Syracuse University

SURFACE

Dissertations - ALL   SURFACE

July 2016

Seeing Right from Wrong: A Defense of A Posteriori Ethical Intuitionism

Preston John Werner
Syracuse University

Preston J. Werner
Seeing What's Right and Wrong: A Defense of *A Posteriori* Ethical Intuitionism
Dissertation Abstract

This dissertation develops and defends *A Posteriori* Ethical Intuitionism, the view that our epistemic access to moral properties is grounded in perceptual experience. It contains two parts. In part I, I present the epistemic access problem for realist moral epistemology and argue against several *a priori* attempts to resolve it. In part II, I defend *A Posteriori* Ethical Intuitionism and its ability to resolve the epistemic access problem.

Part I begins by arguing that evolutionary debunking arguments are best understood as claiming that we lack epistemic access to mind-independent moral facts. The remainder of part I rejects several attempts to respond to this challenge. I first argue that even sophisticated versions of reflective equilibrium, as coherentist theories, fall victim to a "garbage in, garbage out" objection. I conclude that a proper solution to the epistemic access problem must be foundationalist. I then survey other influential *a priori* approaches to epistemic access. I conclude that each either fails to ensure epistemic access, or succeeds by providing epistemic access to the wrong properties—viz. not the robustly normative ones.

Part II begins with an argument that we can have perceptual experiences of moral properties by making use of a "method of contrast" argument of the sort found in the philosophy of perception literature. Next, I defend the foundationalist credentials of *a posteriori* ethical intuitionism against the objection that moral experiences are

epistemically dependent on (moral) beliefs. I then discuss the role of emotions in generating moral perceptions. If emotions are themselves perceptual, their role in producing moral perceptual experiences is not problematic. Some have argued that emotions cannot be perceptual since we can be morally responsible for having (or failing to have) certain emotional experiences. I argue against this claim, and in favor of a perceptualist conception of emotions. Finally, I return to the issue of epistemic access. I argue that the perceptual view provides an explanation of our epistemic access to the moral properties, and that such a view is compatible with the non-causal nature of those properties.

Seeing Right From Wrong: A Defense of *A Posteriori* Ethical Intuitionism

by

Preston John Werner

B.A. University of Iowa, 2007
M.A. University of Nebraska, 2012

Dissertation
Submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in *Philosophy.*

Syracuse University
July 2016

This dissertation is for Steve Swartzer and Clare LaFrance, for teaching me how to see,

and to Nikki Fortier, Emily Garton, and Mose, for teaching me how to feel.

TABLE OF CONTENTS

Chapter 1

How to Understand the Epistemic Access Problem for Non-Naturalist Moral Realism

Introduction

A longstanding objection to non-naturalist moral realism (henceforth "non-naturalism") is that it is incompatible with any plausible explanation of how we can come to have knowledge of moral truths.[1] As most non-naturalists accept, our beliefs about these entities could not be caused by them.[2] It is mysterious, the objection goes, how we could have epistemic access to these non-causal entities in order to form justified beliefs about them. I call this kind of epistemological objection the *access objection*.[3] Though the objection is frequently gestured at, it is not often developed in much detail. Nevertheless, it is an objection generally taken seriously (and rightly so).

The purpose of this paper is to clarify and develop the various forms of the access objection and propose a unique formulation that I believe is superior to those currently found in the literature. The key in developing a proper understanding of the *access objection* is to do so in a way that is charitable both to the proponent of the objection as well as the non-naturalist. This means that the objection should neither be

---

[1] Arguably, this objection goes at least back to Mackie (1977), Ch. 1. Harman (1977), Ch.1 presses an arguably similar worry, though it is in terms of explanation rather than causation. Two more recent versions of this objection, which emphasize evolution as particularly damning for non-naturalist epistemology are Joyce (2001) and Street (2006).

[2] Huemer (2008), Ch.5, Scanlon (forthcoming), Lectures 1-2, Dworkin (2011), Ch.4, and Ridge (2008), Section 3.

[3] As will become clear below, there is no single *access objection*, but many substantively different ways of pressing a general epistemological concern about our access to objective moral truths. So I mean "*access objection*" to apply to a range of epistemological objections to non-naturalism.

formulated in a way that is clearly misguided, nor in a way that begs the question

against the non-naturalist. I develop my formulation of the objection by taking a step

back from metaethics to epistemology more generally. I do this in order to motivate a

general epistemological principle I call the *access requirement*, which requires that an

adequate epistemology for a given domain explain how our justified beliefs can be in

some way responsive to the facts of that domain.[4] The *access requirement* is closely

related to the fact that, as Lawrence BonJour says, "one crucial part of the task of an

adequate epistemological theory is to show that there is an appropriate connection

between its proposed account of epistemic justification and the cognitive goal of *truth*."[5]

If this is right, then a theory of justification for any given domain, ethics included, must

meet the *access requirement*.

1. Benacerraf's Problem

What I am calling the *access objection* to moral realism can be traced back to Paul

Benacerraf's (1973) "Mathematical Truth". As its title suggests, Benacerraf was worried

about the epistemology of mathematics for a Platonist conception of mathematical

objects. David Liggins (2010) helpfully reconstructs Benacerraf's argument as follows:

> (B1) If mathematical platonism is true, then we have knowledge of abstract
>   mathematical entities.
> (B2) If we have knowledge of abstract mathematical entities, then we are
>   causally related to them.
> (B3) We are not causally related to abstract mathematical entities.

---

[4] For the more precise formulation, see section 5 below.
[5] BonJour (1985), pp.108-109.

Therefore:
(B4) Mathematical platonism is not true.[6]

The corresponding argument against non-naturalism is straightforward:

(M1) If non-naturalism is true, then we have knowledge of non-natural
     properties. [7]
(M2) If we have knowledge of non-natural properties, then we are causally
     related to them.
(M3) We are not causally related to non-natural properties.
Therefore:
(M4) Non-naturalism is not true.

Each premise of both arguments can be questioned. I'll focus on the moral case since

that is the focus of this paper. While (M1) is denied by skeptical non-naturalists, and

(M3) is denied by at least some naturalist moral realists,[8] most non-naturalists deny

(M2). As a result, the focus of this paper is on (M2) and variations of it. When

Benacerraf formulated his argument against Platonism, the causal theory of knowledge

was in vogue.[9] If the causal theory of knowledge were true, (B2) and (M2) would be true

as well. However, it's now widely accepted that the causal theory of knowledge is

[6] Liggins 2010, p.68.
[7] I refer to non-natural properties throughout my formulations of Benacerraf-style arguments. The argument could also be run in terms of moral facts, moral truths, moral entities, etc. I stick to "properties" throughout for easier bookkeeping, but nothing I say should turn on this.
[8] E.g. Boyd (1988), Sturgeon (1985), Railton (1986). Oddie (2005/2009) is a self-proclaimed non-naturalist who denies (M3), but he is an exceptional case among non-naturalists. Shafer-Landau (2003/2005) believes that moral properties are constituted out of, but not reducible to, natural properties. Whether someone with this view would deny (M3) depends on whether she thinks that an object inherits its causal properties from the parts that constitute it. As far as I know, Shafer-Landau takes no stand on this question.
[9] The *locus classicus* is Goldman (1967, 1976), though he now rejects the view. See Ichikawa &Steup (2012), Section 6.2 for the analysis and its historical development.

inadequate.[10] As a result, Benacerraf's argument and its moral counterpart are unsuccessful as they stand as arguments against Platonism and moral realism.

Despite the rejection of Benacerraf's assumed theory of knowledge, Benacerraf-style arguments have remained influential. Many have thought that Benacerraf was on to something when he formulated his argument, even if he was wrong about accepting the causal theory of knowledge. I agree with this sentiment. The challenge, then, is to articulate a Benacerraf-style argument which neither relies on a contentious theory of knowledge nor begs the question against the non-naturalist. I will discuss three such attempts, explaining why I think each of them is inadequate. Then I will present and defend my own version of the problem with an emphasis on the premise I use to replace (M2), which I call the *access requirement*.

2. Restricting the Causal Requirement to Existential Claims

Colin Cheyne (1998, 2001) offers a first attempt at refining Benacerraf's argument. Cheyne agrees that the causal theory of knowledge, as it stands, is too strong—unless we are willing to significantly weaken what counts as a cause.[11] Instead, he argues that the causal requirement only applies to knowledge of a given entity's *existence*, not for any of its particular properties. More formally, his condition is:

> (CE) "We cannot know that *F*'s exist unless our belief in their existence is caused by at least one event in which an *F* participates."[12]

---

[10] See, e.g. Ichikawa & Steup (2012), Section 6.2 and Cheyne (2001).
[11] See Cheyne (2001), Chs. 5 and 7.
[12] Cheyne (2001), 99.

This would lead us of the following understanding of (M2):

> (M2a) If we have knowledge of the existence of non-naturalist properties, then our belief in their existence is caused by at least one event in which they participate.[13]

Cheyne motivates (CE), and therefore indirectly (M2a), by providing examples from the history of science in which scientists intuitively did not have knowledge of a theoretical posit until they causally interacted with it. One such example is the discovery of germanium. Dmitri Mendeleev predicted the existence of germanium before he had had causal interaction with it. This is because he had a relatively well-confirmed chemical theory that, given a "gap" in his periodic table, predicted that there was an as-yet-to-be-seen element with the properties that germanium has. However, Cheyne argues that, intuitively, Mendeleev did not know of the existence of germanium given his theory alone:

> Mendelee[v] may have believed on the basis of the gap in his periodic table that germanium might exist and we may even agree that it was a reasonable assumption, but it still remained for the element to be discovered. Knowledge that germanium exists required causal contact, however remote, with atoms of germanium.[14]

Given that Mendeleev's theory was relatively well-confirmed, Cheyne seems to acknowledge here that Mendeleev had some evidence that germanium existed, by way

---

[13] Cheyne (2001), 99.
[14] Cheyne (2001), 100.

of an inference to the best explanation. However, he claims that Mendeleev cannot be credited with full-blown knowledge until he has causally interacted with germanium.

I don't believe that plugging (M2a) into our Benacerraf-style argument will result in the strongest version of the argument, for two reasons. First, in order for (M2a) to play a role in a Benacerraf-style argument against non-naturalism, we would need to change the other premises of the argument, lest we end up with a much weaker conclusion. (M2a) is only a claim about the knowledge of the *existence* of non-natural properties, not about their nature. This is an important difference, since Cheyne acknowledges that (CE) is compatible with knowledge of the nature of an entity that we don't yet *know* to exist. So, even if (M2a) is correct, we may still have knowledge of many first-order normative claims, even if those claims are about non-natural properties. In order for the argument to succeed with its intended strength, we would have to change (M1) to:

> (M1a) If non-naturalism is true, then we have knowledge of the existence of non-natural properties.

Unlike (M1), rejecting (M1a) is consistent with the rejection of skepticism about first-order normative claims. This makes the rejection of (M1a) much easier for the non-skeptical non-naturalist. So one important reason why (M2a) is not the best way to reformulate (M2) is that it won't result in a skeptical argument with nearly as wide a scope as Benacerraf-style arguments are supposed to have.

Secondly, there are good reasons to believe that (CE) is false. Since (M2a) is a corollary of (CE), this should cast doubt on (M2a) as well. Cheyne grants that a well-confirmed theory can provide *some* evidence of the existence of a particular entity in order for his theory to be at all plausible, as for example when he agrees that it is reasonable for scientists to search for entities predicted by their theories. However, once it is granted that we can have *evidence* for the existence of *F*s before we have causally interacted with *F*s, it is hard to see how it couldn't be possible to have *knowledge* of the existence of *F*s before we have causally interacted with them: All we have to do is sufficiently ratchet up the evidence. Perhaps Mendeleev in particular didn't have a sufficiently well-confirmed theory in hand to count as knowing that germanium existed before causally interacting with it. But suppose we gave Mendeleev all other non-germanium evidence involving experimental results and theory in modern chemistry. Wouldn't he then be able to infer the existence of germanium via an inference to the best explanation?[15] The conclusion to draw here seems to be that, at the very least, best explanation arguments can in principle give us knowledge of the existence of something, even when we have not causally interacted with it. (M2a) is not an acceptable premise to replace (M2) in a Benacerraf-style argument.

3. Non-Accidentality

---

[15] I thank Luke Elwonger, Adam Thompson, and Mark van Roojen for helpful discussion on this point.

A second way to revise Benacerraf's argument is to speak in terms of a more epistemically ecumenical notion, that of non-accidentality. It's widely accepted that if it is merely an accident that a belief is true, then such a belief cannot be justified. If I believe on a whim that there is an even number of stars in the milky way, I wouldn't be justified even if my belief turned out to be true. This is because it would be sheer luck that I formed a belief that turned out to be true. More generally, the argument goes, if we can't explain why our beliefs about a particular domain of facts, such as moral truths, are non-accidentally true, then we should throw out all of our beliefs about that domain as unjustified (or at least be agnostic about their justificatory status). This provides us with a second attempt at revising premise two in our Benacerraf style argument:

> (M2b) If we have knowledge of non-natural properties, then it is not a mere accident that (at least some of) our beliefs are true.

(M2b) seems to be defended by many philosophers who discuss Benacerraf-style objections. Michael Huemer, for example, after commenting on the flaws with Benacerraf's original argument, notes that

> [t]here is a more general condition on knowledge that everyone in epistemology accepts: I know p only if it is not a mere accident (not a matter of chance) that I am right about whether p. The challenge for the moral realist, then, is to explain how it would be anything more than chance if my moral beliefs were true, given that I do not interact with moral properties.[16]

---

[16] Huemer (2005), 123. Bedke (2009, forthcoming) prefers to understand some kinds of arguments for moral skepticism in this way as well. Enoch (2011, Ch. 7 Sect. 2) talks in terms of establishing correlations between beliefs and facts, but as I understand him he is presenting the non-accidentality reading using different terminology.

(M2b) is superior to (M2a) in that it doesn't beg the question against non-naturalists who believe moral properties are causally inert. Nor does it rely on a contentious epistemological claim. However, though this non-accidentality condition on knowledge is probably *necessary* for meeting a Benacerraf type worry, it is too weak to provide as strong of an argument as possible on the part of the moral skeptic. To see why, consider one thing David Lewis says in response to the Benacerraf worry in *On the Plurality of Worlds*: "[I]f it is a necessary truth that so-and-so, then believing that so-and-so *is* an infallible method of being right. If what I believe is a necessary truth, then there is no possibility of being wrong. That is so whatever the subject matter of the necessary truth, and no matter how it came to be believed."[17]

If, as is widely believed, the moral truths are necessary, then any moral belief falls under this class of beliefs that Lewis points out has "no possibility of being wrong", in some sense. But a proposition's modal status shouldn't make justified beliefs about that proposition so trivial. I hope this is intuitive enough, but I'll say a little bit by way of defending it. For one thing, this would multiply knowledge beyond what is plausible—seemingly all true metaphysical beliefs would count as justified on this view. Imagine the theist defending her belief in God as justified given that if God exists,

---

[17] Lewis (1986), pp.114-115. For what it's worth, even Lewis himself seems ambivalent with this response.

God exists necessarily. This would not be a satisfactory justification for a theistic belief.[18]

Secondly, we want our theory of justification to be able to carve up good reasons from believing a proposition from bad ways of believing a proposition. If any reason whatsoever is at worst an epistemically neutral reason to believe a necessary proposition, it will be hard to make such distinctions. The non-accidentality that we want for justification is not just any non-accidentality, but one that ties the believer's reasons for believing with the fact in question. This epistemology of necessary truths on-the-cheap will not get us the right relationship of non-accidentality.

However, maybe this is an uncharitable reading of (M2b). Maybe we should understand the non-accidentality condition as applying to the outputs of the methods we use to form the beliefs. So, for example, if I decide whether a given action is wrong by flipping a coin, my belief won't be justified even if it turns out to be true, since the coin could have just as easily landed differently.[19] Understanding "not a mere accident" in (M2b) in this different way is a slight improvement. But in the end, this won't provide an adequate fix. This is because some methods will non-accidentally give the same results every time, whether or not they are at all related to the domain in question. If I have some complicated way of inferring moral claims from purely mathematical claims, my beliefs in those moral claims will not be knowledge, even the ones that are

---

[18] A point made by Field (1991), 236. Similar points are made with respect to Safety and Sensitivity conditions on knowledge by Manley (2007, Sections 1-2) and DeRose (2004, Section 7).
[19] See Manley (2007, Sections 1-2).

true, and even though it is no accident that, given my method, I deduced those conclusions. Now let's suppose that, as is plausible, my math-to-morals deductions are no better than chance at resulting in true beliefs. This might suggest that if we were to bulk up the non-accidentality condition in the most charitable way, we would end up with a non-accidentality condition not on particular beliefs, but on the methods used to form those beliefs. In other words, general reliability is the important epistemic feature that non-accidentality is trying to capture. I consider this possibility presently.

4. Reliability

The third way of formulating the Benacerraf-style worry that I will consider before giving my preferred understanding is to state the second premise in terms of reliability. If our beliefs about a particular domain are to count as knowledge, an adequate epistemology of that domain must be able to explain the reliability of the belief forming processes or methods that we use. For example, an adequate epistemology for the Platonist about mathematics must be able to explain the reliability of mathematicians' beliefs about mathematical claims. With all of this in mind, here is a third try at revising premise two of the (moral) Benacerraf argument:

> (M2c) If we have knowledge of non-natural properties, then there is some explanation of why most of our moral beliefs are formed reliably.

Hartry Field, who provides one of the most widely cited Benacerraf-style arguments in the philosophy of math literature, states his argument both in terms of reliability[20] as well as the related but non-modal notion of *actual* accuracy.[21] On the moral side, in a *Stanford Encyclopedia of Philosophy* entry Michael Ridge states Benacerraf style worries in terms of reliability, suggesting that this is at least one standard way of understanding the argument.[22]

I think (M2c) is on the right track, but I don't believe that it is the best way to formulate a Benacerraf-style argument. (M2c) makes the argument too strong and is thus unfair to the non-skeptical non-naturalist.[23] To see why, it will be helpful to notice a disanalogy between the case of mathematics and the case of morality. Consider the mathematical version of (M2c):

(B2c) If we have knowledge of mathematical properties, then there is some
        explanation of why most of our mathematical beliefs are formed reliably.

(B2c) is a plausible demand on an epistemology for Platonists in math because it's uncontroversial that people generally do a pretty good job at forming true mathematical beliefs. Even if this is doubted, the neo-Benacerrafian can at least demand that the Platonist provide an explanation of why most of the mathematical beliefs of professional mathematicians are formed reliably. There is a clear and very serious cost

---

[20] Field (1989), 233-238.
[21] Field (1989, 238).
[22] Ridge (2008), Sect. 3.
[23] As we'll see below, I also think (M2C) is too weak in a way that helps to motivate my preferred Benacerraf type argument.

to the Platonist who would deny (B2c), but if the Platonist is forced to claim that even mathematicians fail to form mathematical beliefs reliably, they might as well embrace skepticism. In short, (B2c), or perhaps something slightly weaker, is a fair demand to make of the mathematical Platonist since it is a datum that mathematicians are generally very good at reliably forming true mathematical beliefs.

However, unlike in the case of mathematics, there is no consensus even amongst non-skeptical moral realists that people are generally reliable when forming moral beliefs. Huemer, who is a non-naturalist, explains moral disagreement by pointing to 14(!) widespread sources of erroneous judgments.[24] He is happy to accept that the majority of our moral beliefs may be false, compatible with his non-skeptical realism. More generally, it has been shown that moral judgments (at least about particular cases) are susceptible to bias, framing, and responsive to other seemingly non-normative features of cases.[25] It's difficult to see how realists would respond to this challenge without admitting that people are prone to unreliability when forming moral beliefs. But this is perfectly compatible with the rejection of skepticism.[26] The upshot here is that

---

[24] Huemer lists bias, miscalculation, confusion, misunderstanding, oversight, hasty judgments, incomplete information, unarticulated assumptions, stubbornness, fallacies, forgetfulness, the intrinsic difficulty of issues, inarticulate evidence, and mental defects as just some common sources of error. See Huemer (2005), pp.137-139. See also Huemer (2008).

[25] Cushman et. al. (2008), Sinnott-Armstrong (2006), Ahlenius and Tánnsjó (2012) are three such examples.

[26] As Mark van Roojen points out in a similar context, "[a] somewhat unreliable process can be part of a more reliable overall process, and the overall process might be less reliable if it lacked the only somewhat reliable sub-part." (van Roojen, forthcoming, 12) Van Roojen develops a non-skeptic friendly explanation of these empirical results by limiting the epistemic force that intuitions have. On his view, intuitions have some positive epistemic status, but not enough to generate full blown justification without relations of coherence to further support their reliability. I take it that his view is compatible with many non-skeptical versions of moral realism.

unlike the mathematical case, there is no need to salvage the overall reliability of moral

beliefs, since even non-skeptical realists can admit that many run-of-the-mill moral

beliefs are unreliably formed. (M2c) places an unfair epistemological demand on the

non-naturalist.

Would a more restricted version of (M2c) lower the epistemological demand on

the non-naturalist to a more reasonable requirement? Consider something like (M2c*):

> (M2c*) If we (human beings) have knowledge of non-naturalist properties, then
> there is some explanation of why most of the moral beliefs *of moral experts*
> are formed reliably.

However, even this more restricted reading is too strong of a demand to place on the

moral realist, for similar reasons. Unlike in the mathematics case, there is no obvious

need to accommodate the existence of moral experts. Professional ethicists may not

provide an adequate counterpart to professional mathematicians.[27] And the existence of

moral experts outside of the philosophical community is highly disputed.[28] Appealing

to hypothetical moral experts would not help either, since that would only show that

justified moral beliefs are possible for beings that don't in fact exist. We are looking for

a moral epistemology for creatures like us. So at best, (M2c*) is only a demand on a

realist who does accept the existence of moral experts. Since the existence of moral

experts is not a requirement on a non-skeptical realism, (M2c*) remains too strong.

---

[27] Though for one interesting argument that professional ethicists *are* the relevant moral experts, see Sharvy (2007).
[28] See, e.g. Tobia et. al. (2012), Cholbi (2007), and Driver (2006).

The considerations just provided are reason enough to reject (M2c) (or M2c*) as the best way to formulate a Benacerraf-style moral argument. However, I want to talk about a different way in which I think (M2c) is a non-ideal premise, in that it is in another way too weak of a condition on moral knowledge. If I am correct about the above considerations, then this further argument is inessential to my case against (M2c) as a premise. However, I want to discuss why I believe (M2c) is also too weak, not only because I think it is an important point, but also because it will help to motivate the key premise in the version of the argument that I defend below.

A reliability condition on knowledge is similar to a non-accidentality condition on knowledge in that it is (at least partially) modal. Plausibly, when our belief forming practices for a given domain are reliable, it means that, at least with respect to many of those beliefs, if the truths that *actually* hold were false, we wouldn't have come to believe them. This condition has come to be called *sensitivity* in the epistemology literature.[29] It might be claimed that this aspect of reliability is irrelevant for moral beliefs since they are necessarily true if true at all. However, as Matthew Bedke argues:

> Quite generally, independent evidence that a target domain of fact holds necessarily does not insulate the particular beliefs in that domain from sensitivity-type tests. To think otherwise is to believe that justifying testimony to the effect that some domain's basic facts hold necessarily insulates those beliefs from insensitivity tests.[30]

---

[29] It originates with Nozick (1981), though he does not call it sensitivity.

[30] Bedke (forthcoming), 25. Manley (2007, Sections 1-2) makes a similar point with respect to *Safety*, an epistemological principle that is related to Sensitivity.

Bedke denies that sensitivity conditions are trivially met for our beliefs of necessary truths.[31] So on a Bedkean understanding, some of our tests for reliability will stretch beyond counterfactual claims to counter-possible claims.[32] [33]

I'm sympathetic to the spirit of Bedke's claim, but I think that the non-skeptical non-naturalist could reasonably balk at the suggestion that moral beliefs must meet some counter-possibility test in order to count as sensitive or reliable. For one thing, she could claim that this condition is an unreasonably high epistemic standard to require beliefs to meet, whether moral or not. But even if it is a reasonable epistemic condition, it's just hard to assess the truth or falsity of counter-possible claims. Even the non-naturalist, who doesn't identify moral properties with any natural properties, can plausibly claim that, e.g., suffering is *essentially* bad (though badness isn't *essentially* suffering).[34] So in order to assess the Bedkean reliability of our belief that suffering is bad, are we required to go not just counter-possible but counter-*essential*? If Bedkean reliability requires our ability to track truths even in counter-essential possibilities, it seems clear that it places too strong of a burden on the non-skeptical non-naturalist.[35] In

---

[31] This isn't quite right. Bedke seems to be willing to concede that *sensitivity* as an epistemic concept is trivially met for necessary truths. The more important point is that the relevant sensitivity-type epistemic condition on justification/knowledge is not trivially met for necessary truths. He calls any belief that doesn't meet this extended sense of sensitivity "oblivious". See Bedke (forthcoming), especially sections 4 and 5.

[32] Bedke talks of "allodoxic possibilities" instead of counter-possibles, but since I'm not worried about the details of his view here, I skip over this complication.

[33] Bedke's argument doesn't actually assume a reliability requirement. Instead, he just focuses on his sensitivity-type condition of "obliviousness". Nevertheless, I take it that if he were convinced that reliability was the important condition, he would want a belief's being oblivious (in his sense) to undermine its claim to reliability.

[34] FitzPatrick (2008) seems to have something like this view.

[35] I thank Mark van Roojen for pressing me on this point.

short, it will be very difficult to test for a belief's Bedkean reliability. So the non-

skeptical realist can reasonably read reliability in (M2c) in the traditional way and rest

easy, at least with respect to the sensitivity aspect of reliability.

I certainly don't think that meeting a Benacerraf-style argument with (M2c)

instead of (M2) is a walk in the park for the non-skeptical moral realist. However, I

think that Bedke is right that something is missing. (M2c) is a too weak. Suppose Steve

is born with the unshakeable belief that water is H_2O. He has no independent evidence

for the belief's truth. Steve's belief is not the result of some evolutionary adaptation; it is

nothing but a random mutation in his DNA.[36] If the belief is essential to Steve (after all,

it is directly genetically caused), then it is a necessary truth that he believes water is

H_2O. And suppose that this is the only belief Steve has which is caused by this hyper-

modulated genetic mechanism. Furthermore, it is necessarily true that water is H_2O. So

Steve believes water is H_2O at all and only the worlds in which water is H_2O. I'm

inclined to think that despite this fact, Steve's belief that water is H_2O is not an

upstanding epistemic denizen. But doesn't Steve's belief count as reliably formed, if we

understand reliability in the traditional way? Part of the problem here is that reliability

is subject to a generality problem, such that it is unclear what range of actual and

counterfactual beliefs that reliability must range over.[37] Even setting this complication

---

[36] My sincere apologies for the blatantly unrealistic pseudo-science involved in the case.

[37] See Goldman (2008), Sections 3 and 4 for a discussion of the generality problem for reliabilism as an analysis of knowledge, and some responses to it (with references).

aside, it's unclear how to test reliability in a case like this—and yet it is something like

this kind of explanation of the reliability of our moral beliefs that some non-skeptical

realists appeal to in responding to Benacerraf-style arguments.[38] If Steve's belief about

water is not in good epistemic standing, we shouldn't accept these kinds of responses as

being enough to show that our moral beliefs are in good epistemic standing either. In

short, there is no non-contentious way of spelling out the reliability condition that

provides a strong enough condition on the non-skeptical moral realist.[39]

5. The *Access Requirement*: A Benacerraf Inspired Epistemic Requirement for Moral Realism

5a. Formulating the Access Requirement

We have now seen some attempts to formulate a Benacerraf-type argument

against non-skeptical non-naturalism. I have argued that all are inadequate in at least

one way. What we are looking for is a formulation of an epistemic demand that avoids

the kinds of problems that beset the other attempts above. Two issues stand out as

especially important. First, the argument can't be formulated in a way such that it leads

to independently epistemologically implausible consequences by granting

justification/knowledge to those that clearly don't have it, such as the case of Steve

given above. Second, the argument can't be formulated in a way that is question

---

[38] See Enoch (2011), Ch. 7, Wielenberg (2010), pp.447-452, and Skarsaune (2011), pp.233-236, for examples of arguments with this structure.
[39] Setiya (2012), pp.89-92 makes a similar point about reliability.

begging against the non-naturalist, such as by claiming that knowledge of a property requires a causal connection to that property. In this section, I begin by proposing and explaining my preferred alternative, what I call the *access requirement* (AR). Then I explain why AR can capture the virtues of the alternative approaches without the vices. In particular, I think that AR can help to explain why the challenge is often understood in terms of reliability while neither relying on the controversial epistemological problems with reliability, nor (falsely) committing the realist to explaining how most of our moral beliefs are true.

> Without further ado, here is my proposed *access requirement*:

> (AR) Any complete non-skeptical epistemology for a particular domain *d* must explain how we have epistemic access to the *d*-facts such that we are able to form beliefs that are responsive to those facts.

Three clarifications: First, (AR) is a general epistemic requirement.[40] This is an advantage, since it helps to illustrate that the demand on moral epistemology is no different from that of any other domain. Second, (AR) is meant to be a *necessary* condition on justification/knowledge of a particular domain, not a *sufficient* condition. Epistemic access is relatively cheap—we could in principle have epistemic access to a domain and yet be pervasively mistaken about the truths of that domain. It thus avoids

---

[40] One general epistemological statement of something in the spirit of (AR) is found in Sosa (1985): "[J]ustification could not possibly turn out to be a property that a belief might possess in complete independence of the truth of its object." (1985, p.13).

committing the non-skeptical realist to reliability amongst normal human beings.[41]

Third, (AR) is an explanatory demand to be met by the moral epistemologist, not a first-order demand to be met by moral agents.[42] It asks how, given some particular ontology of moral facts, our first-order moral beliefs could in any way be responsive to the truths that they purport to be about.

5b. Is (AR) just the causal requirement in disguise?

With its language of "epistemic access" and beliefs' being "responsive" to facts, the moral realist may be concerned that I have merely rephrased Benacerraf's causal requirement on knowledge in (slightly) less causal language. If the only way in principle to have "epistemic access" to the facts of a given domain is via being causally related to it, then (AR) provides nothing but a more sophisticated way of begging the question against the non-skeptical moral realist. I certainly don't intend access or responsiveness to be understood merely causally. Unfortunately, I am unable to spell out precisely what epistemic access and responsiveness are, for two reasons. First, any further explication would continue to look misleadingly causal or spatio-temporal in a

---

[41] Recall from above that there is an important difference between the possibility of knowledge in some domain and actual reliability in that domain. The non-naturalist arguably must only show the former to respond to Benacerraf-style objections.

[42] It is important to stress this point, since "access" is a term used a lot in the epistemological literature concerning internalism vs. externalism about justification/knowledge. The internalist about justification accepts, while the externalist rejects, that whenever a believer is justified in believing something, the basis of that justification is accessible upon reflection to her. (AR) is silent on the dispute between internalists and externalists about justification. In other words, it is a meta-justificatory requirement rather than a justificatory requirement. For the importance giving a meta-justification for the epistemology of a given domain, even for an internalist about justification, see BonJour (1985), pp.8-15.

way I don't intend. I could say that access is about "being in contact with" the truths in question, or that responsiveness is about forming our beliefs in light of the facts they are about, rather than in light of something unrelated. But I'm pretty sure anyone worried about the language of "access" and "responsiveness" will not be any less worried about the "being in contact with" or "in light of" relations. Second, "epistemic access" and "responsiveness" should be understood very widely, and any further precisification of these notions may rule out ways of meeting (AR) which are in principle possible.

Readers so far unconvinced would be right to be unsatisfied by these remarks. So instead of leaving my explanation of epistemic access with potentially unhelpful metaphors about being in contact with truths or truthmakers, I want to try to make it more clear that (AR) is more than just a re-statement of a causal requirement by going through several different kinds of access that meet the (AR) for a given domain of facts:

> 1. *Causal access*. This is perhaps the most widespread and uncontroversial kind of access to the external world. Beliefs with this sort of access are responsive to the facts they are about via a causal relationship, for example perceptual beliefs.

> 2. *Introspective access*. Though the reliability of introspection has been questioned,[43] it's plausible that we have some special access to our own mental states, however fallible it may be. I take it that even though introspection involves underlying causal brain processes, it is from an epistemic standpoint a different kind of access than causal access.[44] On a traditional sort of model of introspection, our introspective beliefs are responsive to the facts they are about via a relationship of direct acquaintance.

---

[43] See, e.g. Schwitzgebel (2010), Section 4.

[44] I don't want to take a stand on whether, or to what extent, introspection should be subsumed under the category of causal access. I include it in the list because it seems to have been thought to be epistemically distinct in some special way by many philosophers, and my intention here is only to give a list of possibly different forms of epistemic access. See Schwitzgebel (2010).

3. *Constitutive access*. We have constitutive access to a truth *t* when something about our coming to believe *t* is partially constitutive or provides evidence for what is partially constitutive of its being the case that *t*. Arguably, many beliefs about response-dependent properties involve constitutive access.[45] Suppose that something is beautiful iff it is believed to be beautiful by all/many/some normal adult human beings.[46] A normal adult human being comes across a Chuck Close painting and comes to believe that it is beautiful. She has constitutive access to the fact that the painting is beautiful insofar as her belief is partially constitutive of that fact.

4. *Conceptual access*. A certain philosophical school of thought claims that we can learn a lot about the concepts we possess competently by conceptual analysis, which may involve reflecting on how we would apply them in various scenarios.[47] This is one potential explanation and defense of analytic knowledge (assuming that there are analytic truths).[48] Again, though this kind of *a priori* reflection would be underwritten by causal (and possibly also introspective) processes, the *access* in question is not causal or introspective, because of the nature of the truths in question. On this view, analytic truths are not causally related to us, nor are they merely facts about our own mental states.[49]

5. *Intuitive access*. On a recalcitrant but controversial view, intuitions provide a source of justification in a wide variety of philosophical disputes.[50] Intuitions have been defended as a source of justification for beliefs about modality, metaphysics, epistemology, logic, and ethics. Even amongst those that accept that intuitions justify, the question of *how* intuitions can provide justification is a matter of dispute. But any such account would have to explain how intuitions provide access to the truths that they are about, in accordance with (AR). Here are two such (oversimplified) stories that, if they can be made to work, would meet the (AR):

> 1. Intuitions involve direct, *de re* acquaintance with properties (usually conceived as Platonic universals). Via *de re* acquaintance with the property *greenness*, and *de re* acquaintance with the property *redness*, we can come to know that nothing can be both green and red all over. In at least some

---

[45] As with introspection, I don't intend to take a stand on the epistemology of response-dependent properties. My intention here is only to give a list of possibly different forms of epistemic access.

