

The Metaethics of Mind

Andrew Y. Lee

Rice University, Dept. Of Philosophy

andrewyuanlee@gmail.com

<http://andrewyuanlee.com/>

ABSTRACT:

I develop and defend *value experientialism*, the thesis that value facts about experiences (such as the fact that pain is bad) are wholly explained by phenomenal facts (such as the fact that pain hurts). I argue that value experientialism aligns with common sense, that denying it leads to counterintuitive consequences, that it explains doxastic asymmetries between different classes of ethical beliefs, and that it accounts for the epistemic defects of hedonic inverts. I also argue that we ought to embrace the surprising result that there are genuine counterexamples to the explanatory gap between descriptive facts and ethical facts. An upshot is that the metaethics of mind warrants special investigation, for there are intimate connections between experience and value.

KEYWORDS: metaethics of mind, value experientialism, hedonic invert, moral epistemology, pain, well-being, fact-value gap, is-ought gap, Hume's Law, explanatory gap, value of consciousness, hedonism

INTRODUCTION

A central question in metaethics is whether it is possible to account for ethical knowledge without giving up on ethical realism. If ethical facts are independent of our evaluative attitudes, how do we acquire knowledge of those facts? At the heart of the challenge is the *explanatory gap*, or the widely held principle that no ethical conclusion is derivable from purely descriptive premises. As a result, the familiar methods by which we acquire knowledge of how things are do not suffice for acquiring knowledge of which things matter. Instead, it seems we must either come up with new epistemological innovations or seek anti-realist alternatives.

This paper argues that there are intimate connections between experience and value, and that these connections yield a solution to the aforementioned challenge for a special subclass of ethical facts. Here is my core thesis:

VALUE EXPERIENTIALISM:

Value facts about experiences are explained 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 that ascribe only phenomenal properties. And by *explain*, I mean a relation that consists of an epistemic component and a metaphysical component. The epistemic component is that some value facts are *a priori entailed* by phenomenal facts. The metaphysical component is that some value facts are fully *grounded* in phenomenal facts.

To get a feel for value experientialism, consider the badness of pain. It is natural to think that pain is bad because of how it feels. It is plausible that one can know that pain is bad simply by knowing how pain feels. It is counterintuitive to think that a subject could know how pain feels yet not be in a position to know that pain is bad. And if someone knows how pain feels but instead concludes on the basis of such knowledge that pain is good, then they seem to be making an epistemic mistake. This previews some of the motivations for value experientialism, which I will develop in detail later.

If value experientialism is true, then some ethical facts are natural (i.e., fully grounded in natural facts), objective (i.e., true independent of our value beliefs), and

immune to the explanatory gap (i.e., derivable from purely descriptive facts). As I will explain, these commitments distinguish value experientialism from nearly all standard metaethical theories. Yet while I will argue that phenomenal facts explain value facts about experiences, I will also presume that phenomenal facts do not explain other kinds of ethical facts (such as moral facts about what one ought to do or value facts about non-experiential things). If that is right, then the metaethics of mind is special: there is a source of knowledge for value facts about experiences that does not yield knowledge of other kinds of ethical facts.

The paper has two main sections: in §1, I explain and motivate value experientialism, and in §2, I defend value experientialism from objections.

§ 1 | VALUE EXPERIENTIALISM

The focal example of this paper concerns 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. Speaking somewhat loosely, I take ‘pain’ to denote a kind of experience that by definition involves suffering.¹ 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.² 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

¹ In ordinary language, we often predicate ‘hurting’ to bodily events (such as a pinprick) or 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.

² 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].

bad, we could substitute in another kind of experience that one takes to be bad (e.g., torturous pain, nausea, or disgust experience) or good (e.g., pleasure, beauty, or gustatory experience).

This paper is about the relationship between certain value facts and phenomenal facts. By *value fact*, I mean a fact that ascribes only value properties; by *phenomenal fact*, I mean a fact that ascribes only phenomenal properties.³ For example, the fact that 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). I will take facts to be categorized by meaning, rather than by extension: even if the fact that pain hurts entails the fact that pain is bad, that does not make the fact that pain hurts a value fact. This is an initial gloss—I will explain these points in more detail in the next section.

What do I mean by *value*? I will use the term ‘value’ as neutral between denoting goodness or badness (so badness is a kind of value). My concern is only with value that is *prudential* (e.g., badness that makes one’s life worse) as opposed to moral, epistemic, or aesthetic, *final* (e.g., badness for its own sake) as opposed to instrumental, and *pro-tanto* (e.g., badness that can be outweighed by other factors) as opposed to all-things-considered.⁴ I will also take for granted *cognitivism*, according to which value claims express propositions.

I take value facts about experiences to be a subclass of ethical facts. The set of *ethical facts* consists not only of facts about the values of experiences but also of facts about the values of other kinds of entities, and not only of facts about what is

³ I assume the properties are non-trivial—for example, the property of being either good or not good does not count as a value property. And although I focus on universal facts (where badness is attributed to properties, such as pain), I take my arguments to generalize to particular facts as well (where badness is attributed to particulars, such as a particular pain).

