (2012). *Constructing the World*. New York: Oxford University Press.
- Clarke-Doane, Justin (2012). Morality and Mathematics: The Evolutionary Challenge. *Ethics* 122 (2):313-340.
- Delon, Nicolas (forthcoming). Strangers to ourselves: a Nietzschean challenge to the badness of suffering. *Inquiry: An Interdisciplinary Journal of Philosophy*.

---

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.

Goldstein, Irwin (1989). Pleasure and pain: Unconditional intrinsic values. *Philosophy and Phenomenological Research* 50 (December):255-276.

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.

Kagan, S. (1998). *Normative ethics*. Westview Press.

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.
- Kammerer, Francois (2019). The Normative Challenge for Illusionist Views of Consciousness. *Ergo: An Open Access Journal of Philosophy* 6.
- 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 (2021). Modeling Mental Qualities. *Philosophical Review*.
- Lin, Eden (2018). Attitudinal and Phenomenological Theories of Pleasure. *Philosophy and Phenomenological Research*.
- Lin, Eden (2020). The experience requirement on well-being. *Philosophical Studies*:1-20.
- 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.
- 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.

---

Parfit, Derek (2011). *On What Matters: Two-Volume Set*. 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, Bertrand (1918). The Philosophy of Logical Atomism, Lectures 1-2. *The Monist* 28 (4):495-527.

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, Geoffrey (1988). Moral Theory and Explanatory Impotence. *Midwest Studies in Philosophy* 12 (1):433-457.

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.
- Sinhababu, Neil (2018). Ethical Reductionism. *Journal of Ethics and Social Philosophy* 13 (1):32-52.
- Sinhababu, Neil (forthcoming). The reliable route from nonmoral evidence to moral conclusions. *Erkenntnis*:1-21.
- 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 der Deijl, Willem (2019). Is pleasure all that is good about experience? *Philosophical Studies* 176 (7):1-19.
- van der Deijl, Willem (2020). The sentience argument for experientialism about welfare. *Philosophical Studies*:1-22.
- 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.

Wilson, J., 2014, "No Work for a Theory of Grounding", *Inquiry*, 57(5-6): 1-45.