[46] There are many complications I am ignoring here, not the least of which is how to define "normal" in a non-circular way.

[47] Jackson (1998, 2001), Audi (2008b), Russell (2011).

[48] Robert Audi (2005, 2008a), for example, believes that at least some substantive moral knowledge is conceptual, in the sense that the wrongness of certain actions is "contained in" the moral concepts alone.

[49] For a book length defense of analyticity, see Russell (2011).

[50] For an overview with an extensive bibliography, see Pust (2012).

cases, perceptual experience may ground our grasp of the relevant properties, but perceptual experience alone cannot justify the synthetic truth in question. Our access to these truths is met through our *de re* acquaintance with the properties.[51]

2. On a related view, intuitions involve grasping relationships between our concepts. However, according to such a view, this grasping can provide us with more than mere conceptual knowledge. The grasping of our concepts can (indirectly) provide us with knowledge of the relations between properties (again, usually conceived as Platonic universals). It provides us with this knowledge not through some *de re* acquaintance relationship, as in (1), but through the fact that our concepts *can't help but* pick out some property, since there are properties for any consistent concept. We might call this *Description-Constituting Access* (not to be confused with constitutive access), since how we construct our concepts determines which (mind-independent) property we are picking out.[52] (This view has the upside of avoiding the mystery of (1), but the downside of relying on a controversial metaphysical commitment to the abundant theory of properties.)

Two important caveats about this list: First, I don't intend it to be exhaustive. There may very well be other kinds of access that I am overlooking. My intention here is only to provide ample evidence that (AR) isn't smuggling in the causal assumptions of some Benacerraf-type arguments. Second, I don't want to take a stand here on which of these kinds of access can in fact provide justification for some subset of beliefs. This is compatible with (AR), since (AR) is only a *necessary* condition on justification, not a *sufficient* one. There may be no such thing as an analytic truth, in which case there will

---

[51] I think this view is more widely held than it is explicitly endorsed, though it is not a popular view. It is most strongly endorsed by Russell (1912), but something like it seems also to be held by Bealer (1998), Bonjour (1998), and Chudnoff (2011, forthcoming).

[52] "*Description-Constituting Access*" is my phrase. This is the view that Huemer (2005, Ch.5) seems to endorse. Balaguer (1998) uses this abundant view of properties in an attempt to solve the Benacerraf problem for mathematics: Roughly, since there are so many mathematical objects, any consistent mathematical practice will pick out some of them, and so our mathematical practices determine which mathematical objects to which we are referring; so we can't but have knowledge about some mathematical objects or others..

be no instances of conceptual access. Introspection may be altogether too unreliable to provide justification, despite the possibility of introspective access. Many philosophers are skeptical that any account of intuitions could avoid serious metaphysical and (non-Benacerraf) epistemological objections. Again, the point here is only to illustrate that there is nothing essentially causal about (AR), so it is not begging the question against the non-skeptical non-naturalist in the way that Benacerraf's original argument is. Interestingly, all five of these forms of access could plausibly be appealed to as part of a moral epistemology, though perhaps not all as part of a *realist* moral epistemology.[53]

5c. (AR) in relation to the non-accidentality (M2b) and reliability conditions (M2c)

I have tried to show that "access" and "responsiveness" need not be understood in a causal way. Instead, *access* is about bearing some kind of positive epistemic relation to the domain of truths (or truthmakers) that our beliefs are about. I hope it is clear enough given the above discussion that (AR) is more epistemically ecumenical than some kind of causal requirement on justification/knowledge, and thus is superior to the traditional Benacerraf argument (or Cheyne's related causal version of the argument).

---

[53] Causal access is most obviously compatible with naturalist realism, but Oddie (2005/2009), who is a non-naturalist, also argues that our access is (at least partially) causal. Presumably, introspective access best fits with some form of subjectivism. (What about a view that proposed that we are all born with innate knowledge of the moral truths, because, e.g. God placed them in us upon birth? A view like this would offer an introspective methodology, but our access would be causal, not introspective, since our connection to the moral truths would be through God, not through introspection directly.) I believe we would have constitutive access to the moral truths on some versions of metaethical constructivism and some response-dependent accounts. Conceptual access may not seem like a plausible way of getting access to substantive normative truths, but on one reading this is what Audi (2008a) has in mind by his *containment* relation. Intuitive access is probably the most common answer to Benacerraf-type worries given by non-naturalists—see Huemer (2005), Parfit (20 11), Vol. II Ch. 32.

Now I want to show how (AR) is a superior alternative to *non-accidentality* (M2b) and

reliability (M2c) in formulating a Benacerraf argument that both (a) provides a real (and

legitimate) challenge to the realist moral epistemologist while (b) not simply begging

the question against her. I'll focus my discussion on *reliability*, but I will try to make

clear how I think what I say about (AR)'s improvements over *reliability* extend to *non-*

accidentality as well.

Recall from above that the biggest problem with the *reliability* premise is that it is

too strong, in that it demands the non-skeptical realist to defend what she need not be

committed to: That the moral beliefs of most people, or at least serious ethicists, are

reliably formed. Certainly the non-skeptical realist *could* accept this, and these so-called

"preservationists" about morality may have to meet this stronger reliability premise in

order to defend a fully adequate moral epistemology.[54] But we are looking for a fully

general Benacerraf-style argument, so this won't do.

(AR) avoids this weakness, since *access* doesn't imply reliability—in fact, it is

compatible with pervasive error about a given domain. For example, psychological

research has provided much evidence that memories can be influenced and altered in

many epistemically unsavory ways.[55] If this evidence isn't enough to convince you that

we should consider memory unreliable, just suppose that erroneous memories were

---

[54] The term "preservationist" is Unger's (1996, p.11).
[55] For just a few examples, see Loftus and Greene (1980), Powers et. al. (1979), and Loftus et. al. (1978).

even more widespread than they in fact are. Still, and whether or not in such a situation

we would want to count memories as justifying, they would meet the (AR), since at

least some memorial beliefs would bear a causal relation to the truths they are about

(namely, past events). It's conceivable that intuitions are pervasively erroneous, perhaps

due to the many kinds of errors that intuitionists like Huemer mention. For example,

suppose that people are pervasively overconfident that their concepts are determinate

and complete, and thus form all kinds of (false) beliefs about the relations between the

properties that their concepts purport to represent. Nevertheless, so long as the

intuitionists' epistemic story about intuitions and their responsiveness to (mind-

independent) properties is correct, intuitions would meet the (AR).[56]

The inference from the rejection of skepticism to the acceptance of general

reliability is easy to make, but it should be resisted, both by non-skeptical realists and

by those interested in pressing a charitable epistemic objection to realism. The rejection

of skepticism is compatible with openness for massive revision of our moral beliefs.

(AR) is capable of avoiding this illicit inference, since (to speak metaphorically,) foggy

access is still access.

On the other hand, I argued above that reliability is also too weak, at least when

it is understood modally. Some beliefs may meet the reliability test simply because their

---

[56] This is not to endorse the intuitionists' epistemic story, of course. The claim is only that, if those stories could be made sense of, then the intuitionists will meet (AR). The devil is in the details of demonstrating a relationship between intuitions and mind-independent properties.

contents are necessary, which is not (directly) epistemically relevant. Non-accidentality

conditions suffer a similar problem. And if we read reliability as Bedkean reliability, the

non-naturalist can plausibly point out that it is very difficult to make sense of counter-

possible claims. (AR) avoids these problems. Though certain modal conditions can

presumably provide evidence for *access* (for most cases, *reliability* and *access* will go

together), (AR) itself is not a modal condition. (AR) only depends on some relationship

being met in the actual world.[57]

    For similar reasons, (AR) is an improvement over the *non-accidentality* condition

(M2b). Recall that the problem with the non-accidentality condition is that lucky beliefs

can meet the condition merely as a result of the metaphysical nature of the facts in

question. Since moral (and mathematical) facts are necessary truths, any method

whatsoever that results in true moral beliefs will count as non-accidental, since there is

no way those beliefs could be false—there is no possible world in which the moral

truths are different, so the beliefs are non-accidentally true. Non-accidentality alone, I

concluded, is too weak. Since (AR) is not modal, the modal status of the facts in

question won't unacceptably affect whether beliefs about a given domain will meet the

condition. Whether (AR) is met for a given set of beliefs about a particular domain rests

exclusively with the epistemic relations that the beliefs have, which is just what we

---

[57] This also makes the epistemic demand of (AR) weaker, though compatible with, Manley's (2007, Section 2) *Revised Safety*.

should want.[58] (Recall that something like this is what motivated the move from reliability to Bedkean reliability.)

6. Conclusion

It's widely thought, amongst both moral realists and antirealists, that there is something worrisome in the spirit of Benacerraf's argument, even if his statement of it relied on a contentious epistemological claim. Several philosophers have attempted a better restatement of the argument. Crucially, an adequate formulation of the argument depends on how to properly characterize the epistemological requirement we place on the justification of beliefs in a particular domain. Ideally, we want a formulation of the epistemological argument that lies between two extremes. At one extreme, the argument must be formulated in such a way as to avoid begging the question against non-naturalists of a certain sort. At the other extreme, the argument must not be formulated in such a way as to be too easily met, granting justification where it intuitively should not be granted. I have argued that, though the literature on Benacerraf-style arguments in ethics contains several ways of articulating the argument, each articulation is too extreme in at least one of the above ways. On the other hand, (AR) provides us with an expression of the epistemological requirement which does not lie at either extreme. Plugging (AR) into a Benacerraf-style argument provides us with

---

[58] This is not to say that the metaphysical status of the truths in a given domain are irrelevant to whether (AR) is met, since the metaphysical status will constrain the possible epistemic relations a believer can have to the domain in question. For example, if moral properties are causally inert, *causal access* is not an epistemic relation that can possibly be met.

an argument that is a legitimate challenge to the (non-skeptical) non-naturalist without being a dialectical non-starter:

> (M1) If non-naturalism is true, then we have knowledge of non-naturalist properties.
> (M2) If we have knowledge of non-naturalist properties, then our moral beliefs have epistemic access to the non-natural properties in the sense given in (AR).
> (M3) Our moral beliefs do not have epistemic access to the non-natural properties in the sense given in (AR).[59]
> Therefore:
> (M4) Non-naturalism is not true.

Opponents of realism should endorse this argument as stated since it provides a serious challenge to the non-naturalist moral epistemologist. Proponents of non-naturalism should endorse this statement of the argument because refuting it instead of a weaker statement of the argument ensures that their epistemic toil is as fruitful as possible.

---

[59] How can the anti-non-naturalist support (M3)? She could argue that no forms of access so far given can work, and then give an inductive argument to support (M3). Most of the debate over (M3) centers on some particular alleged form of epistemic access. I know of no "master" argument in favor of all possible forms of epistemic access that does not beg the question against the non-naturalist.

Chapter 2: Scanlon's Reflective Equilibrium and the Epistemic Access Problem

Introduction

A non-skeptical moral epistemology must tell us how we can be in principle justified in some of our moral beliefs. Call this the *Justificatory Requirement,* or JR. Depending on one's metaethical views, meeting JR could involve an appeal to intuitions,[1] narrow or wide reflective equilibrium,[2] claims of entitlement,[3] moral perception,[4] or some combination. But this is not the only condition that a theory of moral justification must meet. The moral epistemologist must also explain how it is that her preferred route(s) of justification ensures some non-accidental connection between justified moral beliefs and the moral truths. Call this the *Meta-Justificatory Requirement*, or MJR.[5]

Placing MJR as a condition on an adequate non-skeptical moral epistemology is not to single out morality as specifically epistemically dubitable. Rather, it is an adequacy condition on a theory of justification for any domain. The challenge that MJR presents for moral epistemology is especially salient for non-naturalist moral realists.[6] This is because according to non-naturalists, moral facts are not identical to any natural

---

[1] Audi (2005, 2008a), Huemer (2005, 2008, 2009) Shafer-Landau (2005).

[2] Daniels (1996), Scanlon (2014).

[3] I know of no one who defends entitlement specifically with respect to moral epistemology, but it has been used elsewhere. See, e.g. Boghossian (2000), Burge (2003), and Silins (2012).

[4] Audi (2013), Cullison (2010), Jacobson (2005), McBrayer (2010), and McGrath (2004).

[5] See, most famously, Street (2006) for a worry in this area. How best to characterize the *meta-justificatory requirement* is a question which has sprouted a small literature in itself. How the requirement is framed is important, since some meta-justifications will meet MJR only when formulated in certain ways. The arguments of this paper will not presuppose any particular way of framing MJR.

[6] Huemer (2000) disputes this claim. I cannot address those arguments here.

facts. Furthermore, for many non-naturalists, the moral properties are causally inert.[7]

Non-naturalists cannot easily extend their meta-justificatory explanation of our beliefs

about the natural world to our moral beliefs. Non-naturalists, then, must provide a way

of meeting MJR that is unique to moral facts, or at least unique to morality as well as a

few other (alleged) *a priori* domains such as mathematics and metaphysics.

In *Being Realistic about Reasons,* Thomas Scanlon attempts to meet MJR by

appealing to his particular conception of Reflective Equilibrium (RE).[8] Though reflective

equilibrium (RE) is widely endorsed as a methodology in ethics, it is infrequently

appealed to as a way to meet MJR.[9] Many who endorse RE as a normative methodology

nevertheless appeal to one of the other approaches above in order to meet MJR. In this

chapter, I argue that Scanlon's explication of RE does not provide the resources to meet

MJR. Since RE cannot meet MJR, it cannot provide a wholly adequate moral

epistemology on a realist understanding of moral facts.

I begin by discussing Scanlon's conception of RE and its use in meeting MJR. I

then consider whether any of the individual steps in RE could provide us with non-

accidental epistemic access to the moral facts, concluding that we should be skeptical

that they could. Next, I consider the possibility that the process of RE as a whole can

explain how MJR is met, even if none of its components can. I argue that this is also

[7] Oddie (2005/2009) is an exception here.
[8] Scanlon (2014), Lecture 4.
[9] Two notable exceptions are Daniels (1979) and Lillehammer (2011).

unsuccessful. Finally, I consider whether Scanlon's domain-relative minimalist metaphysics of moral facts can help to explain how RE could meet MJR. I again conclude that the answer is negative. In short, Scanlon has not shown that RE can meet MJR, whether we are minimalists or not about moral facts.

1. Scanlon's Defense of Reflective Equilibrium

In *Being Realistic about Reasons,* Thomas Scanlon addresses MJR directly. His response is two-pronged: First, Scanlon argues that it is mistaken to require some kind of causal or perceptual access to the moral realm in order to meet MJR. Second, he argues that we can provide a full epistemological story with a realist-friendly development of Rawls' notion of reflective equilibrium. I do not dispute Scanlon's first prong. I focus instead on his second claim, arguing that he is mistaken that reflective equilibrium can sufficiently meet MJR.

Since Scanlon and I agree about the first prong, I can be brief. Scanlon argues that the requirement of some causal relation to the moral facts relies on an illicit extension of the epistemology of spatiotemporally located objects into the non-spatiotemporal realm. As he says, in the case of spatiotemporally located objects, "If information is to get from them to us, how is this to happen except by their having a causal impact on our sensory surfaces?"[10] However, in the case of moral and mathematical facts, which are (on his view) not spatiotemporally located, this story is implausible: "if these [moral and

---

[10] Scanlon (2014), 70.

mathematical] facts are 'outside of space and time' the problem of explaining how information could get from them to us seems even greater than in the case of empirical truths. No causal link can bridge the gap".[11] The mistake is in thinking that the causal epistemology of spatiotemporal objects should extend to the epistemology of morality and mathematics.

However, Scanlon admits, even once we've rejected this illicit maneuver, "we are still left with the question of how we discover truths about [moral] matters."[12] To answer this question, we turn to the second prong of Scanlon's moral epistemology: reflective equilibrium (RE). For Scanlon, reflective equilibrium proceeds as follows. First, we formulate considered judgments, "judgments that seem clearly to be correct and seem so under conditions that are conducive to making good judgments of the relevant kind."[13] These considered judgments can be conceptual truths as well as substantive claims that "seem obviously true" and are "clear and apparently unproblematic."[14] Second, we attempt to "formulate general principles that would 'account for' these judgments."[15] A good general principle is one that has great explanatory power without "generating any implausible conclusions, let alone contradictions."[16] Third, if no consistent and intrinsically plausible set of general

---

[11] Scanlon (2014), 70.
[12] Scanlon (2014), 70.
[13] Scanlon (2014), 77.
[14] Scanlon (2014), 73.
[15] Scanlon (2014), 77.
[16] Scanlon (2014), 73.

principles can be found to match the considered judgments, "one must decide how to respond to the divergence between these principles and considered judgments".[17]

It's worth noting that Scanlon argues that RE also provides an adequate epistemology for mathematics, and he seems to believe that this provides some further support for the use of RE in ethics. The thought seems to be that, since it is relatively non-contentious that we have mathematical knowledge, showing that the epistemology of morality works in the same way provides some support that such an epistemology is acceptable and not mysterious. Without getting too sidetracked, I want to briefly explain why this shouldn't provide comfort to the non-naturalist moral realist. First of all, the epistemology of math is not nearly as uncontroversial as Scanlon implies. Benacerraf's epistemological worry originally applied to mathematical Platonism, and Benacerraf-style worries are alive and well amongst philosophers of math.[18] Second, even if RE is an adequate epistemology for the domain of mathematical truths, this is of no help to the normative realist addressing MJR unless mathematical Platonism is also true. Scanlon's mathematical analogy only provides support for his moral epistemology if RE gives an adequate epistemology of math *for the mathematical Platonist*. Just pointing to the way that mathematicians get knowledge of mathematical facts does nothing to establish mathematical Platonism. The proponent of the Benacerraf-style argument

---

[17] Scanlon (2014), 77.

[18] Field (1989) is still the most widely cited argument in the literature. For some recent discussions of the problem, see Kasa (2010), Liggins (2010), and Linnebo (2006).

could accept that RE provides an adequate mathematical epistemology given some constructivist, formalist, or error theoretic ontology of math. So Scanlon's mathematical analogy can only provide support for moral epistemology assuming mathematical Platonism *and* RE in mathematical epistemology. Neither are safe assumptions to make.

Of course, showing that the mathematical analogy does not help Scanlon is not to show that RE cannot meet MJR. So let's turn directly to that question. Scanlon's conception of RE is a (recursive) three step process:

Process	Epistemic Inputs
1. Formulate considered judgments "under conditions that are conducive to making good judgments of the relevant kind."	(a) Epistemologically basic.
2. Formulate general principles that would "account for" considered judgments	(b) Considered judgments
3. Are general principles intrinsically plausible and consistent with all considered judgments? If no, balance and unify conflicting judgments, and repeat step 2.	(c) Considered judgments, general principles (from step 2), and general coherence constraints.

There are two ways that this conception of RE could get us epistemic access to the moral facts. First, one of the steps in the process could alone give us access to the moral facts.[19] Second, there could be some way in which, although none of the steps by itself provided access, when the steps are combined properly, MJR is met. I consider each approach in turn.

---

[19] It is also possible that more than one of the steps gives us access, though in different ways.

Begin with the first approach—can any of the steps themselves meet MJR? I think it's relatively straightforward that steps 2 and 3 can't by themselves meet MJR. To see why, assume for the moment that considered judgments do not meet MJR—they are nothing but arbitrary biases, a result of cultural indoctrination, evolution, and/or caused by something else equally epistemically unsavory. According to step 2, we would take these epistemically arbitrary biases and try to formulate inductive generalizations and general principles. But the general principles we formulate are *completely* based on the considered judgments, which are *ex hypothesi* nothing more than biases. Generalizing from hopelessly bad data does not get us access to the correct data.

Step 3 is of no help fixing the bad data. Step 3 only pressures us to place our considered judgments and generalized principles into a fully coherent set of principles and considered judgments. This procedure can result in a more coherent set of moral beliefs, but coherence alone does not provide access.[20] If the considered judgments do not meet MJR, step 3 will only result in an arbitrary coherent set of beliefs, unresponsive to the domain of facts in question.

None of this is to discount steps 2 or 3 of Scanlon's conception of RE. They are plausible parts of a moral epistemology. But unless we have epistemic access prior to steps 2 and 3, they will be of no assistance. Garbage in, garbage out.

---

[20] BonJour (1985), Ch. 6 is helpful on this point.

All of the weight of meeting MJR, then, is on the formation of considered judgments in step 1. Scanlon actually agrees with this point, pointing out that it is nothing "that a defender of the method of reflective equilibrium should deny, or be embarrassed by."[21] So by Scanlon's own lights, if MJR is a legitimate requirement on moral epistemology, it must be met by the considered judgments. How, then, do considered judgments meet MJR? The Benacerraf-style skeptic objects that considered judgments can't provide us with epistemic access to the non-naturalist moral facts. Responding to an objection along these lines,[22] Scanlon says that the objection rests on:

> [A] misunderstanding of what it is for something to be a considered judgment. In order for something to count as a considered judgment about some subject matter it is not enough that that judgment be very confidently held. It is necessary also that it should be something that seems to me to be clearly true *when I am thinking about the matter under good conditions for arriving at judgments of the kind in question*. My belief, sitting in my armchair with no information about the conditions on the moon, that there is a rock there with my name on it, does not count as a considered judgment in the required sense, no matter how certain I may be that it is true.[23] (emphasis Scanlon's)

Considered judgments are not merely strongly held opinions. In order for a judgment to count as considered, it must be a judgment made under certain epistemically favorable conditions. What conditions, the skeptic asks? Scanlon's reply is that the conditions must be "good conditions for arriving at judgments of the kind in question."

---

[21] Scanlon (2014), 82.
[22] Scanlon is responding to the objection as presented by Kelly and McGrath (2010), Williamson (2007), pp.244-246, and Rippon (2010).
[23] Scanlon (2014), 82.

On a first pass, the Benacerraf-style skeptic may appear to be embarrassed. She demands that Scanlon give an explanation of how considered judgments provide epistemic access to moral facts. Scanlon replies that it is built into the nature of a considered judgment that it is made in the kinds of conditions that are "good" for arriving at judgments of the kind in question. And presumably, when we arrive at judgments about some domain as a result of good epistemic conditions, we'll be in a position of having epistemic access to the facts in question. Take, for example, what I'll call a "considered visual judgment". A visual judgment counts as considered just in case it is a judgment made in good conditions for arriving at visual judgments. Good conditions for making a visual judgment will involve proper lighting, being wide-awake, having one's eyes open, etc. In these conditions, I look at my desk and come to judge, based on visual experience, that there is a pencil on my desk. Since I've met the relevant conditions, my visual judgment counts as considered. Furthermore, considered visual judgments will meet MJR since the conditions on a considered visual judgment ensure that I'll have (causal) access to visible objects in my immediate environment.

Scanlon's idea, then, seems to be that when the Benacerraf-style skeptic asks how considered judgments can meet MJR, she is asking a confused question, since it's *constitutive of a considered judgment* that they provide epistemic access to the moral facts.

This is why Scanlon says that the objection rests on a "misunderstanding of what it is for something to be a considered judgment."[24]

The Benacerraf-style skeptic should be suspicious. Scanlon's response raises an immediate question: Just what *are* the good conditions for making a moral judgment, and how do *these* conditions ensure epistemic access? Without an answer to this question, Scanlon's response is stipulative rather than substantive. The skeptic can just ask a further question which is in the spirit of the Benacerraf-style challenge, namely: What reason do we have to believe that there *are* any *considered judgments*, or at least that human beings have them, if they are stipulated to be just those judgments that meet MJR?

Scanlon explicitly rejects as mysterious the thought that a quasi-perceptual idea of intuitions can provide us with a source for considered judgments. He also rejects any purported causal connection with the moral facts.[25] While these are helpful at explaining what the conditions for a considered (moral) judgment are *not*, Scanlon does not say very much about what conditions are good ones for making judgments about the domain of morality. But there are two suggestions worth discussing.

The first positive thing that Scanlon says in this context is that "one thing one needs to ask, in deciding whether something that seems true should be treated as a considered judgment, is whether it has any implausible implications or

---

[24] Scanlon (2014), 82.
[25] Scanlon (2014), 70.

presuppositions. Identifying something as a *considered* judgment involves reaching a

preliminary conclusion that this is not the case."[26] This is presumably not the only

condition on what makes a given judgment a considered judgment, but it's a start.

However, even if it is an acceptable condition on a considered judgment, it is not the

kind of condition that gets Scanlon any closer to meeting MJR. For this condition is

merely one of coherence—it only says that one necessary condition on a judgment's

counting as *considered* is that it doesn't conflict with other (considered?) judgments. In

order to know whether a particular judgment had any "implausible implications or

presuppositions", we would already need some idea of what moral claims or

presuppositions are implausible.[27]

Secondly, Scanlon says that a judgment can be rejected as not *considered* when the

general principles that would account for it are implausible. Even if correct, this

condition does not come any closer to explaining how judgments could meet MJR and

thus count as *considered*. It is merely a coherence condition, albeit one that applies to

judgments of a different level of generality.[28]

This encompasses all of Scanlon's positive comments about the conditions which

render a judgment *considered* (in the normative case). I don't think Scanlon takes himself

to have provided a complete characterization of the conditions that need to be met for a

---

[26] Scanlon (2014), 84.
[27] One exception may be judgments which lead to strict contradictions, but presumably this condition alone is not nearly strong enough to ensure a judgment counts as considered.
[28] Scanlon (2014), 84.

judgment to count as *considered*. But I'm not sure about this, since he does clearly take himself to have resolved any concerns about epistemic access.[29] In some sense he is right about this, since *considered judgments* will, by definition, meet MJR. But if this is all that Scanlon can show, it is a hollow victory, since given this conception of a *considered judgment*, it is an open and difficult question whether or not human beings ever do have such things.

2. Is Reflective Equilibrium Greater Than the Sum of its Parts?

By Scanlon's lights, in RE, the weight of epistemic access is placed primarily on the considered judgments. I've argued that either (a) considered judgments can't do the work required of them, or (b) human beings may never have considered judgments. But even if none of the three steps that make up Scanlon's conception of RE can themselves meet MJR, the defender of RE could still argue that when all of the steps are combined together, epistemic access is achieved. In other words, the whole of RE is epistemically greater than the sum of its parts. This alternative is worth briefly considering, even though I don't think it is one that Scanlon endorses.

It is relatively easy to see how RE as a whole could be in certain respects epistemically superior to any of its parts. For example, suppose that considered judgments are highly fallible, only correct 51% of the time. If this were the success rate of considered judgments, they would not by themselves be on very good epistemic

---

[29] See, e.g. Scanlon (2014), Lecture 4, Section 4.

standing—surely not enough to themselves be justified. However, even if having a considered judgment alone is not enough to be justified in a particular moral belief, it could turn out that considered judgments provide "justification enough". What I mean by this is that they may provide enough justification on their own such that the extra epistemic boost from applying the second and third coherence steps of RE may be enough, when coupled with their minimal level of justification, to provide justification full stop. This is because it is reasonable to expect the coherence steps of RE to weed out at least some of the outlier intuitions while systematizing the accurate ones, thus providing further inferential support for the (foundational) considered judgments.[30]

This is all right, so far as it goes. But even though it is possible (perhaps even common) for the entirety of a process to be epistemically superior to its parts in certain ways, it is more difficult to see how the entirety of a process could be superior to any of its parts *at meeting MJR*. Coherence measures alone can in some circumstances improve the epistemic standing of a set of beliefs, but they cannot alone provide us with epistemic access to *mind-independent* facts. The mind-independent qualification is important here, for there is one way that the process of RE could as a whole meet MJR without any of its parts doing so. It could do so if the truth-makers for moral claims was *constitutively tied* to the results of RE, properly applied. This is, in fact, the metaphysics of moral facts that some proponents of RE could give, and in such cases, I agree that

---

[30] This isn't exactly the picture developed in Mark van Roojen's (forthcoming) response to skeptical worries about intuitions, but I owe a direct debt to that paper for helping me to see this point.

MJR is met. In fact, this is how I understand Rawls to construe the relationship between

RE and the moral facts, at least in some moods.[31] Rawls interpretation aside, this is the

limiting case in which RE in its entirety can meet MJR without any of its parts

individually doing so—call such accounts *constructivist*.[32]

Scanlon once endorsed constructivism.[33] But he now explicitly rejects it, at least

as a global thesis about reasons in general.[34] When Scanlon considers the constructivist

view that for something to be a reason is for it to remain at the end of the process of RE,

he says:

> "[T]he normative status conferred on a judgment by its being in a set that is in
> reflective equilibrium depends on the quality of the decisions that are made in
> arriving at that equilibrium—decisions about what to count as a considered
> judgment at the outset and about what to modify in situations of conflict. So the
> most that could be said is that *p* is a reason for *x* to do *a* if the judgment that it is
> such a reason would be among *x*'s evaluative judgments in reflective equilibrium
> *if the judgments x made in arriving at this equilibrium were sound*. So understood,
> however, this is not a constructivist account of reasons…The process of seeking
> reflective equilibrium in one's beliefs about a subject matter is therefore not a
> characterization of the facts about that subject matter but rather a method for
> arriving at conclusions about that subject matter".[35]

Scanlon is clear, then, that he does not intend to be a constructivist, even of the sort

which constitutively ties the reason-facts to his favored method of RE. Now, of course,

one could just depart from Scanlon on this point and embrace constructivism. And, as I

---

[31] On this, see Kelly and McGrath (2010), Section 4 and Bagnoli (2011), Section 1.

[32] See Bagnoli (2011).

[33] Scanlon (1998), though Aaron Elliott has convinced me that given his other commitments in that work, he was already committed to some form of mind-independent realism. See Elliott's (ms) "Realist Contractualism".

[34] See Scanlon (2014), Lecture 4 for his extensive discussion of various kinds of constructivist theories. There may, however, be locally constructed sets of reasons, e.g. if justice-reasons are constructed out of the reasons of agents in Rawls' Original Position.

[35] Scanlon (2014), 103.

noted above, I think this would provide a way of meeting MJR. However, embracing

constructivism would be to give up *stance-independence,* a crucial feature of what makes

a particular metaethical view realist. A view counts as *stance-independent* when "*the*

moral standards that fix the moral facts are not made true by virtue of their ratification from

within any given actual or hypothetical perspective."[36] Though Scanlon doesn't use the term

"stance-independence", he does endorse it.[37] Since Scanlon's concern, and the concern

of this paper, is how the realist can meet MJR, this would be to concede the argument in

question.

3. Metaphysical and Epistemological Minimalism?[38]

Although the MJR is an epistemological requirement, its relationship to the

metaphysics of moral properties/facts should by this point be clear. The available

strategies open for a moral epistemologist to meet the MJR are highly contingent on the

particular commitments to the nature of moral properties/facts that she has. With this in

mind, I think it's important to address the particular conception of moral facts that

Scanlon proposes in the earlier lectures of *Being Realistic about Reasons*. Since Scanlon's

metaphysical commitments are somewhat unique, it's worth assessing whether there

may be epistemological strategies open to him that may not be open to more

---

[36] Shafer-Landau (2003/2005), 15, emphasis his.
[37] Scanlon (2014), 52.
[38] Scanlon disavows the labels "quietist" and "minimalist", but I'm not sure what he would prefer to have his view called instead. Wright, who Scanlon sees himself as heavily influenced by, calls himself an anti-realist about ethics. But presumably Scanlon would also disavow this label as well.

metaphysically "robust" realists.[39] I'll first give a brief overview of Scanlon's

metaphysical quietism. Then, I will argue that quietism will not be of assistance with

meeting MJR, at least if we are to understand RE in the way that Scanlon does.

Scanlon's moral metaphysics resembles in certain respects the domain-relative

metaphysics of Carnap, Wright, and Tait.[40] Scanlon argues that common metaphysical

objections to moral realism rest on the mistake of assuming that the metaphysical

standards of science should apply to all domains of inquiry. The objector claims that if

non-natural moral facts don't figure in the best scientific explanations, then they are

metaphysically superfluous, and thus should be cut from our metaphysical

commitments. This is a mistake, Scanlon argues. Instead, we should take domains of

inquiry as basic. The claims of a particular domain,

> …insofar as they do not conflict with the statements of some other domain, are
> properly settled by the standards of the domain that they are about.
> Mathematical questions, including questions about the existence of numbers and
> sets are settled by mathematical reasoning, scientific questions, including
> questions about the existence of bosons, by scientific reasoning, normative
> questions by normative reasoning, and so on.[41]

Given this domain-relativism about metaphysical questions, we shouldn't demand that

moral properties/facts figure in causal explanations, and we shouldn't be concerned

that moral properties/facts are causally inert. Instead, the standards of a particular

---

[39] Scanlon's view shares some affinities with Parfit (2011) and Skorupski (2013). By more "robust" here, I mean non-naturalist moral realists who take their commitment to moral realism to be more metaphysically weighty and in need of defense. (I hope it is clear that I don't intend to be making anything but a fast and loose distinction here.)
[40] Scanlon (2014) 19, 25, and Carnap (1950), Wright (1992, 2003), Tait (1986).
[41] Scanlon (2014), 19.

domain "typically consist, in part, of substantive principles about the domain, such as mathematical axioms, moral principles, and scientific generalizations…this reasoning is itself internal to the domains in question."[42]

Take mathematics, for example. Mathematical entities are arguably causally inert. This is a cause for concern amongst metaphysicians that Scanlon believes mistakenly hold the domain of mathematics up to the standards of science instead of the standards of mathematics. Instead, the relevant standards for good reasoning in the case of mathematics are the method of proof. If a statement in the domain of mathematics is provable (from obviously true axioms), then it is true *simpliciter*. There are no external questions to ask about the existence of mathematical objects, at least so long as the claims of the domain of math do not conflict with the claims of some other domain.

Turning now to the moral case, it appears that RE is supposed to play the role in normative reasoning that proof (along with RE) does in mathematical reasoning. (The distinctive entities that make up the normative domain are reasons or things reducible to reasons.) So the natural analogy with the mathematical domain here is that a claim about reasons is true iff it survives the completed process of reflective equilibrium. But this account conflicts with two related things that Scanlon explicitly says about the normative domain. First, he says that the truths of the normative domain are not

---

[42] Scanlon (2014), 20.

reducible to the truths of any other domain, as they would be on this account.[43] He

stresses that the difference between his view and that of sophisticated expressivists such

as Gibbard and Blackburn is that for him, the correctness of the normative truths are

independent of those commitments themselves.[44] Second, Scanlon explicitly denies this

specific account, since the process of RE can only justify *if done correctly*. He says that the

only way to fix this reductionist account of reasons is to add that *x*'s performing RE will

only lead to true beliefs "*if the judgments x made in arriving at this equilibrium were

sound.*"[45] And, as Scanlon himself points out, adding this "soundness" condition is to

give up the reductionism.