⁴ Though my focus is on prudential value, I believe the arguments generalize also to neutral value (i.e., value that makes a world better or worse). See Nagel [1986, p. 156–163] for discussion of the generalization from prudential value to neutral value.

good or bad for individuals but also of facts about what is morally right or wrong.⁵ Though value experientialism is strictly speaking neutral on whether other kinds of ethical facts are explained by phenomenal facts, I will presume that the answer is 'no'. More specifically, I will take for granted that value facts about non-experiential things and moral facts of any kind are not explained by phenomenal 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.

While I will talk about value facts about experiences without qualification, it is best to understand value experientialism as a restricted universal claim. To see why, consider a pluralist about value who agrees that some value facts about experiences are explained by phenomenal facts but who also thinks that other value facts about experiences are not 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 basically all the arguments in this paper support these kinds of pluralist positions (along with monist positions), it is best to include such views within the scope of value experientialism. In light of this, value experientialism 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 just using the term 'value facts about experience' (with the quantifier restriction implicit).

Though I frame my discussion in terms of value, some might prefer talking in terms of reasons. For example, where I talk about the fact that pain is bad, some might prefer to instead talk about the fact that subjects have reasons to avoid pain. Then we could address *reason experientialism*, according to which reason facts about

⁵ Though I presume here that there are other kinds of ethical facts, value experientialism is consistent with taking value facts about experiences to be the only basic ethical facts.

experiences (such as the fact that one has reason to avoid pain) are explained by phenomenal facts (such as the fact that pain hurts). Though I am sympathetic to reason experientialism, I focus only on value in this paper, and I take value experientialism to be officially silent on the relationship between value and reasons.

WHAT IS VALUE EXPERIENTIALISM?

As mentioned, value experientialism consists of an epistemic thesis and a metaphysical thesis. The epistemic thesis is that value facts about experiences are a priori entailed by phenomenal facts; the metaphysical thesis is that value facts about experiences are fully grounded in phenomenal facts.⁶

Throughout the paper, I will talk about the *explanatory gap*, or the principle that no ethical conclusion is derivable from purely descriptive premises.⁷ More precisely, I will interpret the explanatory gap as the thesis that no ethical fact is a priori entailed by any set of purely descriptive facts.⁸ Taken at face value, the explanatory gap strikes many as extremely plausible: how could facts about what matters be derived merely from facts about what is? In light of this, the most radical

⁶ Though I frame the metaphysical thesis in terms of grounding, it could also be framed in terms of other metaphysical relations (such as realization or constitution). The main theoretical role I attribute to grounding is that it is an asymmetrical metaphysical dependence relation that captures the sense in which some facts are nothing over and above other facts.

⁷ The term 'explanatory gap' is arguably better than more frequently used alternatives. The alternative terms include 'is-ought gap' (but my concern is primarily with value rather than oughts), '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 reductive analysis rather a priori entailment).

⁸ This is a common way of understanding explanatory gaps (especially in the philosophy of mind: see, e.g., Chalmers & Jackson [2001]), though philosophers sometimes characterize explanatory gaps in terms of concepts (e.g., Mehta [2019], in a recent discussion of explanatory gaps across multiple domains, appeals to conceptual families).

consequence of value experientialism is that it denies the explanatory gap, at least for value facts about experiences.⁹

Though I will argue for both the epistemic thesis and the metaphysical thesis, my main focus will be on the epistemic thesis. The epistemic thesis is more controversial: while any ethical naturalist will accept the metaphysical thesis, few philosophers nowadays deny the explanatory gap. In fact, a core argument against the metaphysical thesis appeals to the epistemic gap: if ethical facts are grounded in natural facts then there should be no epistemic gap, but there is an epistemic gap, so ethical facts are not grounded in natural facts. The argument is compelling because it is hard to see how one could endorse the metaphysical thesis without also endorsing the epistemic thesis: if value facts about experience are something over and above phenomenal facts about experience, then how could those phenomenal facts a priori entail those value facts?

If P a priori entails Q, then Q is knowable solely on the basis of P. But I will not argue for the epistemic thesis by arguing that value facts about experience are a priori simpliciter. Instead, my view is that phenomenal facts provide epistemic support for value facts about experiences, in the sense that value facts about experience are a priori knowable *on the basis of* phenomenal facts. Consider, as an analogy, how macrophysical facts are a priori knowable on the basis of microphysical facts but not a priori simpliciter. In general, I will presume that if P a priori entails Q, then P provides justification for Q, P is evidence for Q, it is inconceivable that P & ¬Q, and it is ideally rationally impermissible to believe the conjunction P & Q.¹⁰

⁹ There are well-known counterexamples to simple formulations of the explanatory 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], Fine [forthcoming], and Sparks [forthcoming]. For broader discussion of normative explanation, see Väyrynen [2013].

¹⁰ In Section 2 (under THE SELF-EVIDENCE OBJECTION), I address the objection that value facts about experience are in fact a priori simpliciter.