So unlike the standard minimalist account of the mathematical domain, Scanlon

does not want to identify truth in ethics with the results of any particular method, at

least without adding *substantive* constraints on that method to ensure truth. It might be

claimed that proof in mathematics also has substantive constraints, so there is no

important difference here. For example, mathematical proofs must obey certain rules of

what counts as deductively valid. But for one thing, it is unclear that these standards

are more than merely *formal* rules. But secondly, and perhaps more importantly, the

rules in the domain of mathematics are well-defined and uniform across cases. The

process of RE, if it is to be plausible at all, is all but well-defined and uniform across

---

[43] Scanlon (2014), 2.
[44] Scanlon (2014), 52.
[45] Scanlon (2014), 103.

cases. Any conflict resolution done in the process of RE is done on a piecemeal basis—any such general rule that said we must always reject high-level/mid-level principles or particular cases would result in a process of RE that is both implausible and in conflict with Scanlon's own account of it.

Scanlon must constrain the formal process of RE in order to ensure that the process has both plausible results and at least relatively widespread convergence. However, the only way to do this is with substantive restrictions on the kinds of judgments and refinements that are acceptable within the process. We saw above that a judgment only counts as *considered*, and thus only justified, insofar as it is made under conditions likely to result in true beliefs. But there is no internal method for assessing the truth of a particular judgment—a judgment's merely seeming true on reflection is not proof of its being true or counting as *considered*. Similar restrictions apply to the process of refining our considered judgments to eliminate conflict. Those refinements only count as justified insofar as they are done *correctly*, where correctness is a substantive claim, rather than one that can be determined by engaging in some internal procedure. When I am worried that a particular mathematical proof is fallacious, there is a formal procedure to determine whether or not I am correct. Not so in the moral case.

Once we have opened up this gap between the methods of reasoning in the normative domain and the truths of that domain, the skeptic now has an opening to

press for some explanation of how RE alone can meet the MJR. Of course, meeting the MJR doesn't require a guarantee that the methods of a particular domain must lead to truth in that particular domain. It only requires some account of how a method is responsive *at all* to the truths in question. If the reasons-facts are identified with the conclusion of the process of RE, it is easy to see how pursuing RE could be responsive to the reasons-truths. But once Scanlon is committed to substantive claims about how best to pursue RE, claims which are not formally discoverable from within the process of RE, we have a metaphysical gap between the reasons-facts and the process of RE which is epistemically inexplicable. What Scanlon owes us is a reason to believe that the judgments we do form on reflection are likely to be *considered judgments* more often than not and a reason to believe the refinements we make in the process of RE will by and large be the substantively correct refinements to make. He has offered us no such reasons.

4. Conclusion

Non-naturalist moral realists have long had to contend with an epistemological objection presented against their view: If the moral facts are non-natural and causally inert, then it is mysterious how we could come to have access in order to form non-accidentally true beliefs about them. Thomas Scanlon, in his recent Locke Lectures *Being Realistic about Reasons*, addresses the *Meta-Justificatory Requirement* (MJR) head-on. Scanlon believes that, when coupled with the right metaphysical picture of normative

facts, the method of reflective equilibrium can provide a moral epistemology which adequately meets MJR. I have argued that Scanlon is mistaken on this account. While RE may be enough for constructivists or anti-realists of a certain sort to meet MJR, it does not do the work that it needs to for a non-naturalist moral realist such as Scanlon in meeting MJR. Once we have substantive and stance-independent normative truths, we need some substantive story about how RE could make us responsive to those truths in the proper way. Scanlon has provided us with no such account.

Chapter 3: Why Conceptual Competence Won't Help the Non-Naturalist Epistemologist

Introduction

We have seen that an adequate non-skeptical moral epistemology must not only tell us how we can have justification, in principle, for moral beliefs. We've also seen that the moral epistemologist must further explain how her preferred route(s) of justification ensures, in at least some cases, a non-accidental connection between justified moral beliefs and the moral truths. Call this *the Meta-Justificatory Requirement, or MJR*.[1]

MJR is widely accepted as a burden for the non-naturalist—though of course there is also widespread disagreement about whether that burden can be met.[2] But there is a wrinkle in this challenge that is not widely recognized: Non-naturalists must explain not just how our moral beliefs are non-accidentally connected to some facts or other, but how our moral beliefs are connected to the robustly normative facts—that is, those facts that have genuine normative authority. Non-naturalists, or at least those I will be concerned with here, believe that there is a unique (or nearly unique) set of robustly normative properties that constitute or ground the robustly normative facts.[3]

---

[1] How best to characterize the *meta-justificatory requirement* is a question which has sprouted a small literature in itself. How the requirement is framed is important, since some meta-justifications will meet MJR only when formulated in certain ways. The arguments of this paper will not presuppose any particular way of framing MJR (though see "How to Understand the Epistemic Access Problem for Non-Naturalist Moral Realism" for my view on the issue).

[2] For arguments that it can't be met, see e.g. Street (2006), Fraser (2014), and Bedke (2014). For attempts to meet it, see e.g. Enoch (2011, Ch.7), Parfit (2011), and Vavova (2014).

[3] A small plurality of true (robustly) normative systems may be compatible with the spirit of non-naturalist realism, but there must at least be a very wide swath of candidate moral systems that do not fall into this category. From

So non-naturalists must not only ensure that we have non-accidentally true moral beliefs; they also must ensure that the content of these beliefs are robustly normative. This wrinkle is non-trivial. As I argue below, at least three recent attempts to meet MJR fail because they overlook this fact. As a result, these three moral epistemologies are inadequate, because they provide no explanation of how the epistemic relations that they propose hold between moral beliefs and facts connect us to the robustly normative properties. The lesson is clear. Non-naturalists must keep in mind their commitment to a unique set of robustly normative facts when providing an explanation of how MJR is met. They must do this not only to ensure that their account of the epistemic relation between moral beliefs and moral facts is of the right kind, but also to ensure that the relation features the right kind of relata: robustly normative facts.

I begin in section 1 by reviewing a distinction between robust and formal normativity originally found in Copp (2005) and McPherson (2011).[4] In section 2, I explain why this distinction is of importance to all non-naturalists worried about moral epistemology, regardless of the robustness of their metaphysics. In particular, they must ensure that the concepts or properties that figure in our beliefs are the robustly

---

here on out, I ignore this qualification and just refer to the "uniquely" correct moral system. I don't believe this affects the arguments given.

[4] McPherson uses the distinction en route to a metaphysical argument against quietist non-naturalism. My argument differs in being (a) epistemological, and (b) intended to be of relevance to all non-naturalists, quietist or not.

normative ones.[5] In sections 3-5, I argue that three recent attempts to meet MJR fail

because they overlook the considerations adduced in section 2. In section 3, I discuss

Huemer's (2005) *a priori* moral epistemology. In section 4, I discuss Schroeter (ms), who

provides a metasemantics for moral terms which ensures that MJR is met. In section 5, I

discuss Cuneo & Shafer-Landau's (2014) moral epistemology, which hinges on the

claim that some substantive first-order moral truths are conceptual truths. I draw some

general conclusions in section 6.

1. Robust and Formal Normativity: A Distinction

In chess, certain moves are forbidden. For example, when you're playing chess,

there's a real sense in which you (normatively) can't move your pawns backward.

Insofar as there is a set of rules, chess is formally normative. All it means to say that

chess is formally normative is to say that it's "possible to play an incorrect chess

move."[6] Similarly with other activities: They are formally normative insofar as it is

possible to perform those activities incorrectly. Formal normativity isn't that interesting

or mysterious. There is nothing special about chess, for example. "Schmess" is a game

which is "identical to chess except that in schmess one is permitted to move one's

Knight diagonally."[7] Schmess is every bit as formally normative as chess, since it is

---

[5] I don't mean here to be committing to any particular view of propositions—I use the language of a structured proposition theory (see King 2011), but only because it is most convenient in this context. I take my arguments to be compatible with any view of the nature of propositions.
[6] McPherson (2011), 232.
[7] McPherson (2011), 232.

possible to incorrectly play schmess as well. And we could indefinitely multiply examples of chess-like games which are formally normative.

Within a given formally normative activity, we can make mistakes. But what is crucial (at least for present purposes) about formal normativity is that formally normative activities exhibit *normative symmetry* amongst each other. Though chess players and schmess players may run into practical difficulties when they play each other, there shouldn't be any serious concerns about who is playing the "right" game, and who is failing to do so. It is in an important sense arbitrary whether people decide to play chess or schmess. They can make appeals to which game is more fun or fair, but these are practical questions, not questions about the single "true" way to play chess. There is nothing intrinsically authoritative about formally normative properties.

On the other hand, for the non-naturalists that I'll be concerned with here, moral properties are supposed to have some distinctive authority that (merely) formally normative properties do not.[8] One way of bringing this out is to consider how we might respond to someone who was analogous to the schmess player above. If someone claimed that there was no sense obeying the rules of morality when we could just as easily obey the rules of "schmorality", we would see her as making some kind of

---

[8] Derek Parfit puts the non-naturalist's point concisely when he claims that reductive views of reasons "claim to be describing normative reasons. But on such views…there aren't really any normative reasons…Things matter only in the sense that some people care about these things, and these concerns move people to act." Parfit (2011), Vol. I, 110. (Though Parfit talks in terms of normative reasons rather than normative properties, I take it the robust/formal distinction remains the same.)

serious mistake, or at least as engaging in a non-terminological disagreement. Morality

is, for the non-naturalist, intrinsically binding in a way that merely formally normative

activities are not. There is what we might call a normative asymmetry between morality

and schmorality that does not exist between chess and schmess (or other merely

formally normative properties). This normative asymmetry illustrates that morality

exhibits something more than merely formally normative activities: It is robustly

normative.[9]

We can make use of the robust/formal normativity distinction with respect to

different kinds of entities. Robustly normative properties are those properties (such as

being-morally-wrong) which are intrinsically binding and authoritative in the way that

merely formally normative properties (such as *being-a-legal-chess-move*) are not. Robustly

normative concepts are those concepts that are either intrinsically binding or those that

refer to robustly normative properties. Robustly normative propositions are those that

have robust properties or concepts as one or more of their constituents. Robustly

---

[9] Something's being robustly normative does *not* mean that it is completely normatively overriding. The robust normativity of morality is compatible with other robustly normative properties, e.g. standards of rationality, outweighing their force. The relevant difference here is that we cannot create robust normativity out of thin air in the way we can with formal normativity.

It's also worth noting that *robustness* is not synonymous with *categoricity* on at least some understandings of the latter, as the case of etiquette (see Foot 1972) shows. Etiquette, for Foot, is categorical in the sense that agents are criticizable from the standpoint of etiquette regardless of their desires. This is compatible with the criticism in question's being merely formal criticism. However, on a more Kantian understanding of categoricity, where a fact is categorically binding just in case it applies to rational agents as such, is even more difficult to tease apart from robustness. I stick to the language of robustness throughout as to not assume any particular understanding of categoricity is correct.

On a common conception, *reasons* are taken to be robustly normative considerations. On this understanding, we could say that a property/concept/standard is robustly normative just in case it generates or grounds genuine reasons. (Thanks to Nicole Dular and Hille Paakkunainen for discussion here.)

normative standards are those that are either intrinsically binding or those that inherit their robustness via a relation between them and some robust entity. In short, we might say that

> An entity is *robustly normative* iff it is either fundamentally intrinsically binding (in the way that formally normative entities are not), or not fully explicable without appeal to some fundamentally intrinsically binding entity.

Since I don't want to assume a particular view of which non-natural entity is fundamental, I stick where possible to the distinction between robustly and formally normative truths. However, when discussing the moral epistemologies of some philosophers below, I apply the robust and formal distinction to concepts, moral terms and properties as well. I hope that the fact that the distinction is applicable to all of these entities is relatively clear.

2. *Robust* and *Formal* Normativity: A Lesson

The formal/robust normativity distinction is not my creation.[10] However, what has not been pointed out is its relevance to moral epistemology. I now illustrate the significance of the distinction for non-naturalists attempts to meet MJR.

Suppose that the following two claims are true:

Metaphysical Non-Naturalism. There are irreducible robustly normative truths.[11]

---

[10] As noted above, it has been codified from Copp (2005) and McPherson (2011), though I am sure it has been around much longer. (See, e.g. Foot (1972).)

[11] Sometimes "robust" non-naturalism is used in opposition to "quietist" non-naturalism, where the former but not the latter are committed to some metaphysically weighty conception of moral properties or concepts. I am using "robust" in this context in a way that is compatible with quietism: "Robust" in this context only refers to the normative asymmetry between, e.g. morality and schmorality, as discussed above. Quietists such as Scanlon (2014) and Parfit (2011) are committed to Metaphysical Non-Naturalism since they do accept that there are robustly normative truths, though they don't think such truths carry any domain-independent metaphysical commitments.

Formal Realism. There are (merely) formally normative truths.[12]

I assume that Metaphysical Non-Naturalism is one of the core commitments of non-naturalism. Furthermore, setting perhaps some unorthodox views aside, most everyone, non-naturalist or not, accepts Formal Realism. So Metaphysical Non-Naturalism and Formal Realism seem like relatively safe assumptions to make within the context of doing non-naturalist epistemology.

As it's been traditionally understood, the epistemological objection to non-naturalism requires that the non-naturalist explain how the following can be true:

> *Non-Accidentality:* At least some of our justified, first-order,[13] and paradigmatically normative beliefs are non-accidentally true.

So understood, if non-naturalists can explain how *Non-Accidentality* is true, they have made some important epistemological progress. However, the robust/formal normativity distinction can help to illustrate why the truth of *Non-Accidentality* alone is not enough to vindicate a non-naturalist epistemology. Even if Metaphysical Non-Naturalism and *Non-Accidentality* are both established, it remains an open question whether our paradigmatically normative beliefs contain robustly normative contents. In order to vindicate an anti-skeptical non-naturalism, we need our paradigmatically

---

[12] Note that Formal Realism is neutral with respect to the nature of the formally normative truths, or their reducibility to other truths.

[13] By first-order I mean that the justified and paradigmatic beliefs in question are about the extension of the normative property or concept, such as that expressed by "Lying is morally wrong," or "Actions are good insofar as they increase happiness" (where this is the 'is' of predication). This is contrasted with metaconceptual beliefs such as that expressed by "*Wrongness* is a robustly normative property."

normative beliefs to be beliefs containing robustly normative propositions. An

epistemological story that only shows us how we can have non-accidentally true beliefs

of formally normative propositions would show that *Non-Accidentality* is true, but it

would not vindicate non-naturalism. This illustrates that *Non-Accidentality* isn't the full

story about what a non-naturalist epistemology requires.

 Instead, in order to defuse epistemological objections, non-naturalists must also

meet:

 Content Success At least some of our justified, first-order, and paradigmatically
 normative beliefs contain robustly normative contents.

The need for the non-naturalist to meet *Content Success* falls out of the discussion above.

Explaining how some of our first order normative beliefs could be non-accidentally

true—that is, meeting *Non-Accidentality*—involves illustrating some positive epistemic

relation between our normative beliefs and some stance-independent facts.[14] But this is

not enough for a wholly adequate non-naturalist epistemology. For all that's been said,

Non-Accidentality can be met even while a subject's normative beliefs contain merely

formally normative concepts or properties. Non-accidentally true normative beliefs are

a necessary component of a non-skeptical non-naturalist epistemology, but they are not

sufficient. What is also required is that our non-accidentally true normative beliefs are

of robustly normative propositions. In other words, *Content Success*, along with *Non-*

---

[14] Shafer-Landau (2005), 15.

Accidentality, is required for the non-skeptical non-naturalist to fully discharge her

epistemological burden.

  Non-Accidentality and *Content Success* are distinct conditions on an adequate

moral epistemology. But it is important to keep in mind that the non-naturalist must

explain how both can be met with respect to many of the very same beliefs. Showing

how some paradigmatically normative beliefs of type *A* are non-accidentally true and

showing how some paradigmatically normative beliefs of type *B* are of robustly

normative propositions is not enough. It must further be shown that there is some

overlap between *A* and *B*. In other words:

> *Overlap.* At least some of our justified, first-order, and paradigmatically
> normative beliefs are *both* non-accidentally true and contain robustly normative
> contents.

As I argue below, I think *Content Success*, and its relationship to *Overlap*, have been

overlooked. Both are crucial to a proper understanding of how non-naturalists must

respond to the epistemological objection. This interplay between *Non-Accidentality* and

Content Success, as spelled out in *Overlap*, is an important one to keep in mind.

  The robust/formal distinction, then, is important for non-naturalist

epistemologists to keep in mind, regardless of their metaphysical commitments. This is

because *Content Success* must be met by anyone committed to Metaphysical Non-

Naturalism, which is a core commitment of non-naturalists, quietist and non-quietist

alike. I don't mean to claim that no non-naturalist epistemologies on offer can

adequately meet *Content Success*; rather, I mean only to claim that *Content Success* is a largely overlooked requirement on a non-naturalist moral epistemology. I turn now to illustrating, by way of some examples, the importance of recognizing *Content Success* as a crucial non-naturalist epistemological claim to be met.

3. Huemer's *A Priori* Moral Epistemology

Michael Huemer (2005) has recently attempted to explain our epistemic access to non-natural moral facts via a general theory of *a priori* knowledge. Huemer identifies properties with universals, sometimes aligning himself with Armstrong's (1980) immanent realism.[15] Furthermore, as with most non-naturalists, he accepts that "the moral realm is causally inert."[16] And though he admits that moral properties, if they exist, would be metaphysically strange, he doesn't see this as evidence against their existence.[17] Finally, he accepts the received view that moral properties supervene on natural properties, though they do so in a non-reductive way.[18] An evaluative belief that *p*, on Huemer's view, has some initial positive epistemic status so long as it is grounded in an intellectual seeming that *p*.[19]

---

[15] Huemer (2005) 124-125 contains an argument for realism about universals. Huemer (ms., Sect. 4.2) endorses immanent realism, citing Armstrong 1980. As we'll see below, Huemer accepts (what I call) *Plenitude*, the view that there is a universal for every possible consistent and determinate concept. This is puzzling, since there can plausibly be concepts for non-instantiated properties, which immanent realists by and large are forced to deny. It is unclear what Huemer would say about this potential inconsistency, but this issue will not be relevant for the arguments below, so I set it aside.

[16] Huemer (2005), 122.

[17] Huemer (2005), 200-201.

[18] Huemer (2005), Ch. 8.

[19] Of course, Huemer also allows for inferentially justified evaluative beliefs, e.g. an evaluative belief that *q* based off of an intellectual seeming that *p* and a prior belief that if *p*, then *q*. I stick to the more straightforward case in the text to avoid unnecessary complications.

Huemer's explanation of *a priori* knowledge involves four important claims:

> **(U)** (Mind-independent) universals exist necessarily.
> **(C)** Having an adequate (consistent, clear, determinate) concept *constitutes* the grasping of a (mind-independent) universal.
> **(R)** "Having an adequate grasp of a universal puts one in a position to see that it has certain properties and/or relationships to other universals that you adequately grasp."[20]
> **(A)** "All *a priori* knowledge is, or derives from, knowledge of the properties and relations of universals."[21]

Taken together, these claims can explain how moral beliefs meet *Non-Accidentality*. Here is how Huemer explains it:

> Notice…that the defining characteristics of an adequate grasp are *intrinsic*— consistency, clarity, and determinacy belong to the nature of the concept in itself, as opposed to depending on the relationships between the concept and something else. So the intrinsic characteristics of a concept sometimes are sufficient for its constituting an adequate understanding of the nature of a universal…Therefore, in some cases—namely, when one's intuitions are caused (only) by clear, consistent, and determinate understanding—the internal process by which one forms beliefs guarantees their truth.[22]

Since forming an adequate[23] concept involves meeting purely *intrinsic* criteria, and meeting those very criteria *constitutes* adequately grasping a (mind-independent, causally inert) universal, Huemer has given us an explanation for how we can form non-accidentally true beliefs about universals without standing in some causal or quasi-perceptual relation to them: All we need to do is assess whether the concepts which

---

[20] Huemer (2005), 125.
[21] Huemer (2005), 126.
[22] Huemer (2005), 126.
[23] Note that "adequacy" here is a technical term in Huemer's system.

constitute the belief in question meet the (reflectively accessible) conditions, and we've

got non-accidental truth.

An immediate question arises with respect to (C). Even if my concept *C* is

consistent, clear, and determinate, that doesn't yet explain how *C* is guaranteed to have

a corresponding (mind-independent) universal. Isn't it possible that *C* just doesn't track

any real mind-independent property? As it turns out, for Huemer, the answer is no.

This is because "[t]here is no possibility of one's [concept] failing to refer to anything

(universals are plentiful in this sense, and their existence is necessary)."[24] As long as our

concept is adequate,[25] we can be sure that it refers to a real mind-independent universal.

Huemer's way of meeting the MJR, then, relies on a fifth important claim, which I'll call

Plenitude:

> *Plenitude:* For every possible *adequate* (consistent, clear, and determinate) concept,
> there is a corresponding mind-independent universal.

With this explicit statement of *Plenitude,* we have a complete story about how *a priori*

moral beliefs can explain *Non-Accidentality*.[26] Since we are capable of reflecting on our

moral concepts, adjusting them to become more adequate, and since we know that these

moral concepts have corresponding mind-independent universals (of which our

---

[24] Huemer (2005), 126.

[25] Note that "adequacy" here is a technical term in Huemer's theory, described above.

[26] Assuming, at least, that some of these beliefs based on an adequate understanding are first order normative beliefs. I grant Huemer this point in what follows. (Thanks to Hille Paakkunainen for pointing this out to me.)

adequate concepts *constitute* the grasping of them), we can have non-accidental *a priori* beliefs about at least some of the truths in which our moral concepts feature.[27]

We've seen Huemer's general story about how *Non-Accidentality* will be met. With an eye to *Content Success*, let's turn to a particular case of how epistemic access to moral facts works on Huemer's account. Consider a straightforwardly "good" case of how a particular ethical intuition can guarantee the truth of its contents. Suppose Lucy is reflecting on her concepts of BADNESS and LYING, and on that basis it seems to her that:

(L) LYING is intrinsically BAD.

Furthermore, let's assume that her concepts of LYING and BADNESS are consistent, clear, and determinate.[28] Then, by *Plenitude*, they both refer to mind-independent universals. Let's call the universals they refer to *being-a-lie* and *badness*, respectively. Since Lucy's concepts are consistent, clear, and determinate, and she formed an intuition based on reflecting on those concepts, the following is guaranteed to be true:

(L*) *Lying* is intrinsically *bad*.

So we have gone from a claim about Lucy's mind-dependent concepts to a mind-independent fact about universals. And presumably, since she has met the conditions

---

[27] Interestingly, a similar *Plenitude*-based strategy has been advanced as a way for the mathematical Platonist to meet MJR, developed independently by Balaguer (1998) and Linsky & Zalta (1995).

[28] It should be noted that the conditions on an adequate grasp of a universal are pretty difficult to meet. But I aim to grant Huemer as much as possible here.

that guarantee the truth of my belief that (L*), she *knows* it, or at least is justified in believing it.

So far, so good. But now consider Carol. Carol is also engaged in ethical inquiry, and also reflects on her concepts of BADNESS and LYING. Carol's concepts are *also* clear, consistent, and determinate. But reflecting on these concepts, she has the intuition that:

(-L) LYING is not intrinsically BAD.

By *Plenitude*, then, for similar reasons, the following is guaranteed to be true:

(-L*) *Lying* is not intrinsically *bad*.

Since (-L*) is guaranteed to be true, and Carol has formed her belief in a way that guarantees it to be true, she also *knows* it, or at least is justified in believing it.

It should be obvious at this point that something has gone wrong. (L*) and (-L*) are incompatible with each other, so they can't both be "guaranteed" to be true—at least one of them has to be false. And I certainly don't mean to suggest that Huemer is unwittingly committed to contradictory ethical claims both being true. Rather, there is a subtle but informative mistake in the reasoning given above. Lucy and Carol both grasp their respective concepts clearly, consistently, and determinately. This means that, so long as they are being careful, they won't misapply those concepts. Their concepts have a clear and determinate extension, and they will both be able to classify things as picked out by their concepts or not. But, by *Plenitude*, so long as their concepts have clear and

determinate extensions, there will be properties that their concepts track. So they do

both form true beliefs about the extension of their concepts. However, as the difference

between (L) and (-L) shows, the extension of their concepts of LYING and BADNESS

differ. And this doesn't show us that one of them is mistaken—after all, their respective

concepts are all fully adequate—rather, it shows us that they *have different concepts*

altogether.

The right way to classify the case, then, is to distinguish Lucy's concepts from

Carol's concepts. Call Lucy's concept BADNESS$_L$ and Carol's concept BADNESS$_C$.[29]

With this distinction in mind, we can more precisely characterize their respective ethical

intuitions:

(L$_L$) LYING is intrinsically BAD$_L$.

(-L$_C$) LYING is not intrinsically BAD$_C$.

Given *Plenitude*, both Lucy's and Carol's concepts correspond to mind-independent

properties. And since their concepts are adequately grasped, they will both have

knowledge of the corresponding mind-independent facts. However, since their

concepts differ, the corresponding properties will differ as well. The knowledge that

Lucy has gained from her ethical intuition is not best characterized as (L*), but as:

(L$_L$*) *Lying* is intrinsically *bad$_L$*.

---

[29] It may also turn out that Carol and Lucy have subtly different concepts of LYING, as well. I set this aside for simplicity, as it doesn't affect the strength of the argument.

Furthermore, the knowledge that Carol has gained from her ethical intuition is not best

characterized as (-L*), but as:

    (-L$_C$*) *Lying* is not intrinsically *bad$_C$*.

Once we get clear about what Lucy and Carol believe based on their intuitions, then, we

see that they do not contradict each other after all. This is analogous to the merely

formally normative chess case discussed above. The schmess players truly believe that

knights can move diagonally, while chess players do not. But this is because the

schmess players have beliefs about *being-forbidden-in-schmess*, while chess players have

beliefs about *being-forbidden-in-chess*. If the schmess players and the chess players were

in an argument about this, they would be merely talking past each other.[30]

    We can multiply cases like Lucy and Carol's indefinitely. There are many

different fully adequate possible moral concepts, each of which differs about the

extension of "bad". These concepts are all compatible with one another, so long as we

are clear about whether we are talking about *bad$_1$*, *bad$_2$*, etc. However, unlike in the chess

case, or the case of any formally normative concept, there is a further question in the

moral case: Which concepts adequately characterize the robustly normative properties?

By *Plenitude*, any adequate moral concept will correspond to some mind-independent

---

[30] What if the chess and schmess players had some higher-order agreement that there is a unique fact about the single right way to play a chess-like game? Analogously, what if Lucy and Carol has some higher-order agreement that, whatever the case may be, their concepts aimed at the robustly normative, and thus their disagreement is not merely terminological? Because Schroeter (ms) clearly has machinery available to her to make this move, I address this kind of response below, in section 4. But I think what I say there would apply, *mutadis mutandis* to a similar response made in Huemer's defense.

property or other. Huemer has shown that (e.g.) Lucy's beliefs meet *Non-Accidentality*. Her belief (L_L*) is a non-accidentally true first order normative belief. But he has provided no reason to accept that Lucy's normative concepts actually refer to the robustly normative properties, and thus no reason to accept that her normative beliefs were of the robustly normative facts. In other words, though Huemer has shown how *Non-Accidentality* can be met, he has not shown how *Content Success* can be explained, and so has not provided a wholly adequate non-naturalist epistemology. To paraphrase Street, so long as it hasn't been shown that there is a relation between the moral concepts we happen to have and the normative facts, the appeal to conceptual competency offers no way, in the absence of an incredible coincidence, of showing how our moral beliefs could meet *Content Success*.[31],[32]

Notice that the success of the objection does not rest on the possibility of divergence between two agents such as Lucy and Carol. The case of Lucy and Carol is merely meant to help illustrate the problem. The problem is that a wholly adequate non-naturalist epistemology requires more than the mere having of normative beliefs that track some facts non-accidentally. It must be that those beliefs track the robustly normative facts non-accidentally. Even convergence amongst agents would not be enough to explain how our normative beliefs are responsive to the robustly normative

---

[31] Paraphrasing Street (2006), 124-125.

[32] Couldn't it also be part of Lucy's consistent, clear, and determinate concept of BADNESS that it refers to a robustly normative property? For reasons given in section 4, I think the answer is no.

facts. For all that's been said, and for all we know, our beliefs in such a situation would be converging on some merely formally normative facts.[33]

4. Schroeter's Metasemantic Solution

Laura Schroeter (ms) suggests that the non-naturalist can meet *Non-Accidentality* by appeal to a Chalmersian metasemantics for normative terms. To understand how metasemantics are supposed to assist in resolving epistemological worries for non-naturalist realism, it's first important to remind ourselves what it is that a metasemantic theory is supposed to do. In Schroeter's words, the job of metasemantics is to "specify what it is about someone's psychological states that determines which objects, kinds, or properties her representational states pick out—or more generally, their semantic contents."[34] Such a theory, Schroeter contends, must be able to play "two central theoretical roles".[35] First, the theory must be able to ascribe contents to subjects' beliefs, so that we can assess the success and accuracy of subjects' beliefs and reasoning. Second, the theory should ascribe semantic contents that provide us with the ability to predict and explain subjects' behaviors.

Schroeter uses these two desiderata to provide support for the first of two crucial principles in her overall metasemantic solution:

---

[33] The objection is also not that the privileging of some properties over any others as robustly normative is incompatible with *Plenitude*. There is no reason the Plenitudinous realist couldn't privilege some properties as robustly normative. Rather, the problem is an epistemic one.

[34] Schroeter (ms), 7.

[35] Schroeter (ms), 8.

> *Minimal Charity* (MC) The correct semantic interpretation of the concept expressed by a term must make at least *some* of the beliefs involving that concept come out true.[36]

MC won't alone provide much by way of a response to epistemological objections to non-naturalism. However, it can at least explain why our moral concepts don't refer to properties like *being-a-triangle* or *being-red-or-green*. MC doesn't yet establish that our moral concepts refer to non-natural properties, since MC is compatible with moral concepts' referring to natural properties such as those to which Cornell Realists or other naturalists would have them refer.[37] As Schroeter herself emphasizes, MC can only ensure that "if reference to Realist [Non-Naturalist] moral properties succeeds, then some of our moral beliefs must be true."[38]

A second principle that Schroeter appeals to, both more controversial and able to do more work, involves an idea originally found in Chalmers and Jackson (2001):

> *Ideal Accessibility* (IA) The correct semantic interpretation of the concept expressed by a term 'X' must make the subject's ideal, empirically-informed beliefs about what it takes to be X in any possible world come out true.[39]

The empirical information must be given in an "ideal base-level descriptive vocabulary", rather than in natural language, on pain of circularity. I'll grant for the

---

[36] Schroeter (ms), 7.
[37] Boyd (1988), Brink (1989).
[38] Schroeter (ms), 9.
[39] The wording is Schroeter's (ms, 10). She credits Chalmers (2004) with the revision of Chalmers & Jackson's (2001) original proposal. A further minor complication is whether or not to include phenomenal information—the question turns on whether or not one thinks the phenomenal facts are reducible to the empirical—but settling this dispute is not important for present purposes.

purposes of this paper that such a vocabulary is in principle possible.[40] The basic idea

behind (IA) is that "the reference of a subject's concepts is determined *by the subject's*

own fully informed, ideally reflective verdicts about the reference."[41]

For present purposes, we need not worry whether IA is a correct metasemantic

principle for all terms. In order to resolve epistemological objections to non-naturalism,

it only need be true of *normative* terms. Suppose, as Schroeter argues, that it is. Recall

that *Non-Accidentality*, the epistemological challenge to non-naturalism, demands some

explanation as to how our actual normative beliefs can track the moral properties, given

that those properties are causally inert. IA tells us that a subject's moral beliefs under

ideal conditions are guaranteed to be true, given that those beliefs themselves

determine the referent properties of the moral concepts that figure in them.[42] So our

ideal selves' moral beliefs can straightforwardly explain *Non-Accidentality*. However,

this also means that our actual moral beliefs, or at least some subset of them, are non-

accidentally true as well. This is because

> [I]deal beliefs about X have to be justifiable on the basis of the subject's actual
> beliefs about X—on pain of changing the topic. So Ideal Accessibility constrains
> the subject's actual beliefs indirectly: all of the subject's actual beliefs that would
> figure in justifying her ideal, fully-informed verdicts about what it takes to be X
> must come out true.[43]

---

[40] For much more discussion on such a vocabulary, see Chalmers (2012).

[41] Schroeter (ms), 11, emphasis mine.

[42] I harbor some doubts about whether our ideal beliefs will be guaranteed to be true given IA, but Schroeter makes this claim (ms, 13), and I grant it for the sake of argument.

[43] Schroeter (ms), 11.

Some of our actual moral beliefs must survive the idealization process, lest those beliefs

not really be *our* ideal moral beliefs. But this means that some of our actual moral beliefs

are non-accidentally *guaranteed* to be true, though we may not know which ones.[44] Not

only that, but we also have reason to believe, given IA, that further reflection and

information will help weed out more and more false beliefs. Our giving beliefs up on

receiving more information is very good evidence that they wouldn't survive ideal

reflection, and thus are not true. IA helps establish that some of our actual moral beliefs

must be true, and also that further reflection and information is likely to result in more

true moral beliefs. These results are enough to meet any reasonable meta-justificatory

demand on the non-naturalist, and thus the non-naturalist who endorses IA can meet

Non-Accidentality.

However, once we keep in mind the distinction between formal and robust

normativity, Schroeter's view faces a dilemma, each horn of which is incompatible with

the truth of either *Non-Accidentality* or *Content Success*. Recall that in order for *Non-*

Accidentality, *Content Success*, and *Overlap* to be met, it must turn out that (a) at least

some of our substantive first-order normative beliefs are true, (b) the properties picked

out by our moral terms are the robustly normative properties, and (c), both (a) and (b)

are met with respect to at least some of the very same beliefs. We've seen above how IA

---

[44] I harbor some doubts about this claim. What if all of our moral beliefs are based in false empirical beliefs? But Schroeter may say that even in this case we have true moral beliefs that survive idealization—they just turn out to be conditional beliefs like "If the empirical facts are such-and-such, then the moral facts are such-and-such." This may be correct, but it's not obvious. Nevertheless, I grant Schroeter this point for the sake of argument.

could ensure that the first condition is met: Since our ideal beliefs are a result of some function on our actual beliefs coupled with empirical information, there is some quasi-constitutive connection between our first order normative beliefs and the properties that figure in them.

Return to Lucy's judgment that:

(L) Lying is intrinsically bad.

Suppose for the sake of argument that (L) is one of Lucy's beliefs that *would* survive idealization.[45] Call the property that figures in Lucy's "bad" thoughts *badness_L*. Recall that Carol believes that:

(-L) Lying is not intrinsically bad.

And suppose for the sake of argument that (-L) is one of Carol's beliefs that would survive her own idealization.[46] Call the property that figures in Carol's "bad" thoughts *badness_C*.

As with Huemer, Schroeter ascribes to something like Plenitude, so there is no worry about whether Lucy or Carol's terms fail to refer.[47] However, it should be clear that the same problem arises here for Schroeter's view that arose for Huemer's view.

---

[45] Recall that on Huemer's view, so long as Lucy's belief meet his internalistically specified conditions, then her belief is guaranteed to be correct. On Schroeter's view, Lucy's first order normative judgment may yet be mistaken, since there may be some empirical information currently unbeknownst to her that would lead to her denying (L).

[46] Schroeter may deny that this divergence under idealization between Lucy and Carol can occur, as least so long as they are members of the same community. I return to this point shortly.

[47] Schroeter (ms), 12.

Schroeter has given a metasemantic story which can explain how to avoid a skeptical

moral epistemology, but she has done so in a way which makes meeting *Content Success*

extremely difficult. Lucy and Carol both meet *Non-Accidentality* with respect to each of

their sets of normative beliefs. Some set of Lucy's actual normative beliefs *S* figures in

(partially) determining her idealized normative beliefs. By IA, the members of *S* will be

non-accidentally true. However, if the normative beliefs that make up *S* don't match the

first order *robust* normative truths, Lucy's normative beliefs won't meet *Content Success*.

In other words, while some of Lucy's normative beliefs would be non-accidentally true,

it would be extremely unlikely for those beliefs to be of the robustly normative facts.

But Carol's epistemic situation, and our own, is identical to Lucy's. Schroeter has

provided an epistemology that can explain *Non-Accidentality* but not *Content Success*.

However, the argument provided so far against Schroeter's view is incomplete.

This is because IA actually *does* provide a way for our normative terms to pick out the

robustly normative properties (thus meeting *Content Success*):

> David Kaplan could fix the reference of the name 'Newman1' by stipulating that
> it refers to the first child born in the 22nd Century: if the world cooperates by
> producing exactly one child who fits the description, then the name rigidly
> designates the child. No causal interaction with the child is required in order to
> fix reference…And of course, the child in question is not metaphysically
> constituted by Kaplan's beliefs about it. Similarly, we can stipulatively fix
> reference to a Realist property with an appropriate definite description.[48]

---

[48] Schroeter (ms), 10. Reference to Kaplan (1967) omitted.

The idea here is that moral terms may contain, as part of their metasemantics, a reference-fixing description that they pick out the robustly normative properties, whatever they turn out to be. If I stipulate that my term "wrong" picks out a robustly normative 'not-to-be-doneness' property if it refers at all, then, so long as there is some such property that fits the description, "wrong" will refer to that property. (Compare stipulating that "Julius" refers to the inventor of the zip, if it refers at all.[49] In the moral case the stipulation is that "wrong" refers to the robustly normative not-to-be-doneness property, if it refers at all.)

Now Schroeter does not actually believe that normative terms have explicitly stipulated referents, but the semantics of normative terms may nevertheless contain implicit reference-fixing descriptions which do similar work. Whether this is so depends on whether a subject would consider robust normativity an essential feature of "what it takes to be", e.g., "wrong", after the idealization procedure given in IA. And it's at least plausible that such a judgment would survive the idealization procedure. This would ensure that our normative terms did pick out the robustly normative properties, thus meeting *Content Success*.

So far so good. But remember that what is crucial is for *Non-Accidentality* and *Content Success* to be met with respect to the very same set of beliefs (*Overlap*). I claim

---

[49] Evans (1982), 25.

that the explanation just given of how *Content Success* is met undermines the account's

ability to explain *Non-Accidentality*.

Return to Lucy and Carol. "Badness" refers to a different property for Lucy and

Carol, as their idealized divergence about lying shows.[50] Suppose that Lucy and Carol

realize this. There are two ways they might react:

> *Relativity*: Lucy and Carol accept that "badness" is not univocal, that they refer to
> different properties, and that this is an acceptable end to normative inquiry.

> *Univocal Robustness:* Lucy and Carol both take "badness" to refer to a robustly
> normative property. Since they accept that there is only one robustly normative
> 'not-to-be-doneness'-type property (if any), they take their dispute about lying to
> be a genuine disagreement not settled by empirical information alone (otherwise
> idealization would have resolved the dispute).

Suppose that *Relativity* occurs. In that case, *Non-Accidentality* is met, since Lucy and

Carol both have beliefs which are non-accidentally true. However, it is met at the cost of

giving up on meeting *Content Success*, assuming that there is a unique set of robustly

normative moral properties. By the non-naturalist's lights, there is a unique robustly

normative property of *badness*. But since we could multiply cases like Lucy and Carol's

indefinitely, if we accept relativity then we can't preserve this fundamental

commitment of non-naturalists. So *Relativity* would result in an explanation of how our

normative beliefs could be non-accidentally true, but it would fail to explain how our

moral beliefs were of robustly normative facts—in fact, it would give us positive reason

---

[50] Of course it is also possible that the divergence is a result of their notions of "lying", but I set this possibility aside for simplicity.

to believe that most of our moral terms do not refer to the robustly normative

properties, since the possibility that our terms picked them out would be no better than

chance, for all Schroeter has said.

Alternatively, suppose *Univocal Robustness* occurs. Then it looks like moral terms,

at least for Lucy and Carol, operate more like "Newman1", picking out the robustly

normative properties via something like an implicit reference-fixing description. *Unique*

Robustness, then, is a situation in which *Content Success* is clearly met. However, notice

that this comes at the cost of undermining any reason to believe *Non-Accidentality*. If

picking out the robustly normative properties is a non-negotiable feature of the

semantics of our moral terms, then we'd need some independent reason for accepting

that our first-order normative beliefs—idealized or not—were true, just as Kaplan

would need some further evidence before accepting any first-order beliefs about

Newman1. For example, Kaplan might believe that Newman1 will become a

philosopher, that he will have black hair, and that he will be an excellent knitter. These

beliefs do no work toward fixing the content of "Newman1", and none of them could be

non-accidentally correct (unless Kaplan has some special prescience of the distant

future). They may, for all that we know, all be false and thus rejected under

idealization. Now these rejected beliefs will, in the case of Newman1, be replaced by

true beliefs under idealization, because part of the idealization process will involve

giving Kaplan all the base level descriptive information about the world he needs to

deduce all the facts about Newman1. However, in the normative case, base level

descriptive information will not correct for fundamentally mistaken normative beliefs.

By the non-naturalist's lights, non-normative information alone won't entail robustly

normative facts (in any non-trivial way). Non-normative information alone, for the

same reason, won't help correct an agent's wildly mistaken first order normative

beliefs.

To illustrate this, notice again that we could multiply cases like Lucy and Carol's

indefinitely. Take any first-order normative belief that you have, *B*, which would

survive idealization. There's no principled reason why there couldn't be some other

English speaker who denied *B*, and whose denial of *B* survived idealization as well.[51]

However, suppose both of you accept that moral terms, first and foremost, refer to the

robustly normative properties, whatever they might be. When faced with this

disagreement, then, there's no way to adjudicate the dispute, because neither party has

any reason to believe that it is their first-order belief that is true. But *B* is an arbitrary

moral belief, so this argument generalizes.

Notice, as above, that although instances of idealized disagreement help to

illustrate the problem, the objection is not an argument from disagreement, nor does it

---

[51] Why couldn't speakers pack more into their (implicit) reference fixing description? For example, perhaps "wrong" refers, if it refers at all, to the robustly normative 'not-to-be-done' property that forbids lying, murder, etc.? The problem here is that once we pack in more information about the referent into the reference fixing description, we have less and less reason to think that any such referent exists. Even if there is some sound argument for non-naturalism, this wouldn't guarantee the existence of robustly normative properties with particular first-order recommendations. (Thanks to David Sobel for pressing me on this issue.)

require any disagreement, actual or possible, to succeed. The problem is that if our actual normative beliefs do not latch onto the robustly normative facts, then no amount of non-normative idealization will help to meet *Content Success*. And alternatively, if *Content Success* is met via building robust normativity into the semantics of normative terms, then no amount of non-normative idealization will help to meet *Non-Accidentality* with respect to our first order normative beliefs. The lesson, then, is that no metasemantic view alone can tell us what actions, states of affairs, or persons to which the robustly normative properties apply. At best, it could tell us the referents of our normative terms.

5. Cuneo and Shafer-Landau on "Moral Fixed Points"

In an ambitious and exceptional recent paper, Terence Cuneo and Russ Shafer-Landau's crucial claim is that "there is a battery of substantive moral propositions…that are also nonnaturalistic conceptual truths."[52] They go on to argue that if this claim is accepted, it can do a wide-ranging amount of metaphysical and epistemological work for the non-naturalist. (Here I am only worried about the alleged epistemological work that can be done. I remain agnostic on their claim that the view can do metaphysical work.) For example, the following claims are, they argue, excellent candidates for being non-naturalistic conceptual truths:

> It is pro tanto wrong to engage in the recreational slaughter of a fellow person.
> It is pro tanto wrong to humiliate others simply for pleasure.

---

[52] Cuneo & Shafer-Landau (2014, 2).

If acting justly is costless, then, *ceteris paribus*, one should act justly.[53]

First-order non-naturalistic conceptual truths such as these they call the *moral fixed points*. Any normative system which failed to endorse the moral fixed points would thereby not count as morality. This plausible idea can be seen as an extension of Philippa Foot's point that 'it is wrong to run around trees right handed' isn't just false, but nonsensical.[54] Cuneo & Shafer-Landau argue that the moral concepts can provide substantive constraints on moral theorizing. These constraints arise from the moral fixed points.

So far so good. But in light of the arguments I've given above, a natural question arises: What reason do we have to believe that our moral concepts pick out anything mind-independent? Cuneo & Shafer-Landau's answer to this question involves appealing to what they call the "traditional view" of concepts.[55] According to the traditional view, concepts have three important features. First, concepts—rather than the objects or propositions themselves—are the constituents of propositions.[56] Second, concepts serve as the referential devices that "enable thinkers to refer to things such as objects and properties."[57] Finally, according to the traditional view, concepts are

---

[53] Cuneo & Shafer-Landau (2014, 7). They are explicit that, while these are excellent candidates for being conceptual truths, the specific examples they choose are irrelevant to the more general claim that some substantive first order normative claims are conceptual truths.

[54] See Foot (2001, 7).

[55] Cuneo & Shafer-Landau (2014, 11). As they freely admit, the traditional view is heavily indebted to Frege.

[56] Cuneo & Shafer-Landau (2014, 11).

[57] Cuneo & Shafer-Landau (2014, 12).

"abstract, sharable, mind-independent ways of thinking about objects and their properties. As such, they are very much objective, 'out there' sorts of things, extra-mental items whose existence does not depend on our employing them in thought or language."[58] Furthermore, as mind-independent entities, concepts have essences which underlie conceptual truths such as the moral fixed points.[59] It is this last feature of the traditional view that is crucial for answering the question given above. Our question was how our individual moral representations can pick out anything mind-independent. And the answer, once the traditional view is granted, is simple: The concepts themselves are mind-independent, so once we have a competent grasp of a concept (through a mental representation), we already have a competent grasp of the essence of a mind independent entity.

We are now in a position to see how Cuneo & Shafer-Landau can use the traditional view along with the moral fixed points to explain *Non-Accidentality*.[60] Take some moral fixed point *m*. Now suppose I am conceptually competent with respect to the concepts (normative and otherwise) that figure in the proposition *m*. Suppose, furthermore, that upon reflection I accept *m*. My belief will be non-accidentally true: I formed it in light of my conceptual competence, and my conceptual competence

---

[58] Cuneo & Shafer-Landau (2014, 11).
[59] Cuneo & Shafer-Landau (2014, 13).
[60] I should be clear that this is my best reconstruction of what I think they would say. They don't directly address *Non-Accidentality* in their paper, but they address enough related epistemic issues that I am confident that something like this is what they would say. (See Cuneo & Shafer-Landau (2014), esp. Sections 3 and 4.)

ensures that, if I don't make some mental mistake, my belief is true.[61] So my belief that

m is not just true, but will be true non-accidentally. And such an explanation will extend

to the beliefs of all conceptually competent agents when their beliefs are of moral fixed

points.

Cuneo & Shafer-Landau's view can explain *Non-Accidentality*, at least with

respect to the moral fixed points. But can these beliefs simultaneously explain *Content*

Success? In other words, is it the case that our moral conceptual competency is

competency of the robustly normative concepts? This could occur in one of two ways,

depending on the rest of the details of one's normative metaphysics. First, it could be

that concepts, as mind-independent entities, can themselves be robustly normative.

Second, it could be that our moral concepts are robustly normative in that their

referents are robustly normative properties. Both seem to be viable options given

Cuneo & Shafer-Landau's metaphysical claims about the moral fixed points. The

argument I give below does not depend on settling this question, but I will assume the

first approach in what follows. The first approach strikes me as the more epistemically

promising, because separating the robustly normative property from the essence of the

concepts themselves only opens up another potential epistemic gap between the

---

[61] Notice that the epistemological explanation of the moral fixed points will closely resemble the epistemology of conceptual truths more generally.

believer and the fact (about the robust normativity of her concepts) at issue. However, everything said in what follows will apply to either approach one favors.

As we've seen above, on Cuneo & Shafer-Landau's view, the truth-makers for a moral fixed point are the essences of the concepts that make up the fixed point in question. No further "worldly fact" is needed to ground its truth. This is what constitutes the fact that the moral fixed points are conceptual truths. But, importantly, neither are conceptual truths merely analytic truths. Analytic truths are sentences which are true in virtue of the meaning of the words that compose them. Analyticity, as Cuneo & Shafer-Landau understand it, is a linguistic phenomenon, not a metaphysical one. Conceptual truths, on the other hand, are more metaphysically robust in that they involve mind-independent essences.[62] Since concepts have essences, and, as in the moral fixed points, concepts' essences bear relations to each other, it seems like a genuine possibility that moral concepts could themselves be robustly normative, on Cuneo & Shafer-Landau's view. If this were correct, then no further relation between the moral concepts and any non-natural properties would need to hold for robust normativity to get in the picture. The moral concepts would just have robust normativity built into their essences.

---

[62] Plausibly, if words have their meaning essentially, then analyticity may be metaphysically robust. Though they don't say this explicitly, I suspect that Cuneo & Shafer-Landau would claim that words do not have essential meanings, and that any intuition that they do is based in conflating words with the concepts that words refer to. I assume they are correct about this for the purposes of the paper.

We can grant that the traditional view of concepts is correct and that moral fixed points are true in virtue of the essences of the concepts that constitute them. Furthermore, we can even grant that there are concepts which have robust normativity built into their conceptual essence. These three assumptions can clearly support an explanation of *Non-Accidentality*, but what about *Content Success*? It might seem as though *Content Success* can be explained as well, once we've granted that robust normativity can be built into a concept's essence. However, by now the problem is a familiar one. Showing that we have substantive conceptual moral knowledge and that some concepts are robustly normative is not enough: We need to further show that *our* moral concepts are the robustly normative ones. For again, there could be any number of alternative sets of broadly speaking normative concepts which deny some or all of the moral fixed points while affirming others. There are, for example, the *schmoral* concepts, and with them, the *schmoral* fixed points, which are conceptual truths grounded in the essences of the schmoral concepts. (It doesn't matter whether any actual (or possible) agent grasps the schmoral concepts or believes the schmoral fixed points. Since the schmoral concepts are mind-independent entities, the conceptual truths will exist either way.) We could enumerate such non-moral but broadly speaking normative sets of concepts, fixed points, and essences indefinitely. What we need is some reason to believe that our moral concepts, as opposed to any of these other sets,

are the robustly normative ones. This is what explaining *Content Success* requires, and it

is difficult to see how it would fall out of Cuneo & Shafer-Landau's view.

Interestingly, Cuneo & Shafer-Landau acknowledge a related point, asking "Why

think that we have reason to pledge our allegiance to *this* normative system, rather than

another—call it *schmorality*—that fails to incorporate…the moral fixed points?"[63] (I take

it that they mean "reason" here in the robustly normative sense.) In response to this

question, they concede:

> It's an excellent question, but one that we don't propose to answer here…this
> question is a perennial worry for *all* forms of moral realism…And while
> regarding some substantive moral norms as a species of conceptual truth might
> not specifically aid us in explaining the reason-giving power of moral facts-,
> neither does it make our version of realism any less apt to offer such an
> explanation, whatever it may be.[64]

In itself, this is a fair enough concession: A single paper can't defend every controversial

piece of non-naturalist metaphysics. However, this concession has a crucial

epistemological upshot. We have no explanation—by Cuneo & Shafer-Landau's own

lights—as to why we should think our moral concepts are the robustly normative

concepts. Furthermore, there are any number of alternative systems of normative

concepts, each from our epistemic standpoint equally likely to be the robustly

normative concepts. So, for all that's been shown so far, we have more reason than not,

---

[63] Cuneo & Shafer-Landau (2014), 8-9.
[64] Cuneo & Shafer-Landau (2014), 9. They furthermore cite Shafer-Landau (2009) as an attempted answer to this question. I discuss Shafer-Landau (2009) briefly below.

probabilistically speaking, to think that our moral concepts are not the robustly normative concepts. This would mean that without further argument, Cuneo & Shafer-Landau (2014) will have shown that (many of) our beliefs in the moral fixed points are non-accidentally true, but of no more normative interest than our beliefs about the rules guiding how one can move one's pawn in chess.[65]

The upshot at present is just an instance of the more general lesson of this paper: An explanation of *Non-Accidentality* does not wholly resolve our epistemic worries. Any explanation of *Non-Accidentality* must be integrated with a solution to *Content Success*, and this is a non-trivial task. As with Huemer and Schroeter, Cuneo & Shafer-Landau have provided a solution to the former, but one whose ability to be coupled with an explanation of *Content Success* is unclear at best.

6. Conclusion

A core commitment of non-naturalist moral realism is that there is some unique or nearly unique set of robustly normative facts. Moral norms, along with rational norms, share a distinctively binding normative authority that other conventional norms—such as the norms of chess—do not. Non-naturalists also separately acknowledge that, given the irreducible and non-causal nature of normative facts, some

---

[65] Interestingly, Cuneo & Shafer-Landau reference an earlier Shafer-Landau article (2009) for a defense of the claim that there are categorical reasons. As noted above, categoricity is not the same as robust normativity, but they do seem to be closely related. However, whatever strengths Shafer-Landau's argument has in the categoricity case, it won't solve the present problem, because it relies on substantive first order judgments of the very sort which are in question.

explanation must be given as to how at least some of our first-order justified moral beliefs are non-accidentally true. This metajustificatory requirement is traditionally considered independently of the robustly normative nature and uniqueness of the normative facts. The arguments above have attempted to demonstrate that this is a mistake. Establishing that our normative beliefs are non-accidentally true—that is, establishing *Non-Accidentality* —does not establish that the robustly normative properties or concepts feature in those beliefs—that is, *Content Success*. Without further argument, we have no reason to believe that our normative beliefs, even if justified, pick out the robustly normative rather than some merely formally normative properties or concepts.

I have tried to defend the importance of integrating *Non-Accidentality* and *Content Success* by way of three case studies in moral epistemology: Huemer (2005), Schroeter (ms), and Cuneo & Shafer-Landau (2014). Each of these approaches to meeting the metajustificatory requirement, while importantly distinct from each other, appeal to conceptual or semantic competence to attempt to resolve the epistemological worries. And each is subject to the same concerns related to the connection between *Non-Accidentality* and *Content Success*. At best, these views do not explain *Content Success*. At worst, the explanations given for *Non-Accidentality* undermine any chance at also explaining *Content Success*. The lesson here is not that we should give up on non-naturalist epistemology (though that is a possible response). Rather, the lesson is that

non-naturalist epistemologists must pay close attention to the relationship between their explanation as to how our moral beliefs can be non-accidentally true and their explanation as to how our moral beliefs refer to the robustly normative. An adequate non-naturalist epistemology must show substantial overlap between the set of justified moral beliefs and the set of moral beliefs that pick out the robustly normative. For the non-naturalist, *Non-Accidentality* and *Content Success* need to stand or fall together. And while this may not be an impossible task, it is certainly non-trivial.

Chapter 4: Moral Perception and the Contents of Experience

The view that it is possible to gain moral knowledge by directly perceiving moral properties is an unpopular position in moral epistemology. Instead, moral epistemologists have appealed to the use of intuitions, "rational insight", or some other *a priori* access as providing substantive moral knowledge.[1] However, some authors have recently defended moral perception as a source of moral knowledge.[2] Although these defenses have their virtues, most are lacking in at least one of two ways.[3] First, rather than giving a positive argument that moral perception actually occurs, these authors focus on responding to objections against the very possibility of moral perception. Second, even where positive accounts are given, many fail to engage the recent literature in the epistemology of perception and how moral perception may fit (or not fit) with recent accounts of the contents of experience.[4]

---

[1] This kind of view has a long tradition, of course. See Shafer-Landau (2003), part IV, Audi (2004), Zamulinski (2007), Parfit (2011) Part II Ch.32, and Enoch (2011) Ch. 7, for recent defenses of the intuition/*a priori* access view. Singer (2005) differentiates "rational" from emotional intuitions, arguing that the former, but not the latter, are justificatory. See also Greene (2007).

[2] McGrath (2004), Chappell (2008), Väyrynen (2008), Audi (2010, 2013), Cullison (2010a), McBrayer (2010a, 2010b), and Cowan (2013a, 2013b) .

[3] An important exception here is Cowan (2013b), who provides a qualified defense of moral perception while engaging the philosophy of perception literature. The present paper provides a distinct positive argument for moral perception.

[4] Väyrynen (2008, Section IV) defends *a posteriori* ethical intuitionism from a particular objection from Sturgeon (2002), but notes that his reply would fail "if we could not perceive ethical properties as being instantiated" (p.500). Väyrynen admits that his response to Sturgeon is weak insofar as "establishing that we can perceive ethical properties as being instantiated is a tall order." (ibid.) This paper can be seen as attempting to meet that burden.

McBrayer (2010a, 2010b) is also engaged with the philosophy of perception literature, but only uses it to respond to objections to moral perception, not to provide a positive argument for it.

I will give a positive argument in favor of the perception of moral properties by drawing on a recent debate in the philosophy of perception. Using a *contrast argument*, I argue that moral properties can be part of the contents of experience. In section I, I explain the structure of a contrast argument and provide an example of a (non-moral) contrast argument that can be found in the philosophy of perception literature. Sections II-IV contain my contrast argument in favor of moral perception. In section II, I propose and defend a case involving a contrast in phenomenal character between two individuals perceiving the same scene. In section III, I provide some preliminary reasons for explaining the contrast by appealing to the perception of moral properties by one but not the other individual. Since contrast arguments are inferences to the best explanation, I spend section IV considering and rejecting alternative explanations for the contrast. I conclude that the best explanation for the contrast is that moral properties are capable of being part of the contents of experience.

Before getting into the arguments, I should make explicit three background assumptions I will not defend here. First, I assume that some version of *non-skeptical moral realism* is true, where by this I mean that (a) there are at least some moral properties whose existence does not depend on the mental states of human beings—at least not in any interesting sense, and (b) at least some human beings have at least some

non-trivial moral knowledge.[5] [6] Second, I assume that at least some moral properties

supervene on natural properties, where identity is one kind of supervenience.[7] Third, I

assume what Susanna Siegel calls the "Content View"—the view that "experiences have

contents, where contents are a kind of condition under which experiences are accurate,

similar in many ways to the truth-conditions of beliefs."[8]

1. Contrast Arguments and the Contents of Experience

Proponents of the Content View tend to agree that at least some properties such

as colors, shapes, and states of motion can be part of the contents of experience. Call

these properties *low-level properties*. There is much dispute over whether more

complicated properties can be part of the contents of experience as well. Call these

properties *high-level properties*. There are many different kinds of high-level properties:

Natural kind properties, artifact kind properties, the properties of causal chains, etc.

According to *liberals*, at least some high-level properties can be part of the contents of

---

[5] How exactly to characterize realism is controversial. Rather than getting bogged down in those debates, my definition of realism here is intentionally vague. I don't believe how we cash out realism has much importance to the argument that follows. In particular, I don't mean to take a stand on the naturalism/non-naturalism dispute. (With one exception: Shafer-Landau (2003), who carves up naturalism/non-naturalism in terms of epistemic access, and endorses non-naturalism. Since his non-naturalism is defined in terms of *a priori* access, moral perception is conceptually impossible for non-naturalism as he defines it.)
Isn't MP ruled out by the view that the moral properties are causally inert? I think that McBrayer (2010a, Sect. IV) has shown that this natural thought is mistaken, and that MP is compatible with causally inert moral properties, but I cannot defend that here.

[6] Given that the broader project of this dissertation is defending the possibility of moral knowledge against an argument for moral skepticism, one may wonder how this assumption is fair in this context. The project of this chapter is to establish the in principle possibility of moral perception as a source for moral knowledge. Later chapters (especially chapter 7) address the more difficult question of whether we should think the required conditions are met.

[7] I don't believe that the arguments of this paper hinge on any particular conception of supervenience.

[8] Siegel (2010), p.4.

experience. According to *conservatives*, high-level properties cannot be part of the contents of experience;[9] rather, the existence of such properties from our perception of low-level properties is only concluded as a result of some further (possibly unconscious) cognitive act.[10] The difference between low-level and high-level properties is probably best represented as a continuum, rather than as a binary. Nonetheless, it's helpful for bookkeeping purposes to roughly divide them up, so I use these categories throughout despite this qualification.

The thesis I intend to defend below is:

Moral Perception (MP): At least some moral properties can be part of the contents of perceptual experience.

I will take it as uncontroversial that moral properties would be high-level properties, so if MP is true, then liberalism is true. As Robert Cowan points out, this is why, in assessing MP, we should be careful not to "[prejudge] the outcome of an ongoing debate in the philosophy of perception as to what sorts of things can come to be represented in experience".[11] However, we should also keep in mind that liberalism does not commit one to the view that any and all high-level properties can be represented in experience. A defense of the perception of high-level properties can proceed on a case-by-case basis.

---

[9] I take this liberal/conservative terminology from Bayne (2009).

[10] This could result from an unconscious inference, or, as a referee pointed out to me, our perception could cause an intuition (which wouldn't itself count as a perceptual state). There are other possibilities here, but for present purposes what is important is that the high-level properties are not part of the contents of experience.

[11] Cowan (2013b), 12.

There are two views in moral epistemology that entail MP. First, according to

Perceptual Intuitionism, moral properties can be part of the contents of experience in one

of the traditional sense modalities (sight, sound, etc.).[12] On the other hand, according to

Affectual Intuitionism, (a) emotional experiences are a mode of perception (the

perceptual theory of emotions), and (b) moral properties can be part of the contents of

emotional experience.[13] I wish to remain neutral for the purposes of the arguments

below between these two views. As will be seen, I believe that emotions play an integral

role in moral perception. However, the view that emotions play an integral role in

moral perception is in principle compatible not just with *Affectual Intuitionism* (for

obvious reasons), but also with *Perceptual Intuitionism*. This is because it may turn out

that affective states cognitively penetrate the perceptual processing of the traditional

sensory modalities such that they alter the contents of classically perceptual

experiences. So, for example, a visual experience, while remaining strictly *visual*, may

contain moral properties as part of its contents as a result of affective penetration. Of

course, whether affective states penetrate classically perceptual processing is at least

partially an empirical question which I cannot address here.[14] All I am claiming here is

---

[12] Perceptual Intuitionism is defended by McBrayer (2010a, 2010b), and on one reading of Audi (2013).

[13] This view has been recently defended by Roeser (2011), who sees herself as developing the ideas of Thomas Reid. Cowan (2013, section 4) also discusses this view (and does not outright reject it).

[14] Cognitive penetration here, as elsewhere, also raises interesting and difficult epistemological issues which are outside of the scope of the present chapter.

that even though I take emotional states to play an integral role in moral perception, the arguments of this paper remain neutral between *Affectual* and *Perceptual Intuitionism*.

Though both *Affectual* and *Perceptual Intuitionism* entail MP, the entailment does not go in the other direction. This is because according to both of these versions of intuitionism, moral properties contained within the contents of experience can serve the role of justifying epistemically basic moral beliefs, which MP does not entail. The claim that an experience *E* which contains a property *F* can justify epistemically basic beliefs about *F* is plausible, but it is non-trivial in the case of high-level properties such as moral properties. Things are not epistemically straightforward here because moral properties may figure in experience only as a result of the cognitive penetration of prior moral beliefs.[15] If this were the case, the justification of moral judgments formed on the basis of moral perception would depend at least in part on the status of the prior moral beliefs. While I do think that *Intuitionism* can be defended despite these concerns, I intend the argument of this paper to be neutral on the relationship between MP and its broader epistemological consequences. I only mean to claim here that the truth of MP is an important piece of the moral epistemological puzzle.

One typical defense of the presence of some high-level property in the contents of experience proceeds by way of a contrast argument of the type recently defended by

---

[15] Vayrynen (2008) and Cowan (2013) both contain excellent discussions of this issue.

Susanna Siegel and others.[16] Siegel argues for liberalism using what she calls the *Method*

of Phenomenal Contrast. The method involves considering two very similar but

phenomenally distinct experiences. We first consider an experience in which it is

plausible that the high-level property could be part of the contents of that experience.

Then we consider an experience that (a) uncontroversially does *not* represent the high-

level property, but (b) is in other respects as similar to the first experience as possible.

Now we have two similar but uncontroversially distinct phenomenal states. The

method now proceeds by inference to the best explanation. There must be some

explanation for the difference between the two experiences. The liberal regarding a

particular property contends that the best explanation is that one experience represents

the high-level property while the other does not. A conservative about the property in

question must provide some superior alternative explanation as to why the two

experiences are distinct. This could involve, for example, appealing to differences in

low-level properties of the two experiences, an appeal to a difference in cognitive

phenomenology, or perhaps a denial that the phenomenal difference is a

representational difference. Since the method of phenomenal contrast will always

involve an inference to the best explanation, any particular application of the method is

---

[16] Siegel (2006, 2010, forthcoming), Bayne (2009). For an auditory application of a contrast argument, see O'Callaghan (2011). Cowan (2013b, 5n16) mentions contrast arguments in the context of a discussion of moral perception, but does not himself provide such an argument.

bound to be contentious. This is perhaps to be expected, since it's hard to imagine what

a plausible deductive argument for or against liberalism would look like.

It may be helpful to consider an example. Here is Siegel:

> Suppose you have never seen a pine tree before and are hired to cut down all the pine trees in a grove containing trees of many different sorts. Someone points out to you which trees are pine trees. Some weeks pass, and your disposition to distinguish the pine trees from the others improves. Eventually, you can spot the pine trees immediately: they become visually salient to you…Gaining this recognitional disposition is reflected in a phenomenological difference between the visual experiences had before and those had after the recognitional disposition was fully developed.[17]

The target property in this example is the kind property of being-a-pine-tree. Our two

experiences are someone looking at the trees before and after developing the ability to

recognize them. There is allegedly a phenomenal difference between experiencing a

pine tree before and after someone can recognize them. We have a phenomenal contrast

in need of an explanation. Siegel goes on to argue that the best explanation for this

phenomenal contrast is that one's coming to be able to recognize pine trees gives them a

new ability to represent pine trees in the contents of experience.[18] In the course of her

book, Siegel gives similar contrast arguments about recognizing other kind properties,

semantic properties, and causation.

---

[17] Siegel (2010, p.100). Peacocke (1992, Ch.3) uses a similar example.

[18] Siegel (2010, Ch.4) considers and rejects other possible explanations of the contrast. I take no stand on whether she is correct about the best explanation in this particular case.

The method of formulating contrast arguments should be accepted by liberals and conservatives alike. This is because so long as the proponent of a contrast argument can provide sufficient evidence that a phenomenal contrast exists between two experiential states, there is a difference for which some explanation needs to be provided. But contrast arguments do not in principle favor liberal explanations over conservative-compatible explanations, so they are not unacceptably question-begging. However, it should also be noted that since contrast arguments trade in appeals to better and worse explanations, their conclusions will rest on a number of different considerations, some of which are bound to be disputable. Nevertheless, I think that they can help us make progress even if they do not compel all liberals or all conservatives in any particular instance.

2. A Contrast Case for (Moral) *Badness*

Now we can turn to MP. In order to test MP, we need a contrast case. Perceiving moral properties, if possible, plausibly involves some kind of recognitional disposition.[19] However, perceiving moral properties is not straightforwardly analogous to learning a second language or learning how to distinguish pine trees from fir trees. For many, the development of our ability to notice moral features in our environment is subtle enough that it will be difficult to provide a clear and uncontroversial example of a phenomenological contrast in an individual before and after developing a

---

[19] See Siegel (2010, Ch. 4).

recognitional capacity for some moral feature or property. This is not to suggest that we could not develop moral capacities which would make us more phenomenologically sensitive to the moral features of certain situations. In fact, I think such cases are possible.[20] However, these cases would probably all be too contentious to provide the strong explanatory burden that a contrast argument places on us. So I propose to set this kind of cross-temporal intra-subjective contrast case aside for the purposes of establishing or rejecting MP.