Note that a priori entailment is a different relation than reductive analysis. Whereas a priori entailment is a relation between facts, reductive analysis is a relation between properties. To say that a phenomenal fact a priori entails a value fact is to say that one can know that value fact solely on the basis of that phenomenal fact. In contrast, 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 endorse the a priori entailment claim, I am skeptical of reductive analysis claims: 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 a priori entail macrophysical facts even though there are no reductive analyses of macrophysical properties in terms of microphysical properties.¹¹

Value experientialism remains neutral on a number of nearby issues. First, the thesis is neutral on whether all value facts (or only a privileged subset) are explained by phenomenal facts. Second, value experientialism 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 normative ethics).

By focusing on the metaethics of the mind, this paper adopts a *bottom-up* methodology, whereby we examine a particular domain and identify metaethical principles that apply to that domain. In contrast, the vast majority of contemporary metaethics adopts a *top-down* methodology, whereby we search for principles and

¹¹ See Fodor [1974] for a classic argument against reductive analyses of macroproperties. See Suikkanen [2016] and Ridge [2019] for discussions of reductive analyses in metaethics (in the context of the open question argument against ethical naturalism).

theories that aim to apply globally to all domains. Though I think the top-down methodology is important for metaethical inquiry, adopting an exclusively top-down perspective risks overlooking the details that distinguish one domain from another. Principles such as the explanatory gap may seem attractive from a top-down perspective, but are less obvious when we consider cases such as the badness of pain.

MOTIVATIONS

Before articulating the motivations for value experientialism, let me make a preliminary remark about this paper's dialectical ambitions. My primary aim is to show that value experientialism is dialectically defensible, rather than to persuade those antecedently skeptical of the thesis. At a certain point, philosophical disagreements reach bedrock, and that point is reached relatively quickly when evaluating value experientialism. Because of this, I grant that some of the motivations I appeal to might not move all philosophers. Nevertheless, I suspect a significant proportion of philosophers will share my intuitions and be attracted to value experientialism.

A first motivation for value experientialism is that it aligns with common sense. If you are asked how you know pain is bad (or why 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 why it is bad. And perhaps you might mention 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 pain is intrinsically bad, there seems nothing more relevant than the fact that pain hurts. In fact, if someone is unconvinced by your answer, then the most convincing method of persuasion may be to remind them how pain feels.¹²

¹² I do not condone violence.

A second motivation is that denying value experientialism leads to counter-intuitive consequences. To deny value experientialism, one must hold either that a subject could know that pain hurts yet not be in a position to know that pain is bad (if one denies the epistemic thesis) or that one could be in pain yet for that pain to not be bad (if one denies the metaphysical thesis). But it is implausible that either situation is possible, at least if we stipulate that the subject possesses the relevant concepts and the right kinds of reasoning capacities and that the pain feels the way that ordinary pains do. As before, if you are insufficiently moved by the case of pain, then we can substitute in a different kind of experience, such as torturous pain, extreme nausea, or horrific disgust. Even if all one knows about those experiences is what it is like to have them, it seems that is already enough to know that those experiences are bad. What more could be required?

There may be temptation to claim that knowledge of certain phenomenal facts necessarily entails whatever factor enables knowledge of 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 yet argue that this is because necessarily anyone who knows that pain hurts has the intuition that pain is bad. However, it is incumbent on such a theorist to explain why there should be a necessary connection between knowledge of how pain feels and the intuition that pain is bad. Since it is preferable to avoid positing necessary connections between distinct existents, it is hard to see what could justify this modal claim. Perhaps it is possible to argue that knowledge of how pain feels grounds the intuition that pain is bad, but such a move feels somewhat desperate: at least, we would need independent motivation for thinking that such a grounding relationship holds.

A third motivation concerns doxastic asymmetries. Though there is little consensus amongst philosophers on which ethical propositions are true, it is nearly

universally believed that pain is bad.¹³ There are exceptions, of course, but the fact that there are outliers does not cast doubt on the genuine pattern. Even philosophers moved by arguments for global skepticism about ethical beliefs tend to be reluctant to give up the belief that pain is bad. And while objective list theorists, desire-satisfactionists, and hedonists have many disagreements about what kinds of things are valuable, there tends to be consensus that pain is bad (and pleasure is good). Though 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, meaning we have better epistemic grounds for believing pain is bad than for believing other kinds of ethical facts.

A fourth motivation concerns hedonic inversion. A *hedonic invert* is a subject that has the same kinds of hedonic experiences as ordinary humans (in particular, their pains hurt just as much as yours and mine), but who forms the opposite value beliefs on the basis of their hedonic experiences. In other words, whereas any ordinary subject would form the belief that pain is bad on the basis of how pain feels, the hedonic invert forms the belief that pain is good on that very same basis. There is a strong intuition that the hedonic invert is not merely mistaken in their beliefs, but that they are also epistemically defective in some deeper way. Value experientialism provides a diagnosis of how the hedonic invert is epistemically defective: in particular, the hedonic invert forms beliefs that run counter to their evidence. Though the hedonic invert has evidence that pain is bad (because they know how pain feels), they use that evidence to form the belief opposite from what that evidence supports.