Instead, I propose that we consider two separate individuals perceiving the same situation. Just to illustrate that there is no in principle problem with a contrast argument of this sort, consider a brief example, two individuals taking an Ishihara Test of Color Blindness. In an Ishihara Test for red-green color blindness, individuals look at a circle made up of much smaller red and green dots of various shades. The dots of one of the colors make up a number.[21] An individual with red-green color blindness will not be able to distinguish the red from the green dots in order to read the number. The circle will look like it is made up of similarly colored dots in no particular pattern. A color-sighted individual will be able to distinguish the red from the green dots and thus be

---

[20] For example, understanding the mechanisms and pervasiveness of subtle sexual harassment may help someone be able to perceive moral properties of an instance of harassment that the same well-intentioned but uninformed individual could not. And such a case may involve a phenomenological contrast in the experiences of the individual before and after learning about sexual harassment. (Fricker's (2006) discussion of the concept of sexual harassment may be helpful here.) I thank Steve Swartzer and an anonymous referee for pressing me to explain why I don't think intra-subjective cases are the best kind for a moral contrast argument. I also thank Steve Swartzer for the example of the perception of subtle sexual harassment, and an anonymous referee for the possible connection with Fricker (2006).

[21] For some examples, and to take the test yourself, see http://www.toledo-bend.com/colorblind/Ishihara.asp.

able to read the number. It is overwhelmingly plausible that these individuals, while looking at the circle of dots, will be in phenomenally distinct states.[22] Just as in intra-subjective cases, it appears that some explanation of the phenomenal contrast is required in cases like these as well—at least where we can clearly establish, using empirical data, a phenomenal contrast between two individuals.

I turn now to the moral contrast case which I will use to defend MP. In this case, one individual is a normally functioning adult human being. The other individual is what I'll call an emotionally empathic dysfunctional individual (EEDI). EEDIs are individuals who have a fully functioning "theory of mind"—that is, they are capable of inferring the mental states of others based on behavioral, vocal, and contextual evidence—but who nonetheless lack affective empathy in the sense that they fail to have "an emotional response to another individual that is congruent with the other's emotional reaction."[23] For present purposes, all that is required is that EEDIs have an "inability to feel empathy with the victim" of a particular distressing experience, one common trait associated with psychopathy.[24] It may be helpful to think of EEDIs as psychopaths; in fact, one popular theory of psychopathy is just that psychopaths *are* EEDIs.[25] I refrain from using the term "psychopath" in my contrast case because there is

---

[22]Explanations of the Ishihara Test even sometimes provide examples of what the circles look like to individuals of various sorts. See, for example, http://en.wikipedia.org/wiki/Ishihara_test (Accessed June 21st , 2013).
[23] Blair (2006, p.4).
[24] Hare (1985).
[25] Blair (2006), Hare (1991), Arnett (1997).

no consensus in the field of psychology about the essential features of psychopathy, and I don't want to delve into those debates here.[26] What is of importance for generating my contrast case is only that there really are EEDIs, not whether they are to be identified with psychopathy or some subset thereof.

Consider two individuals coming across the situation described in Gilbert Harman's famous chapter on moral observation.[27] They each "round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it." Both Norma, the normally functioning adult human being, and Pathos, the EEDI, come across the same scene.[28] There is good reason to believe that Norma and Pathos are in phenomenally different states when perceiving the scene. EEDIs tend to not respond to distress cues, whereas for normally functioning individuals, distress cues can invoke powerful phenomenological states. For example, in one study, adults who scored highly on the *Hare Psychopathy Checklist—Revised* were shown images of sad faces as their skin conductance response (SCR) was tested. SCR is correlated with physiological arousal, and thus with emotional intensity more generally. The individuals who scored highly on the *Psychopathy Checklist* showed heavily reduced SCRs in comparison to the control group.[29] Similar results have been shown for children with psychopathic tendencies.[30]

---

[26] Skeem et. al. (2011).
[27] Harman (1977, 8).
[28] Thanks to an anonymous referee for recommending these heuristically helpful names.
[29] Blair et. al. (1997)
[30] Blair (1999).

Those who score high on the *Psychopathy Checklist* have also been shown to not differ in their "startle responses" (blinking and twitching) when surprised by an acoustic stimulus regardless of whether they are looking at a pleasant image, such as a smiling face, or an unpleasant image, such as a mutilated animal. A control group was shown to have a higher startle response when viewing an unpleasant image.[31] These differences, as well as others, appear to be explained by some kind of amygdala dysfunction in EEDIs.[32]

A difference in SCRs, heart rate, and startle responses does not deductively prove that EEDIs and normally functioning individuals are in states with a differing phenomenal character. These measurements only directly show that EEDIs and normally functioning individuals differ physiologically. But to demand such a proof would be asking for a solution to the problem of other minds. What we have is the best empirical evidence currently available, and it all seems to point in the direction of a phenomenal contrast in cases like that of Norma and Pathos above. For most people, witnessing a cat being lit on fire invokes a very strong phenomenological response. But for EEDIs, it appears that no such strong phenomenological response is present. So we have a phenomenal contrast in need of explanation.[33]

---

[31] Patrick et. al. (1993).

[32] Blair (2006), pp.9-13.

[33] An anonymous referee points out that the difference between Norma and Pathos appears at first blush to be emotional. While it is true that the contrast almost certainly has some relationship to emotion, what is unclear is whether the contrast is explained in terms of the affect, or whether affect is merely causally related to the state (belief, desire, perception, etc.) that the explains the contrast. See Sect 4.

3. MP as the Explanation for the Contrast

The phenomenal contrast between Norma and Pathos's perceptions of the burning cat needs an explanation. I contend that the best explanation is a difference in the perception of moral properties—Norma's experience represents the cat's burning as *bad*, whereas Pathos's doesn't. Norma's experiential state meets three conditions that standardly count as sufficient for a states' representing some property *F*. First, assuming non-skeptical moral realism, Norma has developed a disposition to be in this particular experiential state which more or less reliably tracks *badness*. Second, Norma's relevantly associated phenomenology is counterfactually correlated with *badness* (or at least a particular type of badness) in her local environment. Finally, Norma is disposed to form moral beliefs based on experiential states of this kind. In the contrast case in question, we can safely assume that Norma will spontaneously come to form the belief that the cat's suffering is *bad*. So not only are Norma's experiential states historically correlated with *badness* in her local environment, there is also a link between experiential states of this kind and Norma's forming beliefs about *badness* in her local environment. It would appear then that we have some good preliminary reasons in favor of MP as the best explanation of the contrast in question. This hypothesis would also help to explain how Norma can more or less track and respond to *badness* in her local environment without explicit, conscious rule following, while Pathos, as an EEDI, cannot do so. To be sure, the experiential feature that Norma has that Pathos does not is probably emotional in

some sense. But I don't think this should rule out its being perceptual, for reasons I mentioned above (Sect. 1) and discuss in more detail below (Sect. 4.4).

4. Are There Better Alternative Explanations?

I have proposed an explanation for the Pathos and Norma contrast case. However, in order to fully assess this explanation, it is necessary to consider alternative explanations of the contrast and see why they are unsatisfactory. I think there are four alternative possibilities: (a) the contrast is due to a non-representational difference in phenomenology, (b) the contrast is part of their cognitive phenomenology, rather than their perceptual phenomenology, (c) the contrast is due to their representing different non-moral high-level properties, or (d) the contrast is because of a difference in their representations of internal bodily processes. I'll address each of these in turn.

4.1. A Non-Representational Difference?

The first alternative explanation is that the contrast is best explained by some non-representational phenomenal difference. According to this explanation, despite the fact that Norma and Pathos are in states with different phenomenal character, the contents of their experiences are the same. Norma's experience has some *raw feel*, some purely qualitative non-representational character, which Pathos's experience lacks. I have no knock-down argument against this first alternative, but I can see at least two serious drawbacks. First, this explanation relies on the denial of *intentionalism*, roughly, the view that phenomenal character is identical to or supervenes on representational

content. Intentionalism is a serious and popular view, and it would be ill-advised to reject it only in order to deny MP—especially in the absence of a principled reason. Second, the defender of this alternative explanation must provide some independent reason to believe that Norma and Pathos's phenomenal contrast *in this particular case* is non-representational. Given that Norma and Pathos's difference in experiential states could correspond to differing dispositions to judge the cat-burning as bad, we have some positive reason to believe that the phenomenological difference corresponds to either a difference in perceptual content or in cognitive content, broadly construed. If the thought is that the phenomenal contrast is due to differing background cognitive states, then this response just collapses into the second alternative discussed below.

One motivation for this alternative that may avoid the objections just discussed arises from the combination of (a) the fact that something distinctively affective seems to underlie the phenomenal contrast, and (b) the view that affect is a good candidate for a *raw feeling,* even if intentionalism is true of most other phenomenological states. While I agree that (a) is overwhelmingly plausible, there is good reason to reject (b). Though the idea that affect is a raw feeling is a view that has been defended historically,[34] most philosophers and psychologists reject this view of affect because it cannot account for the fact that affective states have intentional objects.[35] To be sure, current orthodoxy

---

[34] Most famously by William James (1884).

[35] See, for example, DeSousa (2013, Sect. 2), Fox (2008, Ch.2). Goldie (2003, 2009) defends a view according to which affective states are essentially feelings, but allows that feelings can themselves be intentionally directed—which renders them more than mere raw feelings.

could be mistaken here, but we'd need some argument to this effect in order to consider the raw feels explanation a viable alternative.

4.2. A Difference in Cognitive Phenomenology?

The second alternative proposal is that the difference is one of cognitive phenomenology. The relationship between cognitive states and phenomenology is a contentious one, but for the purposes of this paper I will grant that propositional attitudes have associated phenomenological characters.[36] There are a few possible candidate differences between Pathos and Norma's cognitive phenomenology. One difference is in the *beliefs* that Pathos and Norma have. A second possible difference is in the *desires* that they have. A third possible difference is in the *intuitions* that they have (or don't have). I'll address each of these in turn.

Plausibly, Norma will form something like the judgment "*That* is bad" while observing the scene, while Pathos will not form such a judgment. So, one non-perceptual difference between Norma and Pathos is that they have different beliefs about the scene in front of them. If occurrent cognitive states can have an associated phenomenal character, then this difference in cognitive states could lead to a difference in phenomenal character. Couldn't this difference explain the phenomenal contrast?

---

[36] The literature on cognitive phenomenology has exploded in the past decade, but arguably the most influential paper is Horgan and Tienson (2002).

This explanation will not do. Although EEDIs are unable to feel empathy for a victim in distress, they are capable of learning and applying moral norms.[37] Imagine that Pathos has been taught that lighting cats on fire is bad. Further, imagine that Pathos has much practice in the lab at applying such norms swiftly and accurately. When Pathos sees the scene, he will be able to quickly judge that *"That* is bad." In short, Norma and Pathos can have the same judgments about the scene, and thus, if the cognitive phenomenology explanation of the contrast is correct, they should have the same cognitive phenomenology. However, we can be confident that the phenomenal contrast would remain despite the sameness of cognitive phenomenology. Thus, the explanation of the phenomenal contrast in terms of different beliefs is unsuccessful.[38]

Now turn to the possible difference in the desires that Norma and Pathos have. Norma probably has some desire that the cat's suffering end, while Pathos plausibly lacks such a desire.[39] This difference in cognitive phenomenology could explain the contrast as well. But the explanation depends on two questionable assumptions. First,

---

[37] Blair (1995).

[38] One possible objection: It could be claimed that Pathos doesn't really *understand* his judgment that "*That* is bad" in the way that Norma understands it, and thus his judgment differs in content from Norma's. But, the objection goes, if Pathos's propositional attitude really did have the same content as Norma's, then the phenomenal contrast would disappear. Although I cannot rule out this possible explanation, without some positive independent reason to think that Pathos's judgment can't share the same content as Norma's, it seems to be special pleading on the part of the opponent of MP. We could stipulate that Pathos can track badness (at the level of belief) as reliably as Norma, albeit through rule application, thus avoiding this possible explanation. So although this is a possible explanation for the phenomenal contrast, it doesn't seem to be the one that most straightforwardly explains the data.

Alternatively, it could be that Pathos and Norma's cognitive phenomenology differs, though their contents are the same. But then we would need some further explanation for *this* fact.

[39] I thank David Sobel and Aaron Elliott for pressing me on this point.

as above, Pathos very well could desire that the cat's suffering end. Perhaps he has been trained up to believe that he will be paid $500 every time he stops a cat from suffering. He then could have a strong desire to stop the cat's suffering, just as Norma does. But it's implausible to think that this would make the phenomenal contrast disappear.

It could be objected that in the case as just described, Norma and Pathos do *not* share the same desire. Norma desires that the cat stop suffering as an end whereas Pathos only desires it as a means. They are motivated by different goals, so their desires must differ in content. But rather than getting bogged down here, I turn to the second questionable assumption of the current alternative explanation. This alternative explanation assumes that Norma has a desire that Pathos does not, namely the desire that the cat's suffering end. However, Norma need not have this desire. Suppose that Norma knows that these kids lit the cat on fire for a reason—there is an evil mastermind just out of sight who will shoot the children if the cat is not burned alive. Norma is confident that there is no open possibility in which both the cat and the children are able to live. As a result, Norma most certainly does not desire that the cat's suffering end, because this would lead to the children's deaths. Yet we would expect the phenomenal contrast to remain.

Perhaps the defender of this alternative explanation would argue that Norma does have a desire that Pathos does not have, but that it is more fine-grained than I have been assuming. Maybe she has a desire that the cat not have to suffer in order to save

the children. And this may very well be a desire that Pathos does not have. I agree that Norma may very well have this desire, and that Pathos probably lacks it. However, recall that the challenge here isn't just to find some relevant difference between Norma and Pathos's desires, but to find one that more plausibly explains the phenomenal contrast than MP. And it seems implausible that this desire could plausibly explain differences between SCRs, heart rate, startle responses, etc. between Norma and Pathos better than MP. It is more likely that the fine grained desires that Norma has that Pathos does not are cognitively downstream from the immediate phenomenology of their respective experiences.

A third possible difference is in the intuitions that Norma and Pathos have. Plausibly, Norma has the intuition that "*That* is bad", while Pathos does not. If intuitions are distinct from other representational states such as beliefs and perceptions, we have another possible explanation for the contrast. And on at least one understanding of "cognitive", intuitions will count as cognitive phenomenology since they are propositionally structured, unlike experiences.[40] Unfortunately, in order to settle the plausibility of this explanation, we would need to settle the extremely contentious debate about the nature of intuitions. Instead of attempting that daunting task, let me just briefly canvas a few of the answers on offer and consider how they

---

[40] Whether or not intuitions count as "cognitive" is not of present importance. If the reader does not prefer this terminology, she can just consider this as a wholly distinct alternative explanation to MP.

would fare as potential explanations of the phenomenal contrast in the case of Pathos and Norma.[41]

On one view, intuitions are just beliefs.[42] If this is the correct view of intuitions, then this alternative explanation of the contrast just collapses into the belief explanation that I have argued is unsatisfactory above. On a second view, intuitions are not beliefs, but are dispositions to believe.[43] If this is the correct view of intuitions, then intuitions will not provide an adequate explanation of the contrast either. To see why, consider again the possibility that Pathos has been taught that lighting cats on fire is bad. Then he *will* be disposed to form the belief that "*That* is bad" when he witnesses the burning cat, just as Norma will. So they both share the disposition in question. But as noted above, it is implausible to expect the relevant phenomenological contrast to go away just because Pathos has been taught to apply a general rule.

The last account of intuitions I'll consider is the view that intuitions are *sui generis* representational states.[44] According to this view, intuitions are the mental states commonly picked out by "seems" language. Since intuitions are *sui generis* states, it is difficult to provide an in depth account of their nature, other than to give examples to

---

[41] On a fourth kind of view defended by Sabine Roeser (2011), ethical intuitions are perceptual emotional states. I hope it is clear enough that such a view is committed to something like MP, and so does not present an alternative to the view defended in this paper.

[42] Lycan (1988, 154-166), Swinburne (2001, 141-142).

[43] Sosa (2007, Ch.3), Earlenbaugh and Molyneux (2009). Rogers and Matheson (2011), Werner (forthcoming).

[44] Huemer (2001, 2007), Cullison (2010b), Tucker (2010), Chudnoff (2011). This view is associated with Phenomenal Conservatism in epistemology, the view that seemings (intuitions) can provide *prima facie* justification for beliefs.

pick out the phenomenon in question. Mathematical intuitions, linguistic intuitions, and metaphysical intuitions provide three such types of cases. If intuitions are *sui generis* representational states, then they may be a good candidate to explain the phenomenal contrast between Norma and Pathos. The idea would be that it non-perceptually *seems* to Norma that the cat burning is bad, but it does not *seem* this way to Pathos.[45] Norma has an intuition that Pathos does not, and this is why their experiences differ.

There are at least two things to say to this alternative explanation. For one thing, it will only serve as a possible alternative explanation to those who accept the account of intuitions as a *sui generis* representational state.[46] This doesn't render this alternative explanation a non-starter, but it does give some reason to search for a more ecumenical explanation of the contrast, if possible. Second, I see no reason to favor this explanation over MP. Note that such an appeal to *a priori* intuition to explain non-moral but otherwise similar contrast cases would be implausible. We don't think, for example, that *a priori* intuition could explain the contrast between the person who can distinguish pines from firs and the person who cannot. So there is some extra burden on those who would favor this explanation to explain how the moral case is importantly different from the non-moral case.[47] Furthermore, for anyone who rejects or is agnostic about the

---

[45] This scenario is still possible if Pathos believes that the cat burning is bad, since we can believe things even despite seemings to the contrary. For example, in the Muller-Lyer illusion, one line continues to seem longer even though I believe that it is not.

[46] To lay my cards on the table, I don't accept the *sui generis* view of intuitions.

[47] One potential difference: Intuitions are prima facie important for moral theorizing, but not for pine tree theorizing. Fully addressing this possible difference would require much more than I can say here about moral

sui generis view of *a priori* intuition, MP serves as a more ecumenical and more explanatorily simple explanation. I don't mean to claim that these burdens cannot be met. I concede that, without a more fully fleshed out picture of this alternative, we can't fully assess the prospects of the *a priori* intuition route for explaining the case of Norma and Pathos.[48] I only mean to stress that more work would need to be done to render this alternative explanation as a real contender.

4.3. A Difference in Non-Moral Properties Represented?

The third alternative explanation acknowledges that the phenomenal contrast between Pathos and Norma is due to a difference in which properties are perceived. However, according to this explanation, the different properties perceived are not (explicitly) moral properties. On this approach, Norma and Pathos directly perceive different properties, but not different moral properties (since neither of them perceives any moral properties at all). The challenge of this approach is to pinpoint a candidate non-moral property to explain the phenomenal contrast. To see why this challenge is a serious one, consider one candidate high-level (non-moral) property, *being-a-cat-in-pain*.

---

epistemology and moral methodology. I do not think that the use of intuitions in moral theorizing, once properly understood, supports the appeal to intuition to explain the contrast in question, because I believe that a priori intuition cannot provide epistemic access to the moral properties. However, defending that claim is outside of the scope of this paper. My positive claim in the text is only that there is an extra burden to be met by the would-be a priori intuitionist in order for their alternative to be properly assessed. I thank an anonymous referee for pressing me on this issue.

[48] Just to be clear, MP is compatible with *a priori* intuition as a source of moral knowledge—it is only incompatible with a view that says that *a priori* intuition is the *only* source of moral knowledge. For all that MP says, *a priori* intuitions may be the source of the vast majority of moral knowledge (though presumably this would be unlikely if MP is true). So the skepticism expressed in the text is only skepticism about *a priori* intuitions explaining Norma's particular instance of forming a moral belief.

Here is how the third alternative explanation could be cashed out with *being-a-cat-in-pain* as the candidate property: When Norma rounds the corner, she directly perceives the property of *being-a-cat-in-pain*. From this perception in addition to other background beliefs, Norma *infers* (perhaps unconsciously) that the property of *badness* is present in the state-of-affairs in front of her. On the other hand, when Pathos rounds the corner, he fails to directly perceive the property of *being-a-cat-in-pain*. He lacks the recognitional capacity to directly perceive that property. Thus, the contents of Pathos's experience are different from those of Norma's, which also explains why Pathos can't infer, at least not in the same spontaneous way, that the state-of-affairs in front of him is *bad*. What is most important is that the phenomenal contrast is a result of Pathos's failure (and Norma's success) at directly perceiving the (non-moral) property of *being-a-cat-in-pain*.[49]

The problem with this explanation is that there is no reason to suppose that Pathos fails to perceive the property of *being-a-cat-in-pain* if we already suppose that Norma does. And this is what would be required to generate a phenomenal contrast. As noted above, EEDIs are not impaired in their ability to perceive the pain or suffering of others.[50] Nor are they impaired in their ability to perceive any other non-moral properties. Since EEDIs are not impaired in their ability to perceive most non-moral

---

[49] Note that this is compatible with Pathos's having the concept of cats' being in pain. Pathos may be able to infer, based on directly perceiving other low-level properties, that *being-a-cat-in-pain* is being instantiated in front of him. And he may even be able to infer from that that *badness* is instantiated in front of him. However, the important difference for the topic at hand is in the properties *directly* perceived.
[50] Richell et. Al. (2003), Blair et. Al.(1996).

properties, the explanation given with respect to *being-a-cat-in-pain* will extend to other alternative explanations of this third sort.

4.4. A Difference in Representations of Internal States?

According to the fourth alternative explanation, Norma's differing phenomenology may be explained by appeal to a representational state about *herself*. She may be representing herself as, for example being in pain or distress as a result of her affective empathy from witnessing the cat's suffering. On this alternative explanation, Norma and Pathos's phenomenal contrast is explained purely by a difference in their internal mental states—Norma is in distress (or some similar state) and represents herself as such, and Pathos is not in distress (because he lacks affective empathy toward the burning cat). No appeal to a difference in their perception of properties "out there", in the world, is necessary. We only need to appeal to Norma's representation of changes within her own bodily states (increased heart rate, SCR, etc.).

To see one way in which this response might fail, consider another kind of emotional response, fear. Suppose I am hiking and I notice a snake near my feet. This invokes a strong fear response with an associated powerful phenomenological character. I freeze in fear—an appropriate response, since with patience the snake will become disinterested and move away. Furthermore, suppose that this fear response and its associated phenomenological character are more generally causally correlated with danger in my local environment, and that the responses in part invoked by states with

this phenomenological character tend to appropriately respond to that danger. A natural question arises—does the phenomenological state associated with fear represent *danger*, (or perhaps more specifically, *dangerous-to-me*)? Or does it merely represent an internal state, *being-afraid*? Perceptual and cognitivist theorists of emotions have compellingly argued that the feeling of fear is capable of representing *danger*, where this is construed as a relation between an object (such as a bee or a snake) and a subject. In brief, this is because the feeling of fear (a) meets all of the appropriate (Dretske/Millikanian) externalist conditions on representation, (b) is capable of having an intentional object, and (c) disposes us to react in ways appropriate to the danger in our local environment (e.g. my freezing when a bee lands on my arm). This perceptual theory of emotions can also be extended to many if not all other emotions with equal plausibility.[51]

Of course we could reject perceptual and cognitivist theories of emotion, and I don't expect to have convinced anyone with what I have said or will say here. However, as noted above, it is widely agreed that emotions can be directed at objects, and most philosophers of emotion accept the possibility of some kind of external world representation in emotional states.[52] We would need some independent reason for rejecting such theories, or at least the representational nature of emotions that these

---

[51] For a defense of a perceptual theory of emotions, see Prinz (2004), especially pp.67-78. For a defense of a (similar in broad outlines) cognitivist theory, see Nussbaum (2001). See also de Sousa (2013).
[52] See de Sousa (2010).

theories entail. Though I can't defend perceptual or cognitivist theories here, I will say that I don't think such reasons are forthcoming.

The last two paragraphs have explained the contrast in terms friendly to the *Affectual Intuitionist,* according to which emotional experiences are a subset of perceptual experiences. But suppose that the perceptual theory of emotions is false. It still could turn out that emotional states, such as Norma's distress, cognitively penetrate her visual processing and thus affect her visual experience. Furthermore, these changes in her visual experience as a result of this cognitive penetration may very well themselves meet conditions (a)-(c) canvassed above. So even if the perceptual theory of emotions is false, we still have a more holistic potential explanation of the differences between Norma and Pathos's experiences which can make more sense of the role of these experiences and their function in each of the individuals' broader cognitive set up. This is the way in which, as mentioned above, *Perceptual Intuitionism* is compatible with the role of emotion in this contrast case.

Whether or not the perceptual theory of emotions is correct, the aspect of Norma's phenomenal character that Pathos lacks meets analogous conditions with respect to *badness* that my feeling of fear meets with respect to *dangerous*. Though Norma's phenomenal character tracks some change in her internal state, it also tracks a property in her external environment. And, plausibly, this phenomenal character disposes her to act appropriately in response—she may try to put out the cat, for

example. Her feeling of distress either represents *badness* in her local environment in the same way that my feeling of fear represents *danger* in mine, or it causally influences visual processing to represent *badness* in her local environment. If the explanation works in the latter case, there seems to be no reason not to extend it to the former.[53] Of course, if Norma's phenomenal character does not track *badness* in her local environment, then all bets are off. But again, we are assuming non-skeptical realism, and we are assuming that Norma generally does a decent job at picking out *badness* in her local environment. So it appears that this fourth alternative explanation, while plausible, also does not provide the best explanation for the phenomenal contrast.

5. Conclusion

The idea that we directly perceive at least some moral properties is one worth taking seriously. It could provide an alternative, or at least a supplementation of, current intuition-based moral epistemologies. In this paper, I have used a *contrast argument,* commonly used in the philosophy of perception, to provide an inference to the best explanation in favor of MP, the thesis that at least some moral properties can be part of the contents of experience. My arguments have concerned one particular moral property, *badness*. I don't take myself to have established that all moral properties can be part of the contents of experience. Since *contrast arguments* must proceed on a case by

---

[53] Robert Cowan (2013a) argues that emotional experiences must be *epistemically dependent* since they "are, or ought to be, responsive to *reasons*" (12). He is also worried (as are many philosophers of perception) about problems with cognitive penetration. (He develops the former of these concern in more detail in 2013b.) I cannot address these concerns here, though they are certainly legitimate worries.

case basis, this is an open question. However, since MP is an existential claim, if the contrast argument for *badness* given above is successful, then MP is true. Nevertheless, I think some reflection will show that *badness* is not unique and that similar *contrast arguments* could be constructed with respect to many moral properties, though perhaps not all. I leave this exercise for future research. In sum, I think we have some positive reason to believe that MP is true, and thus that the direct perception of moral properties is not just possible, but actual.

Chapter 5
***A Posteriori* Ethical Intuitionism and the Problem of Cognitive Penetrability**

According to *a posteriori ethical intuitionism* (AEI), perceptual experiences can

provide non-inferential justification for at least some moral beliefs. Teasing this apart a

bit, AEI consists of two important epistemological claims. The first claim is that the

structure of moral justification is—in at least many cases—foundationalist. Beliefs

formed on the basis of evaluative experiences are at least *prima facie* justified

independently of any other beliefs. The second aspect of AEI is that the justificatory

force of these foundationalist beliefs comes from perceptual experience, rather than

from some traditionally *a priori* cognitive apparatus. Moral epistemology, for the

defender of AEI, is less like the epistemology of math and more like the epistemology of

tables and chairs.

AEI, while still a minority position, has made something of a resurgence in recent

years.[1] But its unpopularity is not without reason: Many serious objections have been

leveled against it, both historically and recently.[2] Perhaps the most common objection is

that it is implausible that evaluative properties could figure in the contents of

experience, as AEI seems to entail.[3] However, even supposing that rejoinders to this

objection have been successful, what seems to be an equally pressing issue lurks right

---

[1] See, e.g., Cullison (2010), McBrayer (2010a, 2010b), Roeser (2011), Audi (2010, 2013), and Cowan (2013a, 2013b)
[2] See, e.g. Strugeon (2002), Huemer (2008, especially 4.4), Cullison (2010), McBrayer (2010a, 2010b), and Faraci (ms).
[3] Värynen (2008, forthcoming), Huemer (2008, 4.4). Defenses of the claim that evaluative properties can figure in the contents of experience can be found in McBrayer (2010a, 2010b), Roeser (2011), and Werner (forthcoming).

around the corner. The worry is that even if evaluative properties could figure in the contents of experience, they would only be able to do so if prior cognitive states influence perceptual experience.[4] Such cognitive penetration in moral perception would arguably undermine the foundationalist credentials of beliefs formed on the basis of those perceptual states. Even granting that evaluative properties can figure in the contents of experience, cognitive penetration provides a serious threat to the foundationalist aspect of AEI.

In this chapter, I defend AEI against the above objection. Rather than deny that cognitive penetration exists, I argue that some types of cognitive penetrability are actually compatible with the AEI's foundationalist structure. I claim that once we notice that cognitive penetration can take distinct forms, we can see that the force of the objection depends on how—if at all—cognitive states affect evaluative perceptual experience. The kind of cognitive penetrability required for the objection to go through is much more specific than its proponents assume. This means that, despite appearances to the contrary, it is proponents of the objection—rather than proponents of AEI—that are compelled to go out on an empirical limb.

The structure of the paper is as follows. I begin in section 1 by briefly canvassing the minimal commitments for the proponent of AEI. In section 2, I explicate in detail

---

[4] This objection is pressed most forcefully by Cowan (2013b, forthcoming) and Väyrynen (forthcoming). A closely related objection is given by Faraci (ms).

both what cognitive penetration is, as well as how its existence is supposed to cast

doubt on AEI. I illustrate the strongest form of the objection by giving what I call the

Epistemic Dependency Argument (EDA). In section 3, I distinguish justificatory questions

from what I call genetic questions, pointing out that a belief's epistemic dependency

hinges on answers to the former but not the latter sorts of questions. In section 4, I

distinguish synchronic and diachronic cognitive penetration, two importantly distinct

ways that cognitive states can influence perceptual experiences.[5] In section 5.1, I use

distinctions drawn in sections 3 and 4 to illustrate that diachronic cognitive penetration

is actually compatible with epistemic independence. If correct, this means that some

perceptual experiences can foundationally justify moral beliefs despite being

cognitively penetrated. Finally, in section 5.2, I examine what all of the above means for

AEI, and propose a tentative model of foundational moral justification for the

proponent of AEI by analogy to cases of perceptual expertise and their cognitive

structure. I conclude in section 6 with some thoughts about where AEI stands with

respect to the epistemological concerns that arise from cognitive penetration.[6]

1. *A Posteriori* Ethical Intuitionism

As I understand it, AEI consists in two substantive claims. First, AEI is a version

of *Ethical Intuitionism*:

---

[5] I don't mean to be taking a stand here on whether these types of cognitive penetration are natural psychological kinds, or about whether they could be divided into further subcategories.

[6] This is all assuming that the philosophical arguments in defense of AEI are successful, which is, of course, very controversial.

> EI: "[N]ormal ethical agents have at least some non-inferentially justified first-order ethical beliefs."[7]

As stated, EI is just the claim that foundationalism—understood in the epistemologist's sense—is true of the structure of ethical beliefs, and that some ethical beliefs are members of the set of foundational beliefs.[8] However, EI does not entail what we might call *Classical Intuitionism*, the claim that the non-inferentially justified ethical beliefs are grounded in *intuitions*, whatever those turn out to be.[9] EI has its unfortunately misleading name presumably because most defenders of EI have also been defenders of Classical Intuitionism.[10] I continue to label EI as such in order to accord with what has now become standard usage in the literature.

The second claim that constitutes AEI is *Ethical Empiricism*:

> EE: The justificatory force of non-inferentially justified first-order ethical beliefs is grounded in the instantiation of evaluative properties in some kind of perceptual experience.

According to EE, non-inferentially justified ethical beliefs are justified analogously to non-moral perceptual justification, on one sort of traditional foundationalist view. The basic picture is as follows. Under certain circumstances, evaluative properties figure in the contents of perceptual experience. Furthermore, at least sometimes, the evaluative

---

[7] Cowan (2013b), 2.

[8] A foundationalist who rejected EI may claim that the structure of moral belief justification is foundationalist, but that the justification of moral beliefs is always mediated by some foundational non-moral beliefs. Something like this is the view defended by Setiya (2013). Note that such a view may also have problems with cognitive penetration insofar as the non-moral perceptual states which justify moral beliefs may themselves be cognitively penetrated in a way incompatible with foundational justification.

[9] As Cowan (2013b, 2n2) himself notes.

[10] See, e.g., Ross (1939), Audi (2004), Roeser (2011), and Tropman (2011, 2014).

properties that figure in the contents of perceptual experience are capable of providing

non-inferential justification for beliefs about the instantiation of evaluative properties.

This is compatible with the claim that sometimes evaluative perceptual experiences fail

to non-inferentially justify. First, there may be defeaters for the justification that an

evaluative perceptual experience would otherwise provide. Second, some evaluative

perceptual experiences may be epistemically dependent in the sense that they can't

provide justificatory force independently of some prior evaluative belief. AEI only

claims that, in at least some circumstances, neither of these things holds. When they do

not, an evaluative perceptual experience can ground a non-inferentially justified moral

belief.[11]

AEI is not a single view but a family of views. I briefly note two recently

discussed versions of AEI. According to *Perceptual Intuitionism*, evaluative properties

are perceived in the same way that any non-moral but otherwise high-level properties[12]

are perceived. For example, the Perceptual Intuitionist may contend that the property

goodness is capable of triggering a recognitional disposition within the visual processing

---

[11] Of course, a view that held that one or both of these conditions holds for almost all instances of moral perception, while being compatible with the letter of AEI, would go against the spirit of the view. I take it that the advocate of AEI is attempting to provide a moral epistemology which entails that human believers can have justified moral beliefs in at least some every day circumstances, on pain of something nearing skepticism. (Thanks to Nicole Dular for pressing me to clarify this point.)

[12] High-level properties are generally thought to be any properties other than color, shape, or edge-based properties. Some examples include kind properties (tables, chairs), causal properties (X-caused-Y), and at least some relational properties (X is-taller-than- Y). It is difficult if not impossible to provide necessary and sufficient conditions for what makes something a high-level property, but this issue is not important for purposes of this paper so I set it aside.

system.[13] On the other hand, according to *Affectual Intuitionism*, evaluative properties are perceived via affective states, which are themselves construed as either a form of perception or as a submodule[14] within a broader perceptual (as opposed to cognitive) system.[15] AEI is also compatible with a number of views about the metaphysics of moral properties—though perhaps it fits with some views more easily than others. So long as one's metaphysics of moral properties is such that they could figure in the contents of experience, one could in principle endorse AEI. Subjectivists,[16] relativists,[17] and both naturalist[18] and non-naturalist[19] realists have all endorsed some version of AEI. The arguments of this chapter, then, are of interest to a variety of otherwise theoretically disparate metaethicists.

2. Cognitive Penetration and the *Epistemic Dependency Argument*

2.1 Clarifying Cognitive Penetrability

---

[13] Such a view seems to best fit with a reductionist conception of the evaluative properties. See Cullison (2010a).