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

¹³ As before, we can substitute in another experience, such as torturous pain, if one is concerned with phenomena such as pain asymbolia.

There is at least some reason to think that the answer is ‘no’. Some 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).¹⁴ To the best of my knowledge, though, 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. If that is right, then value experientialism provides a natural explanation: even though phenomenal facts explain value facts about experiences, they do not explain other kinds of ethical facts. Consequently, hedonic inverts have reason to deny that pain is good since they know how pain feels, but other kinds of ethical eccentrics do not have analogous reasons for denying the unusual ethical beliefs that they have formed. For purposes of space, I will not argue directly for this asymmetry, but it is worth mentioning for those who antecedently find it plausible.¹⁵

¹⁴ For extended discussion of ideally coherent ethical eccentrics, see Street [2009]. It is worth noting that one of the cases Street discusses is Future Tuesday Indifference, which concerns beliefs about pain. However, that case is structurally different from hedonic belief inversion, and it is possible to accept both my conclusions about hedonic belief inversion and Street’s conclusions about the possible rationality of Future Tuesday Indifference. On the other hand, Street [2006] advances considerations that are in tension with my claims about hedonic inversion. There is not enough room in this paper to discuss those arguments in detail, but my responses to her arguments are contained in my responses to objections in the next section: most notably, under THE DISLIKE THEORY OBJECTION and THE DEBUNKING OBJECTION.

¹⁵ The epistemic asymmetry could also be framed in terms of other epistemic goods. For example, I am inclined to think that we have more evidence that pain is bad than we do for other kinds of ethical beliefs; that we have greater justification for believing that pain is bad than we do for other kinds of ethical beliefs; and that it is inconceivable that pain is not bad while other kinds of ethical falsehoods are generally conceivable.

A final point is that value experientialism is elegant. The picture I have sketched posits no non-natural ethical facts, no epistemically opaque metaphysical relations, and no mysterious knowledge-acquisition processes, all while maintaining a background of ethical realism. In light of this, I think the key question is not whether value experientialism is attractive, but instead whether the thesis is ultimately defensible. The aim of the paper so far has been to explain and motivate the value experientialism, but most of the rest of the paper will focus on defending the thesis from objections.

To recap, I have argued that value experientialism aligns with common sense, that denying it leads to counterintuitive consequences, that it explains certain doxastic asymmetries with respect to ethical beliefs, and that it explains the epistemic defects of hedonic inverts.

CONNECTIONS TO PHILOSOPHICAL LITERATURE

Before moving to objections, let me briefly situate value experientialism within the contemporary philosophical landscape. As far as I can tell, there has been no explicit defense of value experientialism, at least in the form presented here. However, there are several clusters of philosophical literatures that have significant connections to value experientialism, and there are a few philosophers who have defended nearby positions.

A first connection is to the literature concerning the nature of ethical facts. On this issue, value experientialism departs from the most popular contemporary theories. In contrast to anti-realists, value experientialism holds that there are objective ethical facts. In contrast to non-naturalists, value experientialism holds that value facts about experience are explained by natural facts. In contrast to non-reductive naturalists, value experientialism holds that value facts about experience are a priori entailed by natural facts. And in contrast to analytic naturalists, value experientialism denies that value facts about experiences are analytic (a point I discuss

later). Instead, value experientialism is a form of *a priori naturalism*, according to which some ethical facts are a priori derivable from natural facts.¹⁶

A second connection is to the literature concerning ethical knowledge. It is common for theories to compare ethical knowledge to other kinds 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 attitudinal knowledge (where we acquire ethical knowledge by understanding the contents of our evaluative attitudes).¹⁷ In contrast, my view is that we acquire knowledge of the values of our experiences in the same way that we acquire knowledge of the phenomenal characters of our experiences. Hence, I endorse a *phenomenal model* of ethical knowledge (at least for value facts about experiences).¹⁸

¹⁶ For examples of anti-realism, see Mackie [1977] and Street [2010]. For examples of non-naturalism, see Moore [1903], Shafer-Landau [2003], Enoch [2011], and Parfit [2011]. For examples of non-reductive naturalism (or ‘Cornell realism’), see Sturgeon [1985], Brink [1986], Railton [1986], and Boyd [1988]. For examples of analytic naturalism, see Jackson [1998] and Finlay [2014]. On some precisifications of terms, value experientialism may count as a form of either non-reductive naturalism or analytic naturalism. But since the thesis differs from canonical versions of the views that go under those labels and since ‘a priori physicalism’ in the philosophy of mind is the view that physical facts a priori entail phenomenal facts, I believe the label ‘a priori naturalism’ is more apt.

¹⁷ 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. This lacuna is a reflection on how little positions like value experientialism have been discussed.

¹⁸ This is a somewhat delicate point. To be precise, I am arguing here that value facts about experience can be known through *inference* (from phenomenal facts). However, the more

A third connection—which moves beyond metaethics—are several literatures that relate consciousness and value. This includes the literature on hedonism and utilitarianism, some parts of the pleasure and pain literature, and a recently emerging literature on the value of consciousness itself.¹⁹ However, nearly all of the relevant work in these clusters falls within normative ethics, whereas the focus of this paper is on metaethics. Though philosophers working in these areas sometimes make claims that are suggestive of value experientialism, there has been little work actually developing the metaethical foundations behind such suggestions. In light of this, value experientialism is independent of (but complements) work in these literatures.