[14] I mean "submodule" here in only the lose sense of being part of a structure in the brain, rather than the strict Fodorian sense entailing information encapsulation.

[15] Affectual Intuitionism (which is the view that I endorse) has both benefits and costs in comparison to Perceptual Intuitionism. The cost is that Affectual Intuitionism requires a commitment to perceptualism about affective states, which is a controversial position among philosophers of emotion. (For defenses, see Nussbaum (2001), Prinz (2006), and Zagzebski (2003). For critical discussion, see Deonna & Teroni (2012, Ch.6).) Some potential benefits of Affectual Intuitionism (which I cannot defend here) are: (a) it is well-situated to explain the tight connection between moral judgments and emotions, (b) it seems better placed than Perceptual Intuitionism to explain the felt normativity of evaluative perception, and (c) because affective states arguably have a dual direction-of-fit, it provides a way of rejecting Humeanism about motivation.

[16] Prinz (2011), Zagzebski (2004) (Though notably her subjectivism is tied to affective responses of God.)

[17] Nichols (2004).

[18] Cullison (2010a), and on one reading, Deonna & Teroni (2012).

[19] Huemer (2005), Roeser (2011).

Jill believes that Jack is angry. This belief causally impacts Jill's experience such that she experiences Jack's face as expressing anger. If she hadn't believed that Jack was angry, she wouldn't have experienced his face as angry.[20] Jill's experience is *cognitively penetrated*. What's important here is not just that the belief affects Jill's interpretation of her visual experience; rather, in cognitive penetration, the character of the experience itself is altered by a prior cognitive state. Presumably, no one would deny that prior beliefs can affect how we interpret experiences. The thesis of cognitive penetration is a more contentious claim about the character of perceptual experience itself.

It will help to more precisely pin down what it takes for a given experience to count as cognitively penetrated. As Susanna Siegel has argued, different conceptions of cognitive penetration will be useful for different projects.[21] The present concern is epistemological, more precisely with foundational beliefs. So an understanding of cognitive penetration which is relevant to the foundational status of beliefs formed on the basis of experience will be most useful.

A simple counterfactual condition—if S hadn't had the cognitive state c, then she wouldn't have had the experience e—won't suffice. Cognitive states can alter our behavior, including our local environment and what we attend to, which will in turn alter what experience we have. For example, my wanting to see what is in my

---

[20] This case is originally from Siegel (2012), 202.
[21] Siegel (forthcoming). See also Stokes (forthcoming).

refrigerator (in part) causes me to open my refrigerator and experience its contents.

Cognitive states that affect our experiences in this way meet the simple counterfactual

condition sketched above. Nonetheless, such cases are intuitively not of a kind with

cases like Jill and Jack. The phenomenon of cognitive penetration is best construed—at

least for present purposes—as a causal relationship between pieces of mental

architecture.

In light of this problem with the simple counterfactual condition, cognitive

penetration is better defined in terms of an *internal* causal dependency of some

experience *e* and a cognitive state *c*, where "internal" indicates that the dependency in

question is all "inside the head".[22] A test developed from Fiona Macpherson captures

this rough idea:

> **(CP)** A perceptual experience *e* is penetrated by some cognitive state(s) *c* iff
> it is not possible for a subject who lacks *c* to have a different experience *e** when
> the nature of the proximal stimulus on the sensory organ, the state of the sensory
> organ, and the location of attentional focus of the subject is held fixed.[23]

(CP) doesn't provide us with a neurophysiological understanding of cognitive

penetration. Nor does it build in the idea of informational encapsulation, an idea that is

commonly thought to be of importance for understanding the full range of

philosophical upshots of cognitive penetration. Despite these weaknesses, it does

provide us with a test which seems to get the cases right, and is plausibly connected to

---

[22] Wu (2013, Section 5) is helpful on this point, though he rejects Macpherson's test.
[23] This test is a loose adaptation of a principle found in Macpherson (2012), 29.

those features. For example, if different subjects can have different experiences despite all other sensory features held fixed, this suggests that there is some present or historical information flow between cognitive states and the perceptual system. Rather than getting too bogged down in these issues, I will understand cognitive penetration in what follows in terms of (CP), while acknowledging that a more precise notion will be theoretically required in the long run.

The case of Jack and Jill involved the cognitive penetrability of vision. Affective experience can also be subject to cognitive penetration. Jona Vance provides an example:

> **Whistle Fear.** Winnie believes that her family is out of the house late at night and *if someone else is in the house, they're an intruder.* Winnie then hears someone whistle close behind her. Her background belief that *if someone else is in the house, they're an intruder* (together with the auditory stimuli as of the whistle) plays a role in causing her to feel fear. If she lacked the belief, she would not feel fear in response to those stimuli.[24]

It is widely believed that some affectual experiences, e.g. the experience of fear, can justify beliefs, e.g. that danger is present. If the cognitive penetrability of visual experience has any consequences for the epistemological role of those experiences, it is plausible that similar consequences also hold for affective experiences. Both perceptual and affectual intuitionists should be concerned by the potential epistemological challenges that arise from the existence of cognitive penetration.

---

[24] Vance (2014), pp.270-271.

2.2 The *Epistemic Dependency Argument*

The proponent of AEI claims that perceptual experiences (broadly conceived as to include affective experiences[25]) can at least sometimes provide justification for moral beliefs. Furthermore, these moral beliefs are non-inferentially justified insofar as they are justified independently of any other beliefs. The non-inferential nature of this justification is crucial, since without it the proponent of AEI would require some other source for foundational moral beliefs, and no other non-*a priori* route seems available. So for AEI to be tenable as a position in moral epistemology, some of our moral beliefs must meet the condition that Robert Cowan calls "epistemic independence":

> *Epistemic Independence* (EI): A justified belief *b* is epistemically independent iff there is no cognitive state *d* such that *b*'s justification depends, directly or indirectly, on *d*'s epistemic status.[26]

A belief which depends for its justification on a perceptual experience of a certain sort will count as epistemically independent so long as the perceptual experience alone is enough to do the justificatory work. A belief *b* is *epistemically dependent*, then, when there is some cognitive state(s), *d*, that *b*'s justification depends on. Suppose, for example, that I have a justified belief that I don't have both almond milk and soy milk in my refrigerator. I open my fridge and have the perceptual experience as of there

---

[25] Unless otherwise indicated, I intend "perceptual experience" in this broader way, encompassing affective experiences as well.

[26] I take this conception of epistemic independence from Cowan (2013a), 10—though of course epistemic independence is a pervasive notion in foundationalist epistemology. My formulation of EI is different from Cowan's, but I believe it captures the same idea.

being a carton of soy milk. As a result, I form the justified belief that *I have no almond milk in my refrigerator.* Although the justification for my belief is partially grounded in my experience of the soy milk in my fridge, it is also partially grounded in my belief that I don't have both almond milk and soy milk. In this way, my belief is epistemically dependent.

The case of epistemic dependency just given is uncontroversial. But what should be concerning to proponents of AEI (as well as perhaps foundationalists generally) is that there are seemingly epistemically analogous cases involving cognitive penetration. The justification for Jill's belief that Jack has an angry expression on his face is grounded in her experience of Jack as having an angry expression on his face. But her experience of Jack's having an angry expression on his face is causally dependent on her prior belief that Jack is angry. So it seems that her experience's justification conferring status—if any—is in turn grounded in the justification she has for her prior belief that Jack is angry. So her belief that Jack has an angry expression on his face is epistemically dependent, and so not foundational.

Suppose the reasoning of the previous paragraph is correct. Further, suppose that all or almost all of our perceptual experiences of evaluative properties are a result of cognitive penetration. Then it would seem that all or almost all of the beliefs we formed on the basis of those experiences would be epistemically dependent in the same way that Jill's belief that *Jack has an angry look* would be. This would cast serious doubt

on AEI, since AEI entails that there is some non-trivial set of justified moral beliefs

which are epistemically independent. Without a set of epistemically independent

justified moral beliefs, AEI loses its foundationalist credentials.

It will be helpful to formulate this argument somewhat more formally. Take

some moral property *M* which is part of the contents of some perceptual experience *E*:

> (1) If *M*'s instantiation in *E* is causally dependent on the penetration of some
> prior cognitive state(s) *S*, then any moral belief *B* which is based on *E* is
> epistemically dependent in the sense that *B* is justified (if at all) partially in virtue
> of *S*.
> (2) All moral content in perceptual or affective experience is causally dependent
> on the penetration of prior cognitive states.
> (3) So *M*'s instantiation in *E* is causally dependent on the penetration of some
> prior cognitive state(s) *S*. (From 2, Existential Instantiation)
> (4) So any moral belief *B* which is based on *E* is epistemically dependent in the
> sense that *B* is justified (if at all) partially in virtue of *S*. (1, 3 MP)
> (5) But this reasoning extends to any moral property which figures in perceptual
> or affective experience.
> Therefore,
> **(Con)** AEI is false.

We can call this the *epistemic dependency argument*, or EDA.

What should the proponent of a posteriori ethical intuitionism say about EDA?

(1) looks

plausible, considering the Jack and Jill case discussed above. As (2) is not an

epistemological claim, but a causal one (with epistemological consequences), its status

is primarily an empirical question.[27] Furthermore, it is an empirical question that is far

---

[27] (2) is not a wholly empirical question, since questions about how mental states get their contents are at least partially philosophical. But I set this aside for the purposes of this paper.

from settled.[28] This makes (2) a natural place to put pressure on EDA. But while the

pervasiveness of cognitive penetration—both with respect to morally loaded cognitive

states as well as others—is not yet established, pointing this out is a non-ideal response

on behalf of AEI. (2) remains a serious empirical possibility, so by rejecting EDA on

these grounds, the defender of AEI is making herself hostage to empirical fortune.

However, the rest of the argument follows logically from the first two premises. So if (2)

remains safe from empirical refutation, the argument looks initially to be very powerful.

While research which impinges on the truth of (2) is certainly relevant to AEI, a

better response to the argument would involve providing some positive reason to

doubt one of the premises, rather than just providing reasons to remain agnostic. In

what follows, I argue that, even granting (2), EDA as given still fails, because (1) is not

true. That is, I will argue that not all prior *causal* dependence of a perceptual experience

on a cognitive state thereby renders it *epistemically* dependent. While this move from

causal dependence to epistemic dependence is easy and intuitive to make, once we

make explicit the difference between a perceptual experience's justification conferring

status and the historical explanation for its justification conferring status (what I call its

genetic explanation), a case can be made that some cognitive penetration is compatible

with epistemic independence.

---

[28] The most well-known argument against the cognitive penetrability of perceptual states is Pylyshyn (1999). Raftopoulos (2009) also argues against penetrability. Arguments in favor generally focus on just one or two candidate instances of penetration, e.g. Macpherson (2012), Wu (2013). For an excellent and current overview of the evidence in favor, see Vetter & Newen (2014).

3. Epistemic Dependence and the Justification/Genetic Explanation Distinction

What does it take for a given mental process to have *prima facie* justification

conferring status? This is one of the deepest questions of contemporary analytic

epistemology, and I certainly don't propose to answer it here. The answer may appeal

to facts about the reliability of the outputs of the process,[29] the phenomenology of the

process or its outputs,[30] the relationship (e.g. counterfactual) between the outputs of a

process and the facts,[31] or some combination of these three or other features. Let's call *F*

the set of features, whatever they turn out to be, that a mental process or its outputs

must have in order for it to have *prima facie* justification conferring status.

Having *F* alone will not always be sufficient for the outputs of an otherwise

justification conferring process to have justificatory force.[32] This is because some mental

processes themselves take as their input the outputs of other mental processes. Suppose

I infer from my belief that *P* and my belief that *if P then Q* that *Q.* There is nothing

wrong with the mental process of *modus ponens.*[33] But my belief that *Q* may still fail to be

justified, since the beliefs that fed into the process of *modus ponens* may themselves be

---

[29] See Goldman (1979, 2008).

[30] Huemer (2001), Tucker (2010).

[31] Nozick (1981), Sosa (1999).

[32] For brevity, I from here on out drop "prima facie" from my claims about justification-conferring states. Unless indicated otherwise, I don't assume that the outputs of a mental process are more than prima facie justification-conferring.

[33] I use modus ponens as a mental process for simplicity (and relative theory-neutrality about *F*), not because I endorse any particular view of mental process individuation. I remain agnostic about whether modus ponens counts as a mental process, an instance of some broader mental process, or not a joint-carving mental process at all.

unjustified. Generalizing a point originally made by reliabilists,[34] some processes only

generate justificatory force conditional on the justificatory force of their inputs. Garbage

in, garbage out, as the cliché goes.

When we are asking about how a particular belief is justified, we are looking for

an explanation in terms of how the belief was the output of some process or chain of

processes which all have feature(s) F.[35] However, this need not be the end of inquiry.

The epistemologist may be interested in a further question, about how these mental

processes came to have F in the first place. Our epistemologist could here be asking two

questions. First, she could be asking why it is that feature(s) F (rather than G, H, or J) are

the necessary conditions on a process's being justification conferring. This is a close

relative of the question that I set aside at the beginning of this section, about what

features F are. I mention it again here only to disambiguate from the other question our

epistemologist may be asking. The second kind of question the epistemologist could be

asking is distinctly causal. That is, she wants to know how these particular mental

processes came to have F: Was it through some evolutionary selection? Social

conditioning? Taking a logic course? Call this question the *Genetic Question,* or GQ.

GQ is an intrinsically interesting (and largely empirical) epistemological

question. Potential answers to the question can also be important for first-order

---

[34] Goldman (1979).

[35] I am assuming throughout this discussion that some version of foundationalism is correct. I am also setting aside complications about defeaters which may undercut the *prima facie* justification-conferring status of one or more of the processes or their outputs.

questions of justification as well, because the causal history of a given mental process may debunk that process as unreliable, or as reliable but in a lucky way such that its justification conferring status is undermined.[36] But so long as the explanation of a process's reliability is not debunking in this way, it will not figure in the explanation of the first order justificatory story about a belief formed as a result of that process.

To illustrate this, consider an example of the belief that *that object is green* formed on the basis of a (non-cognitively penetrated) visual experience. At the very least, the explanation of this belief's justification will appeal to its being based in the visual experience, along with the experience's character.[37] The explanation may also include the fact that visual experience (in this agent, in these circumstances, etc.) has feature(s) F (e.g. is a reliable process, has the proper justification conferring phenomenology, etc.). But what isn't required to explain the justification of the belief that *that object is green* is the full causal history of the visual system of that agent or the evolutionary history of the visual system of the species of which the agent is a member. To demand such stories as part of the explanation of the justification of a first order belief is to confuse the question of a belief's justification with genetic questions about the justification conferring process(es) that resulted in the belief. The explanation of first order

---

[36] Such as, for example, if the reliability is not a result of epistemic access to the facts in question (see my "How to Understand the Access Requirement for Non-Naturalist Realism").

[37] For the dogmatist or phenomenal conservative, this will likely be the complete explanation of the belief's justification. See Huemer (2001), Tucker (2010).

justification bottoms out in the appeal to processes which have feature(s) *F*, whatever they turn out to be.

I take this distinction to be relatively uncontroversial, at least to anti-skeptical foundationalists.[38] But as I shall soon argue, it is only by conflating justification and genetic explanation that premise (1) the epistemic dependency argument gains its intuitive appeal. Once we assess premise (1) in light of this distinction, we can see that in order to be true it must be restricted. Once properly restricted, it becomes clear that some beliefs will retain their epistemic independence despite being formed on the basis of cognitively penetrated experiences. But first, we must distinguish four ways that cognitive penetration can take place.

4. The Varieties of Cognitive Penetration

We can functionally carve up types of cognitive penetration along two different metrics, for a total of four types of cognitive penetration.[39] First, we can distinguish between the cognitive penetration which involves representational states vs. cognitive penetration which involves non-representational, or orectic,[40] states. The paradigmatic instances of these kinds of cognitive penetrability involve beliefs and desires,

---

[38] Internalists should accept it on pain of denying that visual experiences can justify for anyone who doesn't have access to how the visual system evolved (i.e. nearly everyone in history). Externalists should accept it, because a process' ability to be reliable or track the truth is an intrinsic property of the process, rather than a historical one.

[39] Of course, we could differentiate types of cognitive penetration along lines other than functional, e.g. on neurophysiological grounds. However, for present purposes a functional division is the most helpful way to illustrate the epistemological implications of cognitive penetration.

[40] The idea (and terminology) of orectic cognitive penetrability is discussed in Stokes (2012).

respectively. I will set aside orectic penetration in what follows. The second important

difference between cases of cognitive penetration is the difference between synchronic

and diachronic cognitive penetration. Synchronic penetration occurs when the

background cognitive state(s) influence the visual system as the visual processing takes

place in generating a visual experience. In other words, the visual experience is causally

dependent on the existence of the penetrating cognitive state(s) at t or immediately

prior to t. In synchronic penetration, the cognitive states' influence on the perceptual

system is co-temporal with the visual processing itself. The case of Jill and Jack

discussed above provides an example of synchronic belief penetration. On the other

hand, in diachronic penetration, rather than being co-temporal, the influence of the

cognitive states on the perceptual system takes place prior to the particular visual

process that is subject to the influence. Research on the Müller-Lyer Illusion (figure 1)

provides an interesting potential example of this. Segall et. al. (1963) illustrated that

susceptibility to the illusion varies cross-culturally.[41] However, it is unlikely that this

variance is due to synchronic penetration, because the illusion doesn't disappear (for

those who experience it) after an agent knows that it is an illusion.

---

[41] Much debate since the original study has arisen amongst psychologists about what explains the difference. See, e.g., Berry (1971), McCauley & Henrich (2006), Davis & Carlson (1970). On one recalcitrant view (see Ahluwalia 1978), people who grew up in "carpenetered" environments (with many square cornered buildings) are more susceptible to the illusion than those who did not. Figure 1 is from http://www.rit.edu/cla/gssp400/muller/muller.html , which also contains an explanation of the illusion.
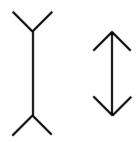
Figure 1

One plausible interpretation of this phenomenon is that one's susceptibility (or lack

thereof) to the illusion is based on past experiences of patterns similar to the lines found

in figure 1 usually favored a particular interpretation of such figures. This favoring led

(presumably unconscious) cognitive processes to reshape the late visual processing

system to tend toward the interpretation that has been correct in the past (for that

agent).[42] This "interpretation" all occurs in visual processing prior to visual experience,

which is why the illusion can't be easily overcome by, for example, shifting one's

attention.[43]

The Müller-Lyer case, and culturally relative visual illusions more generally, are

helpful for illustrating how diachronic penetration may occur even while a visual

experience is not synchronically penetrated. There are many other plausible instances of

diachronic penetration. A relevant set of cases are those that involve perceptual

---

[42] Cecchi (2014, 63) calls this "architectural cognitive penetration", because this sort of cognitive penetration involves "modulating [the perceptual system's] structure…[through] neural reorganization." I choose not to use Cecchi's term here, since, on his conception, some architectural cognitive penetration is synchronic.

[43] For an extended discussion of the relationship between diachronic cognitive penetration and the Müller-Lyer Illusion, see McCauley & Henrich (2006).

expertise, such as a musician who develops perfect pitch or a coffee connoisseur who develops her olfactory and gustatory discriminative abilities. Perceptual expertise arguably (though not uncontroversially[44]) alters perceptual experience, which suggests cognitive penetrability of some sort. The fact that perceptual expertise develops over some period of time suggests diachronic penetration as an explanation. However, an alternative interpretation of perceptual expertise alleges that perceptual expertise involves synchronic penetration. On this alternative, the explanation of the change in experience over time is that the perceptual expert develops her beliefs and knowledge about the phenomenon over time and that these changes in beliefs synchronically penetrate the perceptual expert's experiences. I focused on the Müller-Lyer case above because I think it is one of the clearest cases of diachronic penetration, whereas in the perceptual expertise cases the synchronic penetration interpretation remains empirically plausible. But I mention the perceptual expertise case as well because it provides a helpful possible model of the relationship between an agent's moral beliefs and her affective and perceptual moral experiences according to the proponent of AEI. I will return to this model below (Sect. 5.2).

5. A Restricted Defense of AEI against the Epistemic Dependency Argument

5.1. Synchronic and Diachronic Penetration: Test Cases

Return to the epistemic dependency argument. Recall premise (1):

---

[44] Those who deny cognitive penetrability altogether would deny that perceptual expertise is properly perceptual.

> (1) If *M*'s instantiation in *E* is causally dependent on the penetration of some prior cognitive state(s) *S*, then any moral belief *B* which is based on *E* is epistemically dependent in the sense that *B* is justified (if at all) partially in virtue of *S*.

The clearest case for (1) lies in instances of synchronic penetration. Return to Jill and Jack. Jill believes that Jack is angry. When she looks at Jack's face, this belief causally influences her visual system as it is processing the visual information which will generate her experience of Jack's face as having an angry expression. Information flows between Jill's cognitive system and her visual perceptual system during the time in which the experience is formed. Her belief that Jack has an angry expression is grounded in her experience of Jack's face as having an angry expression. But this particular experience's character depends, at present, on her prior belief. So it seems to be epistemically dependent on her prior belief.

The case of Jill's belief that Jack has an angry expression is not unique; the features of her case will plausibly generalize to all standard cases of synchronic penetration. By definition, synchronic penetration involves a cognitive state's being directly relevant to a particular property's representation in experience. This can be illustrated by considering the nature of the experience if the cognitive system were temporarily blocked from influencing the visual processing system. If this were the case, Jill's visual experience of Jack's face would have differed. This illustrates that the belief is epistemically dependent on Jill's prior belief that Jack is angry. Of course, this isn't to say that Jill's belief that Jack has an angry look isn't justified. Rather, it is to say

that if it is, it is justified partially in virtue of the justification-conferring status of her prior belief that Jack is angry. In sum, synchronic representational penetration is compatible with justified belief, but not with foundationally justified belief—as the proponents of EDA claim.

Next we can assess whether the same line of reasoning applies to diachronic penetration. To test this, return to the proposed explanation of different individuals' susceptibility to the Müller-Lyer illusion. According to one proposal, our susceptibility to the illusion depends greatly upon the architectural environment in which an individual developed.[45] Consider Rebekah, who will be our exemplar stereotypical "Western" perceiver. Rebekah grew up in Generic, USA, an environment filled with pervasively straight lined, right angled architecture. In such an environment, visual inputs which are similar to those found in the Müller-Lyer illusion tend to represent straight lines which obey certain patterns. As one guide explains (from the perspective of a Western observer such as Rebekah):

> In the three-dimensional world, depth perception concerns judging distance. The closer an object is to the retina, the larger it is on the retina. However, in the two-dimensional world of the Muller-Lyer illusion, our brain makes assumptions about the relative depths of the two shafts based on monocular (pictorial) cues. We are used to seeing outside corners of buildings as near to us with the top and bottom of the corner sloping out and away (like the outward slanting fins of the Muller-Lyer illusion). We are used to seeing inside corners of buildings as farther

---

[45] Ahluwalia (1978). This hypothesis has been challenged. See Berry (1971), Segall et. al. (1966), and Stewart (1973) for further discussion.

from us with the top and bottom of the corner sloping in somewhat towards us (like the inward slanting fins of the Muller-Lyer illusion).[46]

What happens to Rebekah, then, is that her visual processing system "infers" certain things about the length of the lines, given the historical tendency of lines like these to obey certain patterns in her environment. Her cognitive system influenced her visual processing system over some pattern of time, priming it to favor the interpretations that are more reliable given her Westernized environment. Crucially, what separates Rebekah's case from synchronic cases of penetration is that, if the diachronic hypothesis is correct, the cognitive influence is not taking place while Rebekah is looking at the illusion. In fact, at least with respect to information relevant to interpreting the length of the lines, Rebekah's visual system is presently wholly informationally encapsulated.[47] Her cognitive system is at present playing no role in the character of her experience of the lines as being different lengths. Rebekah may full well *know* that the lines are the same length, but the illusion will not disappear. If this is a case of cognitive penetrability, it is arguably diachronic.

Suppose, though, that Rebekah has never seen the Müller-Lyer illusion before. She looks at it closely and, based on her visual experience, forms the belief that *the line on the left is longer*. What, if anything, is her belief's justification grounded in? As with all of our cases, her belief would be at least partially justified in virtue of her visual

---

[46] http://www.rit.edu/cla/gssp400/muller/muller.html
[47] Fodor (1983, 66), Pylyshyn (1999).

experience as of the line on the left being longer.[48] But the proponent of EDA must claim

that her visual experience is itself dependent on some background (cognitive) theory

about lines like these and their tendency to obey certain patterns. Of course, this theory

is presumably unconscious. But we can grant for the sake of argument that such a

theory is still a part of Rebekah's broader cognitive system, and that this theory

previously exerted some influence on her visual system. A background cognitive theory

has causally influenced Rebekah's visual system, restructuring it such as to prime it to

favor certain interpretations of data. This is just how diachronic penetration works. The

proponent of EDA must claim that this past restructuring renders Rebekah's current

belief epistemically dependent on her background cognitive theory.

I contend that this is to conflate a justificatory question with a genetic question.

What is relevant to the justification of Rebekah's belief is whether or not the processes

that caused it have feature(s) F (whatever they turn out to be). If Rebekah's visual

experience has F, then it is justification conferring. And her visual experience—in this

case—was produced by her visual processing system, without any cognitive

informational input. Her visual experience, if it is justification conferring at all, is

---

[48] This case is complicated since Rebekah's belief here is subject to the illusion and so false. I am assuming, for explanatory simplicity, that her belief is justified. In most cases, Rebekah's dispositions to interpret Muller-Lyer type appearances as she does will lead to *increased* reliability given her local environment. (This is why she developed the disposition to interpret the lines in this way, after all, given her 'carpentered' environment.) The reader should not get hung up on this disagreement, because the same line of argument given in the text could be given with respect to Rebekah's claim to justification in her non-experimental real world experiences of right-angled walls and buildings. I thank Teresa Bruno-Niño for pressing me to make this clearer.

foundationally justification conferring.[49] Now, the proponent of EDA may protest, 'Of

course her visual experience involved cognitive informational input. For her visual

system was primed to interpret data in this way by her cognitive system. So there was

cognitive informational input, it's just that it occurred in the past.' This is correct, but

not directly relevant to our justificatory question. This past cognitive input may explain

how Rebekah's visual experiences of this type came to have *F*, but that is decidedly a

genetic question about Rebekah's visual processing system. Pointing to the fact that

Rebekah's belief resulted from a mental process with features *F* answers our

justificatory question. But once we have moved beyond this to further explaining *how*

the given mental processes got to have the structure and tendencies that they do, we've

stopped talking about justification and started talking about genetic explanations.

Compare: Suppose I was asked how my belief that *there is a green cup on the table*

is justified. I'll appeal to my visual experience. If my experience is synchronically

penetrated, I may have to appeal to some prior beliefs or recognitional concepts that I

have. But now suppose my interlocutor responds: 'You haven't even begun to show me

how your belief is justified. In order to do that, you have to provide me with a full

evolutionary history of your visual system and how its evolutionary adaptations are

---

[49] Again, this claim is limited to this particular case, since other visual experiences may involve synchronic penetration.

truth-conducive. Only once you've done that can you claim to be justified in believing that *there is a green cup on the table.*'

By demanding the evolutionary history of my visual system, my interlocutor has just changed the subject. If my visual system and any prior beliefs appealed to are justification conferring, then, barring any defeaters floating around, the total justification for my belief has been provided. Asking how the justification conferring processes became justification conferring is surely to ask an interesting question, but not one that directly bears on the justification for my *particular* belief that *there is a green cup on the table.* Of course, in some cases, a mental process' history may provide an undercutting defeater for the claim that the process is justification conferring after all. But, as noted above, so long as no defeaters are present, genetic explanations will have no direct bearing on the justificatory status of some particular output of a justification conferring process. My belief is not rendered epistemically dependent by the fact that my visual system has an evolutionary history.

The demand on behalf of the proponent of EDA is similar to that of the imagined interlocutor above. She claims that the causal history of the visual system's justification conferring status must itself be part of the justification of Rebekah's particular belief. The mistake is more difficult to notice, perhaps because the relevant history is cognitive, rather than un-designed and evolutionary. But once we are careful to distinguish justificatory questions from genetic questions, it becomes clear that the causal history of

Rebekah's visual system is not part of the justification for her particular belief in the

Müller-Lyer case.

5.2. AEI and Degrees of Perceptual Expertise

The case of the Müller-Lyer illusion illustrates that (1) is false. But now we must

ask whether the lesson of the Müller-Lyer illusion generalizes to all cases of diachronic

representational penetration. If it doesn't, the general epistemological worry that lies

behind EDA will remain. So let's consider a kind of cases more amenable to an AEI-

friendly model of foundational moral knowledge: Perceptual expertise. I begin with a

non-moral case of perceptual expertise, before turning to the moral case.

We have some empirical reason to believe that perceptual expertise, in many

cases at least, is an instance of diachronic representational penetration.[50] Interestingly, in

these diachronic models of perceptual expertise, the relevant penetration is wholly

analogous to the case of the Müller-Lyer illusion discussed above. Krigolson et. al.

(2009) explain:

> [T]he ability to make object classifications is dependent upon feedback-related
> reinforcement learning. This begins with feedback signaling from the medial
> frontal areas…to early sensory areas…As the training progresses, these observers
> develop the ability to map the internally generated responses onto their early
> sensory percepts…This may eventually lead to better representations of the
> object categories in the early sensory areas.[51]

---

[50] Though synchronic penetration remains an empirically possible alternative still taken seriously. For a helpful overview of the research, see Wong & Wong (2014). See also, e.g. Krigolson et. al. (2009), Tanaka & Pierce (2009), Cecchi (2014), and Stokes (2014).

[51] Krigolson et. al. (2009), 113.

Just as Rebekah's visual system became primed over time to interpret Müller-Lyer lines in a certain way, becoming a perceptual expert involves the development of a perceptual system, through cognitive feedback, to make sensory discriminations that the perceptual novice cannot make. There is evidence that this cognitive feedback can "[influence] the architecture of the visual system determining the content of the subject's perceptual states".[52] Further support for the diachronic model of perceptual expertise for some domain could be provided by establishing synchronic perceptual encapsulation. For example, if the coffee expert cannot help but taste the hints of citrus in some cup of coffee, even despite belief to the contrary, we would have some reason to believe that there was no synchronic belief penetration. (Such research, has not, so far as I know, been done.)

Cases of perceptual expertise which fit this diachronic model are epistemically analogous to the case of Müller-Lyer illusion above. The upshot is that many of the beliefs formed on the basis of perceptual experience by perceptual experts will be epistemically independent and thus foundationally justified. Which kinds of perceptual expertise will fit this model is an empirical question, but it is empirically promising that many kinds will. This is good news for the foundationalist.

We can finally turn to the upshot of all of this for the plausibility of AEI. Critics of AEI *qua* foundationalism claim that an evaluative property couldn't figure in a

---

[52] Cecchi (2014), 91.

perceptual experience (broadly conceived) unless there were some cognitive penetration occurring. And, as EDA claims, any cognitive penetration of an experience is incompatible with a downstream belief's formed foundational status. However, the discussion above illustrates that there is a possible model of evaluatively-laden perceptual experience that is compatible with AEI.

Here is the sketch of how AEI goes on this model. An agent begins as a moral novice. Still, as a result of affect-laden experience, she has positively and negatively valenced experiences. Over time, the agent develops moral beliefs which influence and restructure her perceptual and affective systems. This influence primes her perceptual and affective systems to trigger positively and negatively valenced experiences which track evaluative properties. The agent forms beliefs about particular cases based on these positively and negatively valenced experiences, e.g. that making *that* face at someone is bad. These beliefs will be justified insofar as the experiences they are based off of have feature(s) F. These beliefs will be foundationally justified insofar as the experiences they are based on are not synchronically penetrated.

Of course, as much as we'd like to think we are expert moral perceivers, this may not be true. Perceptual expertise comes in degrees, and so too would be the case with moral perceptual expertise. So whether a given individual's experience grounded moral beliefs are justified is a difficult question. But it is not a question particular to *moral* expertise, and so no special ground for skepticism about beliefs formed on the basis of

evaluative experience. Just as with all perceptual expertise, then, the justification

conferring status of a particular individual's moral experiences will be a difficult

assessment to make. What is important at present is that we have an empirically

plausible model of moral experience which illustrates the flaw in EDA. As a result, we

have an empirically plausible model of evaluative experience which is compatible with

AEI, which is just what we were looking for.

   One limitation on this model should be emphasized. In order to preserve the

foundational nature of AEI, it must turn out that at least many of the perceptual and/or

affective experiences that ground our particular moral beliefs do not have the character

they do as a result of synchronic cognitive penetration. How frequently, if at all, moral

experiences are synchronically penetrated is an empirical question to which we don't

yet know the answer. The ultimate plausibility of filling out AEI in this way is held

hostage to future empirical research. The most salient alternative to this understanding

of AEI would be the approach which denies the cognitive penetrability of evaluative

experiences altogether. Compared to this alternative, the diachronic penetration model

looks to be far more empirically plausible. The extent and nature of cognitive

penetration is an open empirical question, but it is becoming more and more difficult to

deny its existence altogether.[53] In this way, the alternative to the model defended in this

paper seems far more empirically doubtful. If AEI is to remain a serious contender for

---

[53] Vetter & Newen (2014).

our moral epistemology, then, I think the diachronic penetration model defended here is its best shot.

6. Conclusion

AEI provides an alternative to *a priori*, rationalist moral epistemology, and even provides some *prima facie* advantages or its *a priori* rivals.[54] However, the existence of cognitive influences on perceptual and affective experience threatens to undermine AEI's foundationalist credentials. In this paper, I have explored an argument against AEI based on this intuition, EDA. I argued that diachronic representational penetration is, contrary to EDA, compatible with an experience's foundationally justification conferring status. If the kind of cognitive penetration on evaluative experiences is diachronic, then, AEI can retain its plausibility. On one popular empirical model of perceptual expertise, perceptual expertise involves diachronic rather than synchronic penetration. On one understanding of AEI, evaluative experience involves an analogous form of perceptual expertise. The upshot here is that, if this popular empirical model is correct, at least for evaluative experiences, then the kind of cognitive penetration involved in evaluative experiences provides no threat to AEI after all.

---

[54] Proponents of AEI have appealed to many considerations, none of which are uncontroversial, in support of this claim. It may be thought, for example, that AEI better fits with empirical data on the psychology of moral belief formation, that AEI is more phenomenologically realistic, that it fits better with a naturalistic/empiricist worldview, and/or that it can better account for our epistemic access to the moral properties.

Chapter 6: Perceptualism about Emotions and the Rational Criticizability Problem

Introduction

According to *perceptualism*, emotional experience is a kind of perceptual experience.[1] For example, when I see a snake on the trail, I feel fear. This fear represents the snake as dangerous. One feature of perceptualism is its epistemological component: Since emotional experience is a form of perceptual experience, it *prima facie* justifies beliefs as traditional sense perception can.[2]

Besides being of intrinsic interest, perceptualism is also relevant to Affectual Intuitionism, a version of *A Posteriori* Ethical Intuitionism, in moral epistemology. According to Affectual Intuitionism (AI), emotional experiences provide *prima facie* non-inferential justification for moral beliefs.[3] For example, when Pam senses Jim's discomfort at Dwight's line of questioning, she perceives this as bad, and thus tactfully changes the subject. Pam's belief that Jim is being harmed is justified by her experience of discomfort. For the proponent of AI, then, the role that affective experience plays in ethics is analogous to the role that perceptual experience plays in the justification of run-of-the-mill empirical beliefs.

---

[1] See Soloman (1984), De Sousa (1984), Prinz (2004).

[2] Of course, *how* perceptual experience justifies is itself a contentious issue. But the perceptualist about emotions can remain agnostic on this question—she need only claim that, *however* perceptual experience justifies, emotional experience does in a similar way. See BonJour (2007) for discussion.

[3] See Prinz (2008), Roeser (2011), Cowan (forthcoming-a, forthcoming-b), Werner (forthcoming).

Objections to perceptualism about emotions should cause concern, for both perceptualists and proponents of AI. This is especially true of objections to perceptualism that purport to undermine the claim that emotional experiences can *non-inferentially* justify. Recently, several philosophers have leveled one objection of just this sort.[4] As will be explained in more detail below, the idea is this: Any state that can non-inferentially justify will be immune from rational and moral criticism. In other words, non-inferentially justifying states, such as perceptual experiences, cannot themselves be justified or unjustified—that's what makes them epistemic regress stoppers, after all. However, emotions are criticizable in just this way. So emotions cannot play the epistemic role that perceptualists and proponents of AI claim they do.

The project of this chapter is to defend perceptualism in light of this objection, which I call the *Rational Criticizability Problem*, and thus indirectly defend AEI. To foreshadow, I claim that once we look more closely at rational and moral criticisms of emotional states, it becomes clear that the criticisms target not the epistemological component of emotional experiences but one of a cluster of other properties of the agent which are closely related to the emotion. It shouldn't come as a surprise, given the role affect plays in our cognitive economy, that even if affective experience is a special form of perceptual experience, it will be bound up with action tendencies and values in a

---

[4] Of course this is not the only objection to perceptualism. See also, e.g. Deonna & Teroni (2012), Whiting (2012), and Lacewing (2015). Though I do not find these objections compelling, addressing *every* such objection to perceptualism is outside of the scope of the present paper.

unique way. However, this doesn't entail anything about the epistemic status of emotional experiences.

The structure of the paper is as follows. In section 1, I explain and (briefly) motivate AI and perceptualism as versions of AEI. Then, in section 2, I lay out the *Rational Criticizability Problem* in detail. In sections 3, 4, and 5, I respond to the objection by surveying several case studies in criticism of emotions and their relation to the role that affect plays. Finally, in section 6, I assess where things stand.

1. Affectual Intuitionism and Perceptualist Theories of Emotions

On a familiar foundationalist view, all empirical justification, though it may be mediated by an arbitrary number of other beliefs, bottoms out in perceptual experience (plus some basic principles of reasoning, which may be *a priori*). With respect to empirical justification, perceptual experiences seem well-placed to serve as justificatory regress stoppers.[5]

Suppose foundationalism, broadly speaking, is correct. What the moral epistemologist needs, then, is a plausible story about what could serve as a regress stopper, analogous to perceptual experience, for justified evaluative beliefs. One such story is Affectual Intuitionism:

> *Affectual Intuitionism* (AI): "Normal ethical agents have non-inferential justification for first-order ethical beliefs in virtue of emotional experiences."[6]

---

[5] See Fumerton & Hasan (2010), especially Section 1.

[6] Cowan (forthcoming-a), 11. This is actually Cowan's helpful precisification of a view most prominently defended by Roeser (2006, 2011). As we'll see below, Cowan is at best ambivalent about AI.

To see how AI works, consider a cliché case:[7] Norma rounds the corner to see some

teens lighting a cat on fire, while the cat struggles to get away, clearly (and justifiably!)

afraid. Norma has a strong emotional response of anger, sadness, and perhaps even

disgust at the sight of the cat's suffering and the teens' laughter. Rapidly and

unreflectively, on the basis of this emotional response, Norma forms the belief that the

cat's suffering is (morally speaking) very bad. According to the proponent of AI, two

things are true: First, Norma's belief that the cat's suffering is bad is grounded in her

emotional experience (in the sense of being non-deviantly caused). Second, this

grounding relation also serves to render Norma's belief *justified*. This is because

emotional responses of this sort in Norma more or less reliably track badness in her

environment, and do so in a way which represents badness in some way. In other

words, the evaluative properties are in some sense part of the contents of Norma's

affective experiences.

AI is most naturally paired with what I'll call Perceptualism about Emotions:

Perceptualism: Emotional experiences are a special sort of perceptual experience, in that they are pre-doxastic representational states that serve to *prima facie* justify beliefs about their contents.

It's widely believed, by both philosophers and psychologists, that emotions have

intentional objects.[8] But there are also well-known reasons to not identify emotions with

---

[7] Harman (1977), Ch. 1.
[8] See Fox (2008), Ch. 1, De Sousa (2013), Sect. 3, Goldie (2007).

beliefs.[9] These two features of emotions motivate a broadly Perceptualist view. For

some Perceptualists, emotional experiences are literally perceptions, and the affective

system is a perceptual module to be placed aside the visual, aural, and other perceptual

systems.[10] For others, emotional experiences are not literally perceptions, but are

epistemologically analogous to perceptions in that they can serve to justify beliefs in

their contents in the way, e.g. visual experiences, can.[11]

In what follows, I remain neutral between literalist and non-literalist versions of

Perceptualism. What is relevant at present is the epistemological aspect of

Perceptualism, which literalists and non-literalists agree about: emotional experiences

have (intentional) contents, and these contents can serve to *prima facie* justify beliefs. The

pairing of AI with Perceptualism is, I hope, relatively straightforward. While neither

Perceptualism nor AI entail each other, each one naturally fits with the other into a

cohesive overall affective epistemology.

2. The Rational Criticizability Problem

The Rational Criticizability Problem arises from the mutual inconsistency

between Perceptualism and two other seeming truisms about emotions and perceptual

states:

---

[9] See, for example, Brady (2014), Ch. 2, De Sousa (2013). Though these arguments are more controversial—see, e.g. Nussbaum (2001), Ch. 1.

[10] Johnston (2001), Döring (2003), Prinz (2004).

[11] De Sousa (1984), Roberts (2003).

> *Criticizability*: We can be legitimately criticized for our affective responses.

> *Perceptual Non-Rationality*: Perceptual states can be veridical or not, but they are not subject to rational or moral criticisms (charges of irrationality, charges of unreasonableness, etc.).

First, a word about Criticizability. Although I am calling the problem that arises here the Rational Criticizability Problem (RCP), I intend "Rational" to be read in a very broad way, to encompass criticisms of emotions as irrational, unreasonable, or even immoral. The kinds of criticisms that legitimately apply to emotional states fall under a variety of guises, and Criticizability is intended to capture all of them. (I discuss examples below.)

Perceptual Non-Rationality is also intuitive. Unlike someone who is insensitive to the feelings of others, we don't think that someone who is red-green color blind is thereby subject to rational or moral criticism. Similarly, while I may be criticizable for continuing to be duped into believing it, I am not criticizable because one line continues to look longer than the other in the Müller-Lyer Illusion after I know better. Besides taking actions such as plugging my ears or closing my eyes, perceptual experiences are not subject to rational control in the right kind of way to be criticizable. They are passively exercised faculties, unlike paradigmatically criticizable states such as believing or intentionally acting.

The tension here for Perceptualism should be clear: Perceptual states are not rationally criticizable. Affective states are criticizable. So affective states are not

perceptual states. The problem has been approached from a few different directions.

Peter Goldie discusses the worry in terms of responsibility:

> Intuitively, emotions are passive, and often seem to overcome us, sometimes in spite of our efforts. These considerations suggest that we cannot be directly responsible for our emotions…However, consideration of other examples might suggest otherwise: if I feel a deep loathing of foreigners, I might reasonably be thought to be directly responsible for this, just as much as I am for the actions that I do which are expressive of that loathing.[12]

Robert Cowan, on the other hand, levies the objection in terms of reasons: "Emotions can be justified or unjustified because they are, or ought to be, responsive to *reasons*— e.g. my acting lecherously at the party is a *reason* for me to feel guilty."[13] Finally, Mikko Salmela, in a discussion of recalcitrant emotions (such as a fear of flying even despite knowledge that it is safe), notes that

> [T]he fact that we regard many recalcitrant emotions as well as pathological emotions as irrational rather than *arational*, and try to get rid of them, implies that the problem with recalcitrant emotions is not so much whether they *need* to be revised in light of better knowledge, but rather whether they *can* be so revised.[14]

In short, the RCP takes many forms. Perhaps it is that we are responsible for our emotions but not for our perceptual states. Alternatively, perhaps it is the fact that we can have reasons for or against having certain emotional states, but not for perceptual states. Or it may be a claim about the rational assessability of emotional states. There are

---

[12] Goldie (2007), 932-933.

[13] Cowan (forthcoming-a), 12.

[14] Salmela (2011), 15. In this passage, Salmela is responding to Döring's (2009) discussion of recalcitrant emotions, but I take it he intends his remarks to generalize to all perceptualist views.

probably other ways to think about the problem too, some surely reducible to others. What is important for the RCP to get off the ground is that there is some important and seemingly normative difference between emotional experiences and perceptual experiences: The former are criticizable in some important respect, while the latter are not.

RCP is an objection to Perceptualism. But it should be especially worrisome to the proponent of Affectual Intuitionism. Recall that for AI, affective experiences play an analogous epistemic role with respect to evaluative beliefs that perceptual experiences do with respect to empirical beliefs. If they are to do this on a foundationalist view such as AI, they must provide *non-inferential* justification—that is to say, they must justify without being epistemically dependent on any other mental state for their justifying force. However, the rational criticizability of emotions suggests that they depend on something independent of them for their justificatory power, e.g. other mental states. For example, if, as Cowan says, someone can give me *reasons* to feel guilty, this suggests that emotions are only justified in virtue of something else, e.g. our ability to recognize reasons, or our belief that reasons to feel a certain way are present. The upshot is that, even if AI could be separated from Perceptualism, the RCP would arguably still be a problem for the proponent of AI.

3. Cases of Criticizability

Because Criticizability comes in different forms, it will be helpful to consider a few test cases that a response to the RCP should explain. The simplest case is one when someone's emotions lead them to overreact:

> *Overreaction.* Andy is cut in line at the coffee shop. He throws over a display rack and screams at the line-cutter before storming out.

Andy is criticizable for overreacting. However, he'd also be criticizable even if he didn't react so strongly:

> *Overfeeling*: Andy is cut in line at the coffee shop. Outwardly, he does nothing. Inwardly, he is so upset that he shakes with rage, fantasizing about punching the line-cutter, wishing ill upon his family, etc.

Both *Overreaction* and *Overfeeling* are cases where the strength of the emotion is inappropriate. Even if some limited level of anger is a proper response to someone cutting in line, the amount of anger that Andy feels outpaces what is appropriate.

A third case involves responding to certain features of a situation to the detriment of noticing other salient features. In these cases, one emotional response crowds out all others in a way that is criticizable. For example:

> *Selective Empathy:* A police officer shoots an unarmed Black teen. In the process of getting out of her squad car, she accidentally and painfully slams her fingers in the car door. Jane, who has witnessed the entire interaction, empathizes very strongly with the officer's pain and stress, while feeling little for the shot teen.

The issue here is not merely that Jane feels sympathy for the police officer, for her pain, mental and physical, is genuine.[15] Rather, it is that Jane's sympathy for the police officer is inappropriate given that another person has just been shot.

A final case worth considering is the case of phobias:

> *Phobia:* Doug is powerfully afraid of mice. He is so powerfully afraid that at times it affects his well-being—for example, he refuses to sleep in old houses for fear of them being infested with mice, he will not visit pet stores for fear of seeing mice, nor can he enjoy watching *Ratatouille* with his children. Despite this, he knows that mice are, for the most part, harmless, and that the situations he avoids for fear of mice are not dangerous.

As Salmela notes, phobias are called *irrational* rather than merely *incorrect*. This suggests that phobias involve some kind of rational failing, and are thus criticizable.[16]

4. The Complexity of Emotions

There are three (non-exclusive) options as to how to respond to the *Rational Criticizability Problem* (RCP). The anti-perceptualist has us reject perceptualism in light of Criticizability and Perceptual Non-Rationality (PNR). The perceptualist must reject either Criticizability or PNR. Each approach appears to be seriously revisionary. Giving up Criticizability seems to entail an error theoretic conception of our assessments of

---

[15] It could be plausibly claimed that the police officer here deserves no sympathy, given that her action in shooting the teen was likely seriously blameworthy. This gets into some thorny issues about sympathy/empathy for the blameworthy that are beyond the scope of this paper. The reader should feel free to substitute her favored case of selective empathy.

[16] Salmela (2011), 15.

emotional states. Giving up PNR entails that perceptual experiences are rationally criticizable.[17]

Rebutting the RCP doesn't require anything so bold. Rather, I think that once we carefully separate out the representational component of emotional responses from a variety of other features of emotional states, the force of Criticizability melts away. Criticizability is ambiguous between two readings; on a more general reading it is true but doesn't conflict with Perceptualism, and on a more narrow reading it is in conflict with Perceptualism but lacks support. Clarifying this requires a brief digression on some widely accepted facts about the emotional system.

Whether Perceptualism is true, the emotional system is complex and unique in the roles it plays within our broader cognitive system. Perceptualism has two components: First, it says that emotional experiences are non-doxastic representational states. Second, it says that these representational states serve to *prima facie* justify beliefs in an analogous way to perceptual experiences. Both components are compatible with emotional states functioning in other important ways. In fact, to deny the other roles that the emotional system plays would be to deny widely accepted psychological and philosophical facts.[18]

---

[17] Interestingly, Siegel (forthcoming) has argued that at least some perceptual experiences are rationally assessable. See also Watzl (2014).

[18] I have no ready-to-hand way to carve the emotional system off from other cognitive systems precisely. Broadly speaking, by "emotional system", I refer to the cognitive structures that underlie the generation of affective experiences, and I take this rough and ready characterization to be satisfactory for the claims made in this paper. (Thanks to an anonymous referee for pressing me to be clearer on this point.)

Other than the representational and epistemological features that emotional

states have according to Perceptualists, emotional states, or some component of the

emotional system,[19] play the following roles:

(a) *Motivation.* Emotional states directly cause us to act—that is, they cause us to act without intermediary work by beliefs or desires. The sight of a snake, for example, can cause someone to freeze in fear before any belief that there is a snake present is formed.[20]

(b) *Action-Priming.* As per *Motivation*, emotional states can directly cause us to act. But they can also play a weaker role, by priming us toward certain actions without causing us to act. A feeling of uneasiness about someone may prime someone to turn away, make less eye contact, and stop believing what that person is saying. Or again, a feeling of fear may in many cases prime us to action (freezing, running away, tensing up to prepare for impact, etc.), but not to any particular action.[21]

(c) *Attention Focusing.* Over and above representing their objects as having certain properties, emotions also serve to focus attention on the emotionally-laden features of a situation. This can occur even at the cost of neglecting relevant but not emotionally-laden features. Ben-Ze'ev compares this aspect of emotions to "burglar alarms going off when an intruder appears, emotions signal that something needs attention."[22] From a neurophysiological standpoint, what this comes to is the emotional system playing a role in selecting which information in a perceptual input is to be further processed and brought to conscious awareness.

(d) *Affective Style and Temperament.* Affective style and temperament are notions used by psychologists to describe tendencies to experience affect-laden situations as positive or negative. Individuals classified as neurotic will be more sensitive to experiencing situations as anxiety-inducing and potentially harmful (rather than

---

[19] Whether these features are realized in distinct physiological states from the (purported) representational features of emotion is an interesting empirical question. But what is important for the argument below is only that these things are conceptually separable, and it seems that they clearly are. (Thanks to an anonymous referee for pressing me on this point.)

[20] Fox (2008), Ch. 1.

[21] Fox (2008), Ch. 1.

[22] Ben-Ze'ev (2000), 13. See also Brady (2014), pp.20-22, and Faucher and Tappolet (2002), 114-115.

potentially rewarding), and thus (in light of (a) and (b) above) be more motivated to fear and other avoidance behavior. Two individuals with differences in temperament may both represent the emotionally-laden features of a situation accurately but with differences in the perceived salience of those features, while in other cases an individual with a particularly strong temperament may outright miss certain emotionally-laden features of a situation because of a preoccupation with other features of the situation.[23]

(e) *Character.* A philosophical aspect of the affective system that pairs with the psychological notions of affective style and temperament is the notion of character. Of course, the philosophical issues surrounding character are far too deep to even scratch the surface. What is relevant here is that, arguably, one's emotional dispositions play some important role in making up part of an individual's character. As Angela Smith puts it, "We react with…envy, admiration, resentment, awe, amusement, regret, and gratitude to the people and events we encounter…we regularly take these involuntary responses to have a great deal of expressive significance."[24]

Finally, besides all of the above roles that the emotional system plays in our broader

cognitive economy, there is also an important way in which our background cognitive

states can influence our emotional system:

(f) *Emotional Regulation.* Emotional Regulation (or Self-Regulation) involves an agent exercising cognitive control over the felt strength and motivational force of her emotion. Regulation can be external, as when someone with a fear of snakes avoids the snake exhibit at the zoo, or it can be internal, as when someone consciously directs their attention to something less emotionally abrasive to quell their anger.[25]

It's worth noting briefly that *Emotional Regulation* is wholly compatible with

Perceptualism. One powerful motivation for Perceptualism is that emotional

---

[23] See Fox (2008), Ch. 3 for an introductory discussion with references to the psychology literature.
[24] Smith (2005), 249.
[25] For discussion, see Gross (2008), Roberts (forthcoming).

experiences seem to be passive in a way that beliefs are not. It might seem that this claim about the passivity of emotions conflicts with *Emotional Regulation*. However, the sort of regulation that we have over our emotional experiences is not in conflict with passivity in the relevant sense. In external emotional regulation, an agent alters her emotional experience by consciously selecting her external inputs. Choosing to avoid the snake exhibit at the zoo is regulating one's emotions, not in the sense that one is actively and directly preventing some emotional experience; rather, the agent knows that she won't be able to avoid feeling fear in the snake exhibit *because her emotions are not directly in her control*, and so avoids the relevant situation which would result in her negative emotional experience. Internal emotional regulation works similarly, except that instead of avoiding the emotion-invoking experience altogether, an agent consciously redirects her attention as to avoid the emotional stimulus.[26]

There are analogous kinds of regulation in the traditional perceptual realm. If I hate the taste of mushrooms, I can avoid experiencing it by not eating mushrooms; if norms of politeness dictate that I do, I can distract myself from the taste of mushrooms by directing my attention to the conversation at hand. Neither of these instances of what we might call "Perceptual Regulation" cast doubt on the passivity of perceptual

---

[26] In cognitive reappraisal, an agent consciously and intentionally reinterprets a situation in an effort to suppress a negative emotion or replace it with a positive one (Gross 2002, 2007). Such a reinterpretation does not have a direct analogy to anything in traditional perceptual processing (though attentional shifts, such as in the duck-rabbit illusion, may be similar), because this reinterpretation takes place cognitively downstream, it remains compatible with the passivity of emotional experiences in a particular moment. (Thanks for an anonymous referee for stressing the importance of cognitive reappraisal as an important kind of emotional regulation.)

experience. There is no pressure on the Perceptualist to deny Emotional Regulation either.

It's clear that, Perceptualism notwithstanding, the emotional system does a lot of cognitive work. In section 5, I return to the cases discussed above in light of this fact. As we'll see, a closer look illustrates that the criticizability in each of the cases is not a criticism of the representational nature of the emotional experience, but a criticism of one or more of features (a)-(f) above.

5. Complexity and the *Rational Criticizability Problem*

With the background on the many distinctive (but related) roles that emotions play within our broader cognitive systems, we can return to what this means for the RCP. I begin by providing the response in the abstract before illustrating how this plays out with respect to the cases discussed in section 3.

Let's begin by reminding ourselves of *Criticizability*, one of the key principles underlying the RCP:

Criticizability: We can be legitimately criticized for our affective responses.

Especially as reinforced by the cases given in section 3, *Criticizability* seems undeniable. Rejecting it would be a tough bullet to bite for the Perceptualist. However, as stated, and contrary to the proponent of RCP, the Perceptualist need not deny *Criticizability* at all. Given the plural function of emotions in our cognitive systems, it may very well turn out that emotions are rationally or morally criticizable along many of these metrics,

compatible with their failing to be criticizable *qua* representation. To take just one example, suppose that Carl's annoyance directly causes him to yell obscenities at a child. His annoyance will be criticizable *qua* causing his immoral action. This is compatible with Perceptualism, since Carl's emotional representation of the child as annoying is not the thing which is rationally or morally criticizable. Only the motivational upshot of that emotion is under criticism.

What this shows us is that, if the RCP is really to put pressure on Perceptualism, *Criticizability* is too broad of a principle. Rather, we need to replace it with:

> *Criticizability*:* The representational aspect of our emotional experiences can be legitimately criticized.

Since the affective system is complicated, playing several different (related) roles at once, the narrower principle *Criticizability** is necessary for RCP to be a genuine problem for Perceptualism. The problem, though, is that *Criticizability** is much less intuitive than *Criticizability*, and the common sense sorts of cases used to motivate the latter don't obviously extend to the former. What the proponent of RCP would need are a class of cases where someone's emotional experience are being rationally or morally criticized in a way not subsumable to criticisms of the other features of the emotional system (a)-(f). It is difficult to see what these cases might look like.

In the abstract, once we get clear about what kind of criticizability is really required to cast doubt on Perceptualism, the RCP loses its original intuitive force. But it is worthwhile to see how this plays out in particular cases, such as those discussed in

section 3. Begin with *Overreaction* and *Overfeeling*. In *Overreaction*, recall, Andy explodes

into a violent rage as a result of someone cutting in front of him in line. In this case,

Andy is surely criticizable, morally and probably rationally as well. However, it seems

clear that the criticism is directed at Andy's emotion only insofar as it caused his

immoral and irrational behavior. So here it is the motivational component of Andy's

emotional response that is under censure.

In *Overfeeling*, on the other hand, Andy does not act immorally or irrationally.

Instead he stews in his rage, thinking terrible things and shaking with anger, but

outwardly showing a calm demeanor. While in this case there is no direct action of

Andy's to trace our emotional criticism back to, there seem to be multiple other

candidate explanations of the target of our criticism of Andy here. First, insofar as his

anger is priming him to take drastic action, his motivational dispositions, even if

currently suppressed, seem like legitimate targets for criticism. Relatedly, insofar as

Andy doesn't seem to be distancing or repudiating his rage, instead dwelling on it, he'll

be criticizable from the standpoint of emotional regulation. In this instance, Andy can

choose to regulate and distance himself from his anger by thinking about something

else, finding distractions or otherwise avoiding the sorts of situations that predictably

trigger his rage. Instead, Andy has chosen to endorse at least the felt strength of the

emotion, even encouraging it (both now and in future situations which are similar in kind). For that he is criticizable.[27]

Selective Empathy, recall, involves Jane's strong empathy for the police officer to the expense of having very little sympathy for the unarmed teen. In this case, it's not that her empathy for the police officer is misrepresenting *per se*; rather, it's that her empathy is inappropriate in these circumstances since it crowds out her ability to empathize with the unarmed teen, where empathy would be more legitimately placed in this situation. If we suspect that Jane is implicitly or overtly racist, our criticism of her emotional dispositions in this case could first and foremost be seen as a criticism of her character. If this is not a fluke emotional reaction on Jane's part, but rather represents a pattern, her character seems legitimately criticizable. But even supposing that this is a fluke, Jane may still be criticizable. But the most plausible candidate for criticism here, it strikes me, is that her emotions have shifted her *attentional focus* in a problematic way. Again, notice that the problem with Jane isn't that she represented the police officer as deserving empathy when she doesn't—as I've stipulated the case, the police officer does deserve empathy. So she isn't misrepresenting. Rather, she's *overlooking* crucial features of the situation, and she is doing so because her empathy for the police officer has given

---

[27] We may also criticize Andy from the standpoint of his overall emotional temperament, if he is prone to these sorts of overreactions frequently. But I don't want to belabor the point.

her emotional tunnel vision of sorts. So again, Jane doesn't appear to be criticizable for the representational aspect of her emotion.

The final case study to consider are phobias such as Doug's in *Phobia*. Now, first, I should flag that it is unclear to me that people really are criticizable for having phobias. Surely phobias can be detrimental, even destructive to our well-being. But that merely shows that they are pervasively incorrect, not that they are subject to blame. In defense of the claim that emotions like phobias are criticizable, Salmela points out that "the fact that we regard many recalcitrant emotions…as irrational rather than *arational*, and try to get rid of them, implies that the problem with recalcitrant emotions is not so much whether they *need* to be revised…but rather whether they *can* be".[28] As for the linguistic point, that we label phobias as "irrational" rather than "arational", I don't think this provides much evidence one way or another. What of the point that we all accept that recalcitrant emotions ought to be revised? It is true that we often have emotional dispositions that we want to alter, judging them to be prone to error. But this is no different from perception. If I could easily retrain myself to not fall victim to the Müller-Lyer illusion, without thereby causing other potential visual errors, I would do so.

Emotion may come apart from perception in two ways here: First, insofar as emotion is tightly linked to normative judgments and action, it may be more

---

[28] Salmela (2011), 15.

normatively important to correct dispositions to emotional errors than dispositions to

(traditionally) perceptual errors. Second, as a matter of empirical fact, our emotions

may be more malleable than (traditionally) perceptual systems. This isn't to say that we

directly control our emotions—rather, we can indirectly affect our emotional

dispositions with time and practice. But neither of these differences casts doubt on

Perceptualism; first, these are only differences in degree, not in kind, and second,

neither difference entails any *epistemic* difference between emotions and traditional

perceptual systems, which is the heart of Perceptualism (as I'm understanding it here).

Suppose that all of what I've said in the last paragraph is mistaken, and that

phobias are rationally criticizable after all. Insofar as Doug is criticizable for his fear of

mice, it seems as though the criticism must be of his ability to regulate this fear

adequately. For the fear is, qua phobia, recalcitrant and not in his control. What is in his

control—or at least is presupposed to be, if the criticism is to hit the target—is his ability

to suppress his fear in order to, e.g., watch *Ratatouille* with his child. So again, it appears

as though it is Doug's *Emotional Regulation* (or lack thereof) that is criticizable, not the

representational aspect of his fear. (Of course, his fear is misrepresenting the world; that

is common ground between the Perceptualist and her opponent.)

6. Conclusion

There are *prima facie* good reasons to endorse Perceptualism about emotions,

both for the proponent of AI as well as for independent reasons. On the other hand, the

criticizability—moral, rational, and prudential—of emotional states looks to be in

conflict with modelling emotions on perceptual states. This criticizability is pervasive in

our interactions with others, so denying it would seem radically implausible. However,

I have argued that, once properly understood, the criticizability in question is actually

perfectly compatible with Perceptualism. If what I've argued is correct, we can preserve

both the common sense intuitions about the criticizability of emotions as well as the

well-motivated perceptualist theory of those states.

Chapter 7

Toward a Perceptual Solution to Epistemological Objections to Non-Naturalism

Stance-independent non-naturalist moral realism[1] is subject to two related

epistemological objections. First, there is the metaethical descendant of the Benacerraf

problem in philosophy of math.[2] Second, there are evolutionary debunking arguments.[3]

How best to understand these two arguments, as well as whether one problem is

reducible to the other, both remain controversial.[4] I have defended the claim that each

of these arguments is best understood in terms of claims about our epistemic access to

the moral properties in chapter 1. And, as is no surprise, the success of these arguments

is contentious as well. In chapters 2 and 3, I have argued against several realist attempts

to resolve the problem.

These epistemological objections against moral realism are, perhaps,

unsurprising. What is surprising is that, despite many other differences, the most

widely discussed attempts to solve these epistemological problems have done so

without appealing to any particular moral epistemology. To some extent, this makes

---

[1] Shafer-Landau (2005, p.15) defines *stance independence* as the claim that "the moral standards that fix the moral facts are not made true by virtue of their ratification from within any given actual or hypothetical perspective." "Non-naturalism" is also subject to different understandings. I will not define "non-naturalism" precisely for the purposes of this paper, except to say that on a non-naturalist view, the moral facts are not identical or reducible to natural facts. (This conflicts with the epistemological characterization of "non-naturalism" which Shafer-Landau favors.)

[2] Benacerraf (1973).

[3] The two most famous examples of evolutionary debunking arguments are those of Joyce (2001, Ch.4, 2006) and Street (2006).

[4] See, for example, Enoch (2011, 7.1-7.3), Vavova (2015), and Clarke-Doane (forthcoming), amongst many others.

sense: A response to these particular epistemic concerns that is otherwise neutral about surrounding epistemological issues would be preferable to one only available to those willing on to take on other epistemological commitments. On the other hand, I think that the focus on these epistemologically neutral responses leaves many interesting theoretical stones unturned. Exploring the ability of particular theories in moral epistemology for handling these difficult epistemological objections can help to illuminate strengths or weaknesses within these theories themselves, as well as opening up potentially unexplored avenues for responding to deeply entrenched concerns about our epistemic access to the moral properties.

This chapter can be seen as a case study in the latter kind of project. In particular, I assess the prospects of an empiricist and perceptualist model of moral knowledge, *A Posteriori* Ethical Intuitionism, for responding to the Benacerraf problem and evolutionary debunking arguments against non-skeptical moral realism. I argue that AEI has powerful responses to these objections which are not open to other moral epistemologists. Furthermore, the uniquely perceptualist responses are arguably *more* compelling than the various other approaches to EA that have cropped up in the literature. The upshot is that *if* some version of AEI is correct, then the realist has less to fear from Benacerraf and evolutionary debunking style epistemological objections. Insofar as one is already a committed realist, then, this provides some indirect support

for AEI. The more general suggestion, however, is that non-naturalists look to particular theories of moral justification in responding to skeptical arguments.

The structure of the chapter is as follows. In section 1, I discuss the origins of, and relationships between, the Benacerraf Problem and Evolutionary Debunking Arguments against non-skeptical moral realism. In section 2, I begin by briefly sketching the perceptualist position in moral epistemology. In sections 3-5, I tease out three distinct ways of understanding these sorts of epistemological worries—rather than engaging in exegesis, I formulate what I take to be the three most powerful epistemological problems in the area. (I leave it to evolutionary debunkers to argue that there is some stronger way of understanding the objection that I have overlooked.) In the remainder of the paper, I argue that AEI provides a framework for responding to each of these three distinct epistemological worries. (A full defense requires a brief digression into the philosophy of perception, in section 8.) A non-naturalist proponent of AEI, then, seems to be able to handle these traditional epistemological problems. Whether other substantive non-naturalist epistemologies can do the same remains to be seen—but at the very least, I hope this case study in AEI encourages further thought about this method for responding to epistemological concerns targeted at non-naturalists.

1. Benacerraf's Problem and Evolutionary Debunking Arguments: A (Very) Brief History

As discussed in chapter 1, the so-called Benacerraf problem goes back at least to Paul Benacerraf's (1973) "Mathematical Truth". As its title suggests, Benacerraf was worried about the epistemology of mathematics within a Platonist conception of mathematical objects. The gist of Benacerraf's argument is straightforward. Benacerraf (at least in 1973) accepts a causal constraint on a theory of knowledge. On a Platonist conception, mathematical objects have no causal power. So if Platonism were true, mathematical knowledge would be impossible. An obvious initial worry here is that causal constraints on knowledge (or justification) have largely fell out of favor.[5] But the argument has remained influential, because even if the problem isn't the causal inefficacy of Platonic objects *per se*, there does seem to be a serious mystery about how we could have any epistemic access to abstracta of that sort, causal or not.[6]

While Benacerraf was concerned with the case of mathematical knowledge, a similar sort of problem arises for non-naturalist moral realists, according to which moral properties are not causally efficacious.[7] This parallel was first noticed by Gilbert Harman (1977). And again, as in the mathematical case, even while causal conditions on knowledge have fallen out of favor, there is still a widespread sense that explaining our epistemic access—in terms of responsiveness to the moral facts—is a serious concern for

---

[5] For an overview and references, see Jenkins Ichikawa & Steup (2012).

[6] Field (1989), Cheyne (2001), and Liggins (2010) for discussions of this point and developments of Benacerraf's problem.

[7] E.g. Heathwood (2015, p.3), McGrath (2014, p.186), and Scanlon (2014). Oddie (2005) is an exception to this general rule, advocating a version of non-naturalism according to which the non-natural properties are causally efficacious.

the non-naturalist moral realist. Part of the challenge here is just to understand what kind of explanation is being called for. While there is certainly dispute about that, I think it's overwhelmingly plausible that there is something of real concern in the ballpark of Benacerraf's original worry. Non-naturalists shouldn't be satisfied to wait until a clear consensus emerges on the question of what epistemic access comes to before giving a positive epistemological story which assuages these concerns.

On the other hand, a more specific sort of undercutting defeater for moral knowledge comes in the form of evolutionary debunking arguments (EDAs). Evolutionary debunking arguments are a relative of Benacerraf-style objections. Though they have almost certainly been around much longer, interest in EDAs has skyrocketed largely as a result of the work of Sharon Street (2006) and Richard Joyce (2001, 2006). The details of Street and Joyce's (as well as other debunkers) EDAs differ in important ways. At the most abstract level, the idea of an EDA is to provide a genealogy of our moral beliefs or belief-forming mechanisms in terms of the fitness-enhancing evolutionary nature of those beliefs or belief-forming mechanisms. Such an explanation, the debunker argues, will make no appeal to any mind-independent moral facts. This tells us that the moral beliefs or belief-forming mechanisms that we have are, given the wide variety of (conceptually) possible moral systems, extremely unlikely to be tracking the moral facts. This knowledge of almost certain unreliability serves to undercut any justification we initially had about stance-independent moral facts.