There are a few philosophers who have argued for positions that are closely related to (though still distinct from) value experientialism. In particular, some have argued that phenomenal facts analytically entail value facts about experiences: for example, that the concept BAD is built into the concept PAIN.²⁰ In contrast, I believe we can always conceptually distinguish descriptive properties from ethical properties, and I will argue that the right way to capture the relation between the phenomenal facts and value facts about experiences is in terms of a priori entailment rather than in terms of analyticity. Though this distinction may at first blush seem subtle,

general epistemological 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 reveals why I choose to frame value experientialism as supporting a phenomenal model of ethical knowledge.

¹⁹ See Bramble [2016] and Deijl [forthcoming] for defenses of hedonism and non-hedonistic experientialism, respectively. See Kahane [2009, 2010] for discussions of the nature, value, and epistemology of pain that express sympathies to value experientialism. See Siewert [1998] for argument that conscious experiences are intrinsically valuable, and Lee [2018] and Kriegel [2019] for discussions of whether consciousness itself is intrinsically valuable.

²⁰ See Von Wright [1963], Mendola [1990], Hewitt [2008], and Massin [forthcoming].

it has important ramifications: most notably, only my view entails that there are genuine counterexamples to the explanatory gap.

Finally, there are structural parallels between value experientialism and other theories connecting phenomenal facts to facts of other domains. In particular, phenomenal conservatives hold that certain epistemic facts are explained by phenomenal facts, phenomenal intentionalists hold that certain intentional facts are explained by phenomenal facts, and phenomenal powers theorists hold that certain causal powers facts are explained by phenomenal facts.²¹ Analogously, value experientialism holds that some 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 getting a handle on value experientialism.

§ 2 | OBJECTIONS

In what follows, I will defend value experientialism from a number of objections, which will also serve to further clarify value experientialism itself.

THE EASY KNOWLEDGE OBJECTION

The easy knowledge objection is that value experientialism cannot explain 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, value experientialism entails that everyone is in a position to easily know that pain is bad. Yet some philosophers nevertheless deny that pain is bad.

In response: philosophers believe all sorts of crazy things. 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 knowledge of the external world exists, how folk-

²¹ For phenomenal conservatism, see Huemer [2001]. For phenomenal intentionalism, see Kriegel [2013]. For the phenomenal powers view, see Mørch [2018].

psychology eliminativists deny that beliefs and desires exist, how nominalists deny that properties exist, and how ontic structuralists deny that things exist. I am not suggesting that any of these views is false (except illusionism), but I know of nobody that thinks all of these 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.²²

Philosophers who endorse these radical 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 plausible that theoretical considerations do sometimes outweigh common sense. But it is also plausible that in some cases, one might be mistaken about how strong the countervailing theoretical considerations are and as a result deny claims that one is 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 conclusion that consciousness does not exist, so too there are clever arguments to the conclusion that pain is not bad.

The easy knowledge objection is also undercut by distinguishing ideal rationality and prima facie rationality. On the one hand, value experientialism 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 actual philosopher has an ideally rational set of philosophical beliefs. On the other hand, value experientialism allows that philosophers arguing for these counterintuitive positions still satisfy the standards for prima facie rationality that are more pertinent to ordinary human life. In light of this, we can reconcile the fact that everyone is in a

²² 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. For example, many philosophers think that we are in a position to easily know that there are external objects even though it is difficult to justify the claim that we have such knowledge.

position to easily know that pain is bad with the fact that philosophers who deny that pain is bad are not being irrational by ordinary standards.

THE SELF-EVIDENCE OBJECTION

The self-evidence objection is that it is self-evident that pain is bad. Consequently, it is not that the fact that pain hurts a priori entails the fact that pain is bad. Instead, the fact that pain is bad is knowable absent any premises whatsoever.

What does it mean for a proposition to be self-evident? At first pass, self-evident propositions are propositions such that anyone who understands them is in a position to know they are true.²³ There are two relevant precisifications of this idea that 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.²⁴ Or putting it in terms of concepts, this objection claims that the concept BAD is contained in the concept PAIN. However, it is easy to see that ‘pain is bad’ is not analytic when we apply some standard tests for analyticity. First, the denial of an analytic fact should seem non-sensical, 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 facts can be translated into logical facts through the substitution of synonyms, but there seems no way to translate ‘pain is bad’ into a logical fact through such a method. Finally,

²³ See DePaul & Hicks [2016] for more discussion of self-evident propositions in metaethics.

²⁴ Note that 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].

ethical nihilists are competent users of the term 'pain' but do not deny that pains exist. These observations are strong reason to deny that 'pain is bad' is analytic.

The second version of the 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.²⁵ 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 provides no grasp of how pain feels (but which still enables one to think thoughts that refer to pain). Even if it is a priori that pain is bad, it is plausible that the purely functional concept of pain does not enable one to acquire that a priori knowledge. What more is needed in order to acquire such knowledge? It is plausible that what is required is a concept that provides 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, value experientialism 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). But the conditional premise contains an ethical term (namely, 'bad'), so we do not have a genuine counterexample to the

²⁵ In fact, it is not obvious that it is a priori that pain is bad. In order to know a truth a priori one must have the relevant competency with the relevant concepts, and in order to have the relevant competency with the concept PAIN, one must know how pain feels. But is it possible to know a priori how pain feels? My own view is that it is not—as a consequence, I believe that knowledge that pain is bad is always (ultimately) justified a posteriori. However, discussing these issues in depth would take us too far astray.

explanatory gap. More generally, one might observe that any unconditional ethical conclusion will require a conditional premise containing an ethical term.

Suppose (per reductio) that an explanatory gap occurs between F-facts and G-facts whenever no set of purely F-premises logically entails a G-conclusion. Then explanatory 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 explanatory gap. Since explanatory gaps are not ubiquitous, the supposition that lack of logical entailment is sufficient for an explanatory gap should be rejected.

What about a version of the validity objection that appeals to the disjunction of logical or analytic validity (rather than only logical validity) as the criterion for bridging an explanatory gap? The problem with such a view 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 relevant conditionals: for example, we can determine whether an inference from Fx to Gx is analytic by determining whether the fact $Fx \supset Gx$ is analytic. When we do so, it is easy to find cases of a priori 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 a priori entails Q but implausible that P analytically entails Q .²⁶

The general lesson is that explanatory gaps ought to be characterized in epistemic terms, rather than logical or semantic terms. Because of this, a priori

²⁶ This conclusion aligns with standard understandings of a priority, analyticity, and logical validity. For further discussion, See De Paul & Hicks [2016], Russell [2017], and Rey [2018].

entailment (rather than rather logical or analytic validity) is the relevant relation when investigating explanatory gaps.

THE CLOSURE OBJECTION

The closure objection is that phenomenal facts can explain only other phenomenal facts. In other words, one might endorse a closure principle on phenomenal explanation. To support the objection, one might appeal to the common idea that there are different domains of facts—phenomenal, physical, abstract, normative, and so forth—that are conceptually isolated from one another, and contend that explanation cannot transmit across conceptually isolated domains.²⁷

To assess the closure objection, we need to precisify the criteria for categorizing facts. For the purpose of assessing explanatory closure principles, it is most perspicuous to type facts via the kinds of properties that are ascribed. On this categorization scheme, 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 way of individuating facts preserves the idea that some domains are conceptually isolated from each other: for example, it remains plausible that physical facts do not explain phenomenal facts. Nevertheless, I will argue that there are good reasons for thinking that phenomenal facts can explain facts of other kinds.

I have already mentioned some theories that take phenomenal facts to explain facts of other kinds. In particular, phenomenal conservatives hold that phenomenal facts can explain epistemic facts, phenomenal intentionalists hold that phenomenal facts can explain intentional facts, and phenomenal powers theorists hold that phenomenal facts can explain causal power facts. In order for the closure objection to work, all of these views must be false. Nevertheless, each of these positions is controversial, and I do not wish to lean on them to respond to the objection.

²⁷ See Mehta [2019] for a recent discussion.

Consider instead *structural facts*, or facts ascribing only structural properties (such as parthood, quantity, or dimensional properties). As a heuristic, structural facts are the kinds of facts directly captured by formal representations (such as mathematical models). Since structural facts do not ascribe phenomenal properties, they are not phenomenal facts. But it is plausible that structural facts about experiences are explained by phenomenal facts. Consider, as an example, 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 nearer in the space.²⁸ This formal model captures various structural facts about color experiences, such as facts about their similarity structure. Yet these structural facts are explained by phenomenal facts: it is in virtue of color experiences having their phenomenal characters that they stand in the structural relations captured by the model.

In response, the objector might argue that these structural facts are just further phenomenal facts. But while structural facts about experiences are grounded in phenomenal facts, they are not themselves phenomenal facts. To see why, consider how we can ask whether formal models of color experiences also capture the structural relations amongst physical colors, or 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 at least conceivable that structural properties instantiated by experiences are also instantiated by other kinds of things. But that means those structural properties are not phenomenal properties (since, by definition, phenomenal properties can be instantiated only by experiences). On a side note, the same line of reasoning shows that why taking value facts about experiences to be explained by phenomenal facts does not mean that those value facts are themselves

²⁸ See Clark [2000] and Rosenthal [2015] for discussion.

phenomenal facts (since it is at least conceivable that the same value properties are instantiated by non-experiential things).

At this point, the objector might point out that the relevant structural facts are still facts about experiences. But that is largely irrelevant, since the fact that pain is bad is also a fact about experiences. The relevant way of categorizing facts is by appeal to the properties they ascribe (rather than by their referents). If we were to instead categorize facts via their referents, then the fact that pain is bad would then count as a phenomenal fact (since it is a fact about experiences), meaning that the closure objection could not even get off the ground. In other words, the closure objection appeals to the idea that facts ascribing only phenomenal properties cannot explain facts ascribing other kinds of properties. But that hypothesis is false, so the support for the closure objection is undercut.²⁹

THE DISLIKE THEORY OBJECTION

The dislike theory objection is that pain is bad because 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 explained in terms of dislike).³⁰

²⁹ What if the objector concedes that phenomenal facts can explain facts of other kinds but 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 tantamount to simply denying value experientialism, so independent support is needed in order for the objection to not be question-begging. I have argued against the most natural way of supporting the objection, so the dialectical burden is on the objector to seek alternative support. Furthermore, I suspect that many will find the intuition motivating the objection less compelling when considering value facts about experiences.

³⁰ See Street [2006] for an example of a dislike theory objection.

Let us start with the objection that pain hurts because it is disliked. Suppose that what it is for an experience to feel unpleasant is for it to be disliked. That is consistent with thinking that the fact that pain hurts explains the fact that pain is bad. This version of the dislike theory provides a reductive analysis of what it is for an experience to feel unpleasant, but it leaves open what is explained by phenomenal facts about unpleasant experiences. In light of this, even those who contend that pain is bad because it is disliked can accept that pain is bad because it hurts, so long as they think that pain hurts because it is disliked.

There may be a temptation here to contend that there is no phenomenology of unpleasantness. After all, there is a debate in the pleasure and pain literature on whether unpleasantness is a phenomenal property or an attitudinal property. However, whether all unpleasant experiences have some phenomenal property in common is irrelevant to value experientialism. Suppose there are many different ways for an experience to feel unpleasant and that what makes any given experience count as unpleasant is the disliking of that experience rather than the instantiation of a common phenomenal property. Nevertheless, one could still think that value facts about those experiences are explained by phenomenal facts about those experiences.³¹

The other version of the dislike theory objection claims that pain is bad because it is disliked (and denies both that hurting is a matter of being disliked and that pain is bad because it hurts). Unlike the previous version, this version of the dislike theory is inconsistent with value experientialism. 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.³² On top of that, taking dislike to be independent of hurting commits the dislike theory to some counterintuitive

³¹ See Heathwood [2007], Bramble [2013], and Lin [2018] for discussions of the relation between pleasure and desire. See Kahane [2009] for discussion concerning pain.

³² See Bain [2017] for some discussion of this datum in the context of the pain literature.

consequences: as discussed in §1, the dislike theorist will be forced to hold that it is possible for a subject to be in pain yet to not be in a bad state and 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 undercuts the motivation for value experientialism.³³

Now, in order for a debunking argument to succeed, one needs more than merely the premise that there is an evolutionary explanation for why we have certain beliefs. If that were the only requirement, then nearly any set of beliefs 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.³⁴

³³ See Street [2006] for a classic example of a debunking argument against ethical beliefs.

³⁴ 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. But as far as I can tell, Clarke-Doane's argument equivocates between two 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 $\neg P$ is rulable out a priori). After all, there are plenty of philosophical debates involving peer disagreement over truths that are presumably knowable a priori. For further elaboration of these points, see my responses under THE EASY KNOWLEDGE OBJECTION and THE SELF-EVIDENCE OBJECTION.

For example, debunking arguments against our moral beliefs presume that it is conceivable for the moral facts to have been different from whatever they actually are. 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).

Value experientialism challenges this background premise of the debunking objection. If value facts about experience are a priori entailed by phenomenal facts, then it is inconceivable for the phenomenal facts to hold without the relevant value facts also holding. As discussed previously, claiming that it is inconceivable that pain is not bad (given knowledge of how pain feels) does not mean that those who deny that pain is bad are making nonsensical claims or that they are irrational by ordinary standards. Instead, it simply means that the fact that pain hurts a priori entails that pain is bad. In response, the debunker could deny that this a priori entailment holds. But this simply moves the dialectic back to the question of whether value experientialism is plausible in the first place. The upshot is that there is no debunking argument against value experientialism that does not presuppose its falsity.

THE GENERALIZATION OBJECTION

The generalization objection is that value experientialism fails to properly generalize beyond the core example of the badness of pain. We have all kinds of phenomenal knowledge about our experiences, but we are often unsure of what to think about the values of those experiences.

To begin, observe that in the vast majority of cases we know approximately how good or bad an experience is. For example, consider what it was like for you to last exercise, or what it was like for you to eat the last meal you ate, or what it was like for you to last watch a movie or listen to music or play a game. I suspect you will find it easy to know approximately how good or bad those experiences were for you. In fact, it is hard to think of an experience 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 the experience is. Even when an experience is close to value neutral, we

can often know that it is close to value neutral (even if we cannot know whether it has net positive value or net negative value).

What about exact knowledge of the values of our experiences? Though it is plausible that we rarely acquire exact knowledge of the values of our experiences, it is also plausible that we rarely acquire exact knowledge of the phenomenal character of our experiences. There is controversy over which principles and theories best capture the epistemology of experience, but just about all philosophers of mind nowadays accept that our phenomenal judgments are less reliable when targeting aspects of experiences that are detailed, subtle, or complex. 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 experience is continuous or discrete.³⁵ Since we normally have only approximate knowledge of the phenomenal characters of our experiences, it is unsurprising that we normally also have only approximate knowledge of the values of our experiences.

CONCLUSION

I have explained, motivated, and defended value experientialism. The motivations were that value experientialism aligns with common sense, that denying it leads to counterintuitive consequences, that it explains doxastic asymmetries between different kinds of ethical beliefs, that it diagnoses the epistemic defects of hedonic inverts, and that it provides an elegant metaethical picture for value facts about experience. The objections addressed were the easy knowledge objection, the self-evidence objection, the validity objection, the closure objection, the dislike theory objection, the debunking objection, and the generalization objection.

³⁵ See Block [1995], Schwitzgebel [2006], Williamson [2000], and Lee [2019] for different kinds of arguments that our knowledge of our own experiences is limited.

I began the paper with a central question in metaethics: is it possible to account for ethical knowledge without giving up on ethical realism? Though I have said little to answer this question for many kinds of ethical facts, I have argued that there is a special class of ethical facts—namely, value facts about experience—where the explanatory gap is surmountable. Value facts about experience are explained by phenomenal facts, so we can acquire ethical knowledge on the basis of our phenomenal knowledge. Though the inference is 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 one: 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.

Brink, David O., 1986, "Externalist Moral Realism", *Southern Journal of Philosophy*, 24 (Supplement): 23–40.

Boyd, Richard, 1988, "How to be a Moral Realist", in Sayre-McCord 1988, 187–228.

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

Bramble, Ben (2016). A New Defense of Hedonism about Well-Being. *Ergo: An Open Access Journal of Philosophy* 3.

Byrne, Alex & Hilbert, David R. (2002). Philosophical issues about colour vision. In L. Nagel (ed.), *Encyclopedia of Cognitive Science*. Macmillan.

Chalmers, David J. (2003). The content and epistemology of phenomenal belief. In Quentin Smith & Aleksandar Jokic (eds.), *Consciousness: New Philosophical Perspectives*. Oxford University Press. pp. 220--72.

Chalmers, David J. & Jackson, Frank (2001). Conceptual analysis and reductive explanation. *Philosophical Review* 110 (3):315-61.

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

Clark, Austen (2000). Quality space. In Austen Clark (ed.), *A Theory of Sentience*. Oxford University Press.

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

Deijl, Willem (2019). Is pleasure all that is good about experience? *Philosophical Studies* 176 (7):1-19.

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.

Finlay, Stephen, 2014, *Confusion of Tongues: A Theory of Normative Language*, Oxford: Oxford University Press.

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.
- Lutz, Matthew and Lenman, James, "Moral Naturalism", The Stanford Encyclopedia of Philosophy (Fall 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2018/entries/naturalism-moral/>>.
- Kriegel, Uriah (2013). The Phenomenal Intentionality Research Program. In U. Kriegel (ed.), *Phenomenal Intentionality*. Oxford University Press.
- Kriegel, Uriah (2015). *The Varieties of Consciousness*. Oxford University Press.
- Kriegel, Uriah (2019). The Value of Consciousness. *Analysis* 79 (3):503-520.
- Lee, Andrew Y. (2018). Is consciousness intrinsically valuable? *Philosophical Studies* 176 (3):655-671.
- Lee, Andrew Y. (2019). The Microstructure of Experience. *Journal of the American Philosophical Association*:1-20.

-
- Lin, Eden (2018). Attitudinal and Phenomenological Theories of Pleasure. *Philosophy and Phenomenological Research*.
- Mackie, John L. (1977). *Ethics, Inventing Right and Wrong*. Penguin Books.
- 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 Explanatory 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 (1986). *The View From Nowhere*. Oxford University Press.
- Parfit, D. 2011. *On What Matters*, Oxford: 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.
- Railton, P., 1986, "Moral realism", *The Philosophical Review*, 95(2): 163–207.

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/>>.

Ridge, Michael, "Moral Non-Naturalism", *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2018/entries/moral-non-naturalism/>>.

Rosenthal, David (2015), "Quality Spaces and Sensory Modalities," in *The Nature of Phenomenal Qualities: Sense, Perception, and Consciousness*, ed. Paul Coates and Sam Coleman, Oxford: Oxford University Press, pp. 33-65.

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.

Shafer-Landau, Russ. 2003. *Moral Realism: A Defense*. Oxford: Clarendon Press.

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.

Siewert, Charles (1998). *The Significance of Consciousness*. Princeton University Press.

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.

Street, Sharon (2010). What is constructivism in ethics and metaethics? *Philosophy Compass* 5 (5):363-384.

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

Sturgeon, Nicholas, 1985, "Moral Explanations", in *Morality, Reason, and Truth*, David Copp and David Zimmerman (eds.), Totowa, NJ: Rowman and Allanheld, 49-78.

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

Väyrynen, Pekka (2013). Grounding and Normative Explanation. *Aristotelian Society Supplementary Volume* 87 (1):155-178.

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.