In the above, I've provided nothing more than the briefest of sketches of Benacerraf arguments and EDAs. But I think it's clear that these challenges are, if not the same, at least closely related. Both of them are tightly connected to two ideas. First, if the explanation of our moral beliefs or moral belief-forming mechanisms do not invoke some connection to the moral facts, then we should be skeptical about our prospects for moral knowledge. And second, if moral facts have the metaphysical status that non-naturalists believe that they do, then it is difficult if not impossible to see how an explanation of our moral beliefs or belief-forming mechanisms could invoke the moral facts.

Such is the general outline of the two challenges, but we'll need to precisify before we can fully assess their chances at success. In the next section, I provide three potential ways of precisifying these sorts of challenges. Each precisification is distinct and compelling enough that it requires a distinct response, which I attempt in sections 4-6. But first, I recap the AEI, which will be doing some of the heavy lifting below.

2. *A Posteriori* Ethical Intuitionism: An Outline[8]

AEI consists in two substantive claims. First, AEI is a version of *Ethical Intuitionism*:

> EI: "[N]ormal ethical agents have at least some non-inferentially justified first-order ethical beliefs."

---

[8] This section is largely overlap from my "A Posteriori Ethical Intuitionism and the Problem of Cognitive Penetration" [Ch. 5].

As stated, EI is just the claim that foundationalism—understood in the epistemologist's sense—is true of the structure of at least some ethical beliefs, and that some ethical beliefs are members of the set of foundational beliefs. However, EI does not entail what we might call *Classic Intuitionism*, the claim that the non-inferentially justified ethical beliefs are grounded in *intuitions*, whatever those turn out to be. EI has its unfortunately misleading name presumably because most defenders of EI have also been defenders of Classic Intuitionism. I continue to label EI as such in order to accord with what has now become standard usage in the literature. The second claim that constitutes AEI is *Ethical Empiricism*:

> EE: The non-inferential justification of first-order ethical beliefs is grounded in the instantiation of evaluative properties in some kind of perceptual experience.

According to EE, non-inferentially justified ethical beliefs are justified analogously to non-moral perceptual justification, on one sort of traditional foundationalist view. The basic picture is as follows. Under certain circumstances, evaluative properties (or their instantiations) figure in the contents of perceptual experience. Furthermore, at least sometimes, the evaluative properties that figure in the contents of perceptual experience are capable of providing non-inferential justification for beliefs about the instantiation of evaluative properties. This is compatible with the claim that sometimes evaluative perceptual experiences fail to non-inferentially justify. First, there may be defeaters for the justification that an evaluative perceptual experience would otherwise provide. Second, some evaluative perceptual experiences may be epistemically dependent in the sense that they can't provide justificatory force independently of some prior evaluative belief. AEI only claims that, in at least some circumstances, neither of these things holds. When they do not, an evaluative perceptual experience can ground a non-inferentially justified moral belief.

AEI is not a single view but a family of views. I briefly note two recently discussed versions of AEI. According to *Perceptual Intuitionism*, evaluative properties are perceived in the same way that any non-moral but otherwise high-level properties are perceived. For example, the Perceptual Intuitionist may contend that the property *goodness* is capable of triggering a recognitional disposition within the visual processing system. On the other hand, according to *Affectual Intuitionism*, evaluative properties are perceived via affective states, which are themselves construed as either a form of perception or as a submodule within a broader perceptual (as opposed to cognitive) system. AEI is compatible with a number of views about the metaphysics of moral properties. However, as the focus in what follows is on epistemological objections to non-naturalist moral realism, I assume that version of AEI throughout.

3. Formulation 1: Epistemic Access is Impossible

The first formulation of the epistemological problem, and perhaps the boldest, revolves around the claim that, given the metaphysical status of the non-natural properties, we could bear no justification conferring relationship to them. As Bengson says in a similar context, "What relation does a thinker's mental state…bear to an abstract fact that explains how the state can be non-accidentally correct with respect to that fact, and hence able to serve as a source of knowledge for it?"[9] In short, epistemic access to the moral properties is impossible. Call this the *Impossibility Argument* (IA):

> 1. Any complete non-skeptical epistemology for a particular domain *d* must explain how we have epistemic access to the *d*-facts such that we are able to form beliefs that are responsive to those facts.[10]
>
> 2. There is no explicable mechanism to explain how human beings could have epistemic access to non-natural moral facts such that our beliefs are responsive to those facts.

Therefore,

---

[9] Bengson (forthcoming), 8.
[10] Notice that this is the "access requirement" (AR) discussed at length in Ch. 1.

3. Non-skeptical non-naturalist moral realism is false.

Premise 1 is an attempt to reframe Benacerraf's causal requirement on knowledge in non-causal terms. Epistemic access may be causal, but it may also be through one of several other routes: Introspective, Conceptual, Intuitive, or Constitutive.[11] The rough idea here is that, causal or not, the believed fact should be playing a role in explaining why the belief is held.[12] So premise 1 is an attempt to preserve the intuitive force behind Benacerraf's argument without begging any questions against the non-naturalist. Premise 2 merely states that there is no such epistemic relationship between human agents and the non-natural moral facts to provide epistemic access of the sort discussed in premise 1.[13] So, assuming that premise 1 is correct, the challenge for the non-naturalist is to provide a positive account of what such an epistemic relationship might look like.

4. Formulation 2: Reliability Would Be *Extremely* Unlikely

The proponent of the second formulation of a debunking argument is grants the possibility of epistemic access (at least for the sake of argument). But she claims that this is not enough to rebut skepticism, since it remains extremely unlikely that our actual moral beliefs are latching onto the moral facts. This seems to be what Sharon Street has

---

[11] [See Ch. 1]

[12] For more on this, see Jenkins (2008, Ch.3).

[13] Hartry Field (2005, 77) gives a powerful statement of the sentiment behind premise 2: "[Non-naturalism] postulate[s entities] that are mind-independent and bear no causal or spatiotemporal relation to us, or any other kinds of *relations* to us that would *explain* why our beliefs about them tend to be correct".

in mind when she says that "Allowing our evaluative judgments to be shaped by

evolutionary influences is analogous to setting out for Bermuda and letting the course

of your boat be determined by the wind and tides".[14] Just as it would be *a priori*

extremely unlikely that our boat would make it to Bermuda with no guidance, it would

be extremely unlikely—given all the conceptually possible ways the moral truths could

turn out—that our evolved moral beliefs latched onto the moral facts.

Call this the *Extreme Unlikelihood Argument* (EUA):

1. There is a huge range of logically possible stance-independent moral truths (and systems of truths).

2. Our actual moral judgments are not caused by the moral facts.[15]

Therefore,

3. It is extremely unlikely that a significant subset of our moral beliefs are true, and those that are true are true as a result of luck in a way that undermines their justification.

EUA, as stated, is not logically valid. But it remains a powerful inference to the best

explanation, and so merits some sort of response on behalf of the non-naturalist.

5. Formulation 3: No Good Reason to Think We're Tracking the Moral Facts

---

[14] Street (2006), pp.121-122.

[15] In Street's argument, evolution plays the central role in the genealogy of our moral judgments, but a proponent of EUA could also appeal to culture, individuals' upbringing, or some combination of these three things. Regardless of the details here, most non-naturalists would accept premise 2, so I do not distinguish between these possibilities.

The third and final formulation of the epistemological objection to non-naturalism that I'll consider relies on something like the following plausible epistemic principle:

> No Good: "If you have no good reason to think that your belief is true, then you cannot rationally maintain it."[16]

If No Good is correct, and our moral beliefs are caused by something other than the moral facts, it looks as though we have a defeater for most if not all of our moral beliefs. Take some actual moral belief that we have, such as that *suffering is bad.* The proponent of No Good points out that we have no good reason to think this belief is true—after all, whatever it was caused by, it presumably wasn't caused by the fact that *suffering is bad* (even if it in fact is). So we can't rationally maintain it. A similar argument would apply to any moral belief you choose, so skepticism seems to follow.

Now, the anti-skeptical non-naturalist may try to forestall this sort of argument by pointing out that we do have good reason to think many of our moral beliefs are true. And she can claim this by appeal to any of a number moral belief-forming processes: Her beliefs may be based on moral intuitions, or moral perceptions, acquaintance relations with the moral facts, etc. And she can say that her preferred theory about what grounds our moral beliefs does give her good reason to think they're true after all.

---

[16] This is Vavova's (2014) formulation. Street (forthcoming, p.2) endorses a similar but more sophisticated principle.

However, this kind of response on behalf of the anti-skeptical realist only serves to move the bump in the rug. That's because insofar as No Good is an attractive epistemological principle, so too is:

> No Good Process: If you have no good reason to think that a given belief-forming process generates true beliefs, you cannot rationally maintain beliefs formed on the basis of such a process.

Even granting the anti-skeptical realist's initial response above, then, she'll still owe us a good reason to think that her favored moral belief-forming process(es) generate true beliefs a significant portion of the time. And this is no easy task.

Putting this together, we get what I'll call the *No Reason Argument* (NRA):

> 1. If you have no good reason to think that a given belief-forming process generates true beliefs, you cannot rationally maintain beliefs formed on the basis of such a process.
>
> 2. We have no good reason to think that moral belief-forming processes generate true beliefs.

Therefore,

> 3. We cannot rationally maintain our moral beliefs (i.e. they are unjustified).

Notice the difference between NRA and EUA. Unlike EUA, NRA doesn't depend on any claim about the logical space of possible moral facts or the prior likelihood that we would latch onto the moral facts. It only depends on the claim that we've got no (good) non-circular reason to think that our moral belief-forming processes are reliable. This, along with the plausible claim that lacking such a reason undermines any claim to

justification (or rational belief), is enough to undermine justification to most, if not all, of our moral beliefs.

6. Two Epistemically Neutral Responses to Epistemological Objections (and Their Problems)

Standard realist responses to epistemological arguments like those above have so far tried to remain ecumenical between different moral epistemological views. This makes sense, since it allows the realist to remain otherwise neutral about substantive positive moral epistemology, and thus available to as many anti-skeptical realists as possible. I briefly consider two approaches of this sort, before turning to a discussion of AEI in particular.

A first strategy of this kind is the familiar 'partners in guilt' strategy.[17] In this context, the idea is to argue that these arguments must fail, since we clearly have knowledge in other domains where similar arguments could be given. For example, it's clear that we have mathematical and logical knowledge, but mathematical and logical facts plausibly have a similar ontological status to moral facts, non-naturalistically construed. Metaphysical facts, such as claims of essence, may have a similar ontological status. But we surely have at least mathematical and logical knowledge, if not metaphysical knowledge as well. So whatever other epistemological problems there

---

[17] See, e.g. Brink (1984), Shafer-Landau (2005, 209-211), and Tropman (2008). Some also refer to this as the "partners in innocence" strategy. Shafer-Landau (2005, 209) calls them "partners in crime".

might be for non-naturalism, it cannot be the lack of epistemic responsiveness, luckiness, or a principle like No Good Process.

The partners in guilt strategy assumes virtually nothing about what a positive moral epistemology must look like. Nor does it make any assumptions about the genealogy of justified moral beliefs. In that sense, the partners in guilt strategy has the advantage of being open to a wide range of otherwise diverse moral epistemologists.

On the other hand, the partners in guilt strategy relies on two assumptions which are, I think, both problematic. A partners in guilt strategy is only going to be successful insofar as the partners in question really are (a) uncontroversially epistemically innocent as well as (b) metaphysically analogous. The first problematic assumption that those who appeal to partners in guilt in this context make is that (a) and (b) are met with respect to the domains of math, logic, and metaphysics. The partners in guilt strategy is especially ironic in this context, since the arguments presented above are arguably descendants of analogous arguments given against Platonism in mathematics. Similar concerns have been raised by philosophers of logic[18] and metaphysicians[19] about those respective domains. These concerns are not only taken seriously, but have motivated many to less ontologically robust conceptions of these domains.[20] In short,

---

[18] Cheyne (2001), Liggins (2010), Schechter (2010).
[19] Kornblith (2007), Bengson (forthcoming).
[20] See, for example, Field (1989), and, in the ethical case, Street (2006, 2008). Note that such a move preserves our claim to knowledge in a given domain, and so avoids the skeptical conclusion.

the partners in guilt strategy in this context is unconvincing because the partners *are*

(arguably) guilty.

A second problem for the appeal to partners in guilt here is that it must also turn

out that the partners in question genuinely fail to meet the epistemic criterion laid out

in the three formulations above. Some story about the justification in these contested

domains must exist, if skepticism is false. Proponents of the epistemological analogy

between morality and these other domains assume that a successful epistemology for

any one of them will be easily extendable to each of the others. Such a "one size fits all"

theory of the traditionally *a priori* domains would be nice, but it is non-trivial whether

any such approach is feasible. The partners in guilt response to the epistemic access

argument given above assumes without argument that some "one size fits all"

epistemological theory of these domains exists.

A second and recently popular strategy for responding to epistemological

objections to non-naturalism are so-called "third factor", or "pre-established harmony"

explanations.[21] Proponents of this strategy first argue that an explanation of the truth of

our moral beliefs is enough for moral knowledge, even if those beliefs are not to be

explained in terms of the facts they are about. Then, second, they attempt to provide an

explanation as to why many of our actual moral beliefs are true. For example, Enoch

points out that there's a good evolutionary reason why the belief that *survival is good* is

---

[21] Enoch (2011, Ch.7), and Skarsaune (2011) both take this approach.

adaptive, so it's no coincidence that we happen to have such beliefs. And, since it turns

out that survival *is* good, our belief is true.[22] Such a third factor story gives an

explanation of reliability without positing any positive relation between us and the non-

natural normative facts, so it provides a way of rejecting the arguments above.[23]

One way to respond to third factor explanations is to consider them on a case-by-

case basis. For example, one might deny that the belief in question (e.g. *survival is good*)

is adaptive, or one might deny that the adaptive belief in question is true (e.g. survival,

it turns out, is not good). I want to claim that such a piecemeal approach is unnecessary,

because a mere explanation of reliability is not enough to underwrite positive

justificatory status. To see why, consider:

> *Genetic Luck.* As a result of a very strange genetic mutation, Teresa, who is
> otherwise epistemically normal, is born with the unshakeable belief that
> *Water=H₂O*. Teresa has no independent evidence that this belief is true; she just
> can't help but believe it. Her disposition to form this unshakeable belief is
> essential to Teresa, insofar as it is wired into her DNA.

Teresa believes that *Water=H₂O* at all and only those worlds in which it is true and she

exists. Her belief is not just reliable, but (to borrow a term from Hawthorne (2002))

hyperreliable. We thus have an explanation of the reliability of her belief. Still, it seems

---

[22] Enoch (2011), pp.168-177. More carefully, Enoch points out that his particular appeal to "survival is good" as the particular normative claim that harmonizes with our evolutionarily ingrained beliefs may not be correct. But so long as there is *some normative fact or other* that aligns with our evolutionary dispositions, the general strategy will succeed. See Enoch (2011), 171.

[23] To be clear, Enoch doesn't understand the epistemic challenge to non-naturalism as IA. However, if his third factor response is successful, it may also undermine IA. Enoch himself doesn't understand the epistemic challenge in accordance with IA, since he thinks that anything like premise 1 is mistaken (2011, 152-153). [I defend premise 1 in Ch. 1.]

wrong to say that Teresa is justified. Arguably, this is because neither Teresa's belief, nor her method for forming it, have anything to do with the fact that water is H_2O.

The pre-established harmony between Teresa's belief about water and the facts about water does not, then, seem to be enough to underwrite justification. But the relationship that holds between Teresa's belief and the facts about water are analogous to the relationship between our moral beliefs and the moral facts according to proponents of third factor rejoinders to IA. The purported adaptiveness of certain moral beliefs on such accounts are not explained by their truth—it is just a coincidence that the moral truths happened to also be the things that were adaptive to believe.[24]

While I do think this issue is a serious and often overlooked one for proponents of third factor responses to debunking arguments, there is surely more to be said. However, I propose to set further discussion aside and turn to a distinct approach to these epistemological objections. The strategy is to assess the prospects for responding to the epistemological objections from within a particular moral epistemology, AEI. If successful, and assuming AEI is independently plausible, this would provide a response to these epistemological objections and, in doing so, provide a further reason to support AEI, namely, its ability to explain epistemic access to the normative facts. I'll consider each of the three versions of the epistemological arguments given above in turn.

---

[24] For further discussion of this and related issues, see [Chapter 1].

7. *A Posteriori* Ethical Intuitionism and the Impossibility Argument

Recall the second premise of IA:

> 2. There is no explicable mechanism to explain how human beings could have epistemic access to non-natural moral facts such that our beliefs are responsive to those facts.

The challenge, for AEI or any other positive moral epistemology, is to explain how epistemic access to non-natural moral facts is possible. It may initially seem like the proponent of AEI has a simple but complete answer: Our access is perceptual. We perceive the non-natural moral facts, and after all, there is nothing implausible about claiming that perception can provide access to mind-independent properties. Insofar as premise two of IA is false with respect to other perceivable properties—tables, cats, shapes, etc.—it will also be false for non-natural moral properties. So, according to AEI, the non-natural facts are epistemically accessible, and IA fails.

This response may sound too good to be true, and that's because it is. While it's true that, according to AEI, our access to the non-natural facts is perceptual, just how this is possible is much more unclear than it is in the case of tables, cats, or shapes.[25] This is because, unlike tables, cats, and shapes, the non-natural properties are widely thought to be non-causal.[26] The non-naturalist may, of course, reject this aspect of the view. But setting that aside, it seems as though the quick and dirty response given above just pushes the problem of epistemic access back a step. The skeptic can now ask: How could we have perceptual access to a causally inefficacious property, when perception is essentially a causal relationship? Insofar as this is impossible, the spirit of IA retains its full force.

The proponent of AEI could deny that perception is essentially causal, but without further motivation, this would appear *ad hoc*. She could also deny that non-natural properties are

---

[25] Schroeder (2008, Sect. 4) raises a similar worry for the view that desires are appearances of the good.
[26] See n.7.

causally inefficacious, but that would raise its own problems.[27] It might seem that these two

options, both unpalatable, are the only routes available for the proponent of AEI. And so it may

look as if, initial appearances aside, AEI is not well placed to provide an adequate account of our

epistemic access to the non-natural properties.

What we have, then, is a seemingly inconsistent triad:

Causally Inefficacious (CI). Non-natural moral properties are causally inefficacious.

Perceptual Access (PA). We have epistemic access to non-natural moral properties
through perception.

Causal Condition on Perception (CCP). Perception is an essentially causal relation.

As we've just seen, we should be reluctant to give up either CI or CCP. Rejecting PA, and thus

accepting premise 2—and thus the soundness—of the Impossibility Argument, appears to be the

only option left. But this isn't right. Contrary to initial appearances, the triad above is not

inconsistent. We can simultaneously accept CI, CCP, *and* PA—or so I presently argue.

The appearance of inconsistency arises because perception is essentially causal, whereas

non-natural moral properties are non-causal. However, once we focus on what precisely CCP

says (and doesn't say), it becomes clear that CCP is actually compatible with moral perception,

and thus compatible with PA, even if the moral properties are non-causal. To see this, notice that

the proponent of *a posteriori* ethical intuitionism needn't—and in fact *shouldn't*—deny that

moral perception is causal. If Norma perceives that Tibbles's being lit on fire is bad, and this

perception isn't hallucinatory, she surely must stand in some causal relation to Tibbles.[28] Thus,

Norma's perception is essentially causal; CCP is met. And yet her perceptual experience

represents *badness*, a causally inefficacious property; so we haven't given up CI either.

---

[27] For this strategy, see Oddie (2005).

[28] Emily surely is causally related to Tibbles, but is she causally related to the *badness* being instantiated? Arguably not. See sections 9 and 10 for a response to this worry.

Again, it would be fair to worry that there is some implicit sleight of hand going on—can

epistemic access really be this easy? In a sense, I think the answer is yes, in that I think the

explanation just given is correct. But it's also frustrating, because it seems to respond to the letter

of the epistemological concern without adequately addressing the spirit. Since I'm worried about

the spirit of the objection as well, it is well worth saying a bit about what the underlying

philosophy of perception issues that I think are going on beneath the surface.

8. Interlude: Just What is Essentially Causal about Perception?

There is a vast literature in the philosophy of perception concerning what properties

figure in the contents of perceptual experience. Call *Conservatism* the view that only low-level

properties—such as shapes, colors, and tones—can be represented. Call *Liberalism* the view that

some high-level properties—such as natural kinds, artifacts, and relations—can also be

represented. Conservatives and liberals disagree radically about what properties feature in

perceptual experience, but they nevertheless widely[29] agree that perception is an essentially

causal relationship.

I cannot adjudicate the conservative/liberal dispute here. But it seems safe to assume that

AEI is only going to be even initially plausible to liberals—assuming that moral properties are

high-level properties, conservatives are going to reject AEI from the get go. In what follows, I

will assess how best to understand the essentially causal nature of perception from within a

liberal framework. In the bigger picture, this is a contentious assumption. But since proponents

of AEI are already committed to liberalism, it is a safe assumption to make in this context. The

idea, then, is to home in on the essentially causal nature of perception by considering some

---

[29] Though not universally—see Snowdon (1979).

causally unique cases of properties thought to be perceivable by liberals about perceptual experience.

Begin by considering one natural but strict way to understand the causal constraint on perceptual representation:

Strict CC. Necessarily, if a property *F* is part of the contents of *S*'s perceptual experience *e*, then *F* (or the fact that *F* is instantiated) is at least partially causally responsible for *e*.

Strict CC is a relatively robust causal constraint on perceptual representation. But it's also an initially intuitive way of characterizing the causal nature of perception in a precise way. Nevertheless, I now argue that Strict CC should be rejected by liberals about perceptual experience. I'll argue this by considering three sorts of properties that liberals have defended as perceivable that couldn't be, if Strict CC were true: Absences, Gibsonian affordances, and the mental states of others. I consider each in turn.

Many liberals have recently argued that perceptual experience extends beyond the representation of positive properties to the representation of what we can call *absence properties.* For example, you may perceive a gap in an otherwise predictable pattern of coins arranged on the table,[30] the holes in a slice of Swiss cheese,[31] darkness in the inside of a cave,[32] or the sound of silence.[33] Suppose that these liberals are right—that we do perceive at least some absence properties. It's unclear if this is compatible with Strict CC, since it is unclear that the *lack* of something can figure in a genuine causal relationship. It's plausible that silence, for example, does not involve the existence of some causal property, but rather the lack of any causally efficacious property of a certain sort. The causal (in)efficacy of absence properties is a matter of

---

[30] Farennikova (forthcoming), 2.
[31] Lewis & Lewis (1970), Casati & Varzi (1994), pp.156-158.
[32] Sorensen (2008), Ch.10.
[33] See, e.g., Sorensen (2008, Ch.14, 2009), Soteriou (2011), Phillips (2013), and Simon & Garfunkel (1964).

some controversy.[34] However, the liberal arguments for the perceivability of absence properties do not directly depend on any such metaphysical claim. So, while it's not uncontentious, the perception of absence properties does provide some *prima facie* reason to favor a less robust causal constraint on perception than Strict CC.

A second set of properties which appears to conflict with Strict CC are what I'll call *affordance properties*. The idea of affordances in perceptual experience goes back to the psychological research of James J. Gibson (1977, 1979), but it has also been the subject of quite a bit of recent work in the philosophy of perception.[35] In Gibson's words, affordances are properties that tell an animal what an environment "*offers* the animal, what it *provides* or *furnishes*, either for good or ill."[36] Others—both philosophers and psychologists—following in Gibson's footsteps have attempted to refine the idea of affordance properties in various ways.[37] But paradigmatic instances of affordance properties should illustrate the idea clearly enough for present purposes. For example, an animal's prey may be seen as *to-be-killed,* a cup as *able-to-be-picked-up*, and the liquid in the cup as *drinkable*.[38] In brief, affordance properties relate agents and their abilities to the environment. They represent something like potential actions.

Though affordance properties are surely grounded in causal properties (for example, the structure of the cup underlies its ability to be picked up), they are arguably not themselves causal. However, according to at least many psychologists and philosophers, affordance properties are perceivable.[39] Insofar as this is right, it casts doubt on Strict CC, since the

---

[34] See, e.g. Lewis (2004), McGrath (2005), and Dowe (2009).
[35] Chemero (2003), Prosser (2011), Nanay (2012), and Siegel (2014).
[36] Gibson (1979), p.127.
[37] See, e.g. Reed (1996), Chemero (2003).
[38] Affordance properties appear, then, to come in two levels of strength—some features of objects render things possible, while others render things as appearing (practically) necessary. See Siegel (2014) for discussion of this point.
[39] This claim is far from uncontentious. But so far as I know, no one has rejected the perceivability of affordances on the grounds that they're not causal.

perception of affordance properties is incompatible with it. In short, affordance properties (and their perceivability) give us further reason to favor a less robust causal constraint on perception than Strict CC.

Finally, consider perception of the mental states of others. For example, many philosophers of perception and mind have recently argued that we can literally perceive the affective states of others.[40] Rowland Stout, for example, argues that we can "literally perceive someone's anger" in the sense that this perception is non-inferential.[41] The causal efficacy of mental states is one of the thorniest issues in philosophy. But, as far as I know, no one arguing against the perception of mental states has claimed that the perception of these states hinged on this controversy. Appeals in favor of the claim that we can perceive these states are generally phenomenological and empirical (appealing to modules in the brain dedicated to "mindreading"), not to the causal efficacy of these states. So it seems as though at least many liberals should be friendly to the perception of the mental states of others, regardless of their direct causal efficacy.

If a broadly liberal view of perceptual content is correct, it seems like Strict CC is not the right way to understand the causal constraint on perception. However, given the consensus that there is some causal constraint on perception, some weaker constraint must hold. Unfortunately, without taking controversial stands on the cases above (and others), a full account can't be attempted here. However, if any of the properties discussed above are perceivable, something *at least as* weak as the following must hold:

> *Weak CC*. Necessarily, if a property *F* is part of the contents of *S*'s perceptual experience *e*, then either (a) *F* or (b) some property (or set of properties) *G* which grounds *F* is at least partially causally responsible for *e*.

---

[40] See, e.g. Green (2010), Stout (2010), and McNeill (2012).
[41] Stout (2010), 29.

Depending on what one says about the cases above, Weak CC may remain too strong to be an accurate causal constraint on perception. But I hope to have established that the liberal about perceptual experience should favor something at least as weak as Weak CC, independent of any consideration of the perception of moral properties.

9. Moral Properties and the Reformulated Causal Constraint on Perception

I have argued that liberals about perceptual experience can plausibly reject Strong CC. This argument was necessary because ensuring that causally inefficacious moral properties can figure in the contents of perceptual experience is crucial to rebutting the Impossibility Argument. The weaker the causal constraint that one favors, the easier moral perception will be to accommodate. Suppose everything I've claimed is correct. It then appears that the strongest skeptical argument to be made against moral perception is to (a) grant Weak CC and then (b) argue that moral properties could not meet even this reformulated version of the causal constraint. I propose to grant (a) for the sake of argument, and turn to (b).

Whether moral properties can meet Weak CC turns on the metaphysical question of whether moral properties are wholly grounded in natural properties.[42] This is in principle compatible with a non-naturalist conception of moral properties, since grounding doesn't entail identity or reduction.[43] Consider a candidate moral principle, *suffering is bad*. If the truth of this moral principle is grounded in instances of causally efficacious properties, then the perception of the badness of suffering will be compatible with Weak CC.

---

[42] See Väyrynen (2013) for a helpful discussion about grounding and its potential use in illustrating the relationship between the natural and the normative.

[43] Rosen (2010), Audi (2012). For a non-naturalist metaphysics that fits this mold, see Leary (forthcoming).

Not every non-naturalist view favors this understanding of the grounding relation between the natural properties and the non-natural, normative properties.[44] I cannot hope to adjudicate this complex issue here. What must be flagged, however, is that different conceptions of the nature of the moral properties may still be compatible with the perception of those properties. Furthermore, there may be independent reasons to favor a causal constraint on perception which is weaker than Weak CC. If this is the case, a partial grounding of the normative in the natural may be enough. The weaker a causal constraint that holds, the easier it will be to defend a perceptualist moral epistemology.

10. *A Posteriori* Ethical Intuitionism and the Impossibility Argument: Redux and Recap

We saw above (section 7) that standard responses to the Impossibility Argument attempt to cast doubt on premise 1 (or its variants):

> 1. Any complete non-skeptical epistemology for a particular domain d must explain how we have epistemic access to the d-facts such that we are able to form beliefs that are responsive to those facts.

Though these strategies have the advantage of being neutral with respect to the mechanisms that underlie moral justification and moral knowledge, I've argued that they are nonetheless unsuccessful. This prompted a closer look at premise 2 of the argument:

> 2. There is no explicable mechanism to explain how human beings could have epistemic access to non-natural moral facts such that our beliefs are responsive to those facts.

---

[44] According to the conception given in the text, the truth of a moral principle is grounded in its instances—i.e. *suffering is bad* is a true moral principle in virtue of the distribution of the suffering properties and the badness properties. On an alternative conception, the moral principles play a metaphysical role in the instantiation of moral properties—i.e. the instantiation of badness depends on the interaction between the instantiation of suffering and the truth of the moral principle *suffering is bad.* On such a view, the instantiation of a moral property is only partially grounded in causally efficacious properties. So if the defender of moral perception favored such a view she would have to defend something weaker than Weak CC.

I argued that the proponent of AEI is well-placed to reject premise 2, since for her, perceptual experience provides a perfectly explicable mechanism[45] to provide human beings with epistemic access to non-natural moral facts. The proponent of the impossibility argument accepts that perception provides epistemic access to many empirical facts, so unless she can give us special reason to doubt that this occurs in the moral case, the skeptical argument fails.

As we saw, the moral skeptic might reject perceptual access in the moral case by claiming that a perceptual moral epistemology falls afoul of the consensus in the philosophy of perception that perception is an essentially causal process. Non-natural moral properties are causally inefficacious, so they are not perceivable, even in principle.

We have now seen that this attempt to single out moral properties using the causal constraint on perception is not so simple. The perception of causally inefficacious properties is compatible with the causal constraint on perception, properly construed, as long as those properties are related to causally efficacious properties in the right sort of way. Given some plausible causal constraints on perception, non-natural properties will be perceivable after all. Epistemic access to moral properties is possible, contra premise 2, if we endorse *A Posteriori* Ethical Intuitionism.

11. *A Posteriori* Ethical Intuitionism and the Extreme Unlikelihood Argument

If, as I've argued, the Impossibility Argument fails, then epistemic access to the moral properties is possible. Turn now to the second formulation of the debunking argument. The skeptic may grant that justified moral beliefs are in principle possible, but that it remains *extremely* unlikely that we have them, given the fact that it would be a massive coincidence if

---

[45] This isn't to say that no further details of how this mechanism works need filling in. An adequate account of moral perception must also show itself to be otherwise empirically, metaphysically, and epistemologically plausible. This arguably can't be done without a complete explanation of the mechanism(s) of moral perception.

our moral belief-forming faculties actually tracked moral facts. This is the Extreme Unlikelihood

Argument (EUA):

> 1. There is a huge range of logically possible stance-independent moral truths (and systems of truths).

> 2. Our actual moral judgments are not caused by the moral facts.

Therefore,

> 3. It is extremely unlikely that a significant subset of our moral beliefs are true, and those that are true are true as a result of luck in a way that undermines their justification.

The non-naturalist has three possibilities for responding to EUA. She can reject premise

1, reject premise 2, or, since the argument is not deductively valid, reject the move from

1 and 2 to the conclusion. Premise 1 is overwhelmingly plausible.[46] Premise 2 is also

widely accepted by non-naturalists.

However, the discussion above suggests a strategy for resisting the inference

from the premises to the conclusion. According to *A Posteriori* Ethical Intuitionism, our

moral experiences are caused, in part, by the properties which ground the moral facts.

As long as our moral experiences are caused by the property-arrangements that ground

the moral properties, there will be a non-accidental connection between our moral

experiences and the moral facts. Such non-accidentality would seem to undermine the

---

[46] One strategy for rejecting something like premise 1 is to argue that some substantive first-order moral facts are conceptual truths. This is suggested by Cuneo & Shafer-Landau (2014) and, on some readings, Audi (2008). I argue elsewhere that this sort of view just pushes the epistemological problem back a step.

claim of extreme unlikeliness in any particular case, and so would undermine the

debunker's conclusion of EUA.

12. *A Posteriori* Ethical Intuitionism and the No Reason Argument

"Ok", the skeptic might say in light of the reasoning above, "*if* there's a non-

accidental connection between our moral experiences and the natural facts that at least

partially ground the moral facts, then the debunker would be in trouble. But we have

no reason to think any such connection holds. So we've got no reason to think that any

of our moral experiences are tracking the moral facts." I think this is a perfectly natural

way to respond to the above discussion of the Extreme Unlikelihood Argument. And it

also provides a helpful segue, since it appears to be a restatement of the third

debunking argument, the No Reason Argument:

> 1. If you have no good reason to think that a given belief-forming process
> generates true beliefs, you cannot rationally maintain beliefs formed on the basis
> of such a process.
>
> 2. We have no good reason to think that moral belief-forming processes generate
> true beliefs.

Therefore,

> 3. We cannot rationally maintain our moral beliefs (i.e. they are unjustified).

The No Reason Argument is *prima facie* plausible. The non-naturalist might put pressure

on premise 2 by claiming that we do have good reason to think that our moral belief-

forming processes generate true beliefs. After all, many of our moral beliefs are true,

and they are a result of our moral belief-forming processes.[47] But this kind of response, as Street points out, is question-begging, since it assumes the truth of our moral beliefs, the very things which are presently in doubt. If the non-naturalist wants to reject premise 2, she's going to have to appeal to some *independent* reason for trusting her moral belief-forming faculties, lest her rejection be question-begging. It's hard to see how the non-naturalist can achieve this task.

However, the non-naturalist need not seek for independent verification for her moral beliefs, because, as Katia Vavova has pointed out, premise 1 of the No Reason Argument is false. As she puts it:

> [Premise 1] seems compelling because it raises a familiar sort of skeptical challenge. But it also collapses the debunker's challenge into that more ambitious one for which no empirical premise is necessary and which undermines much more than evaluative realism…Unless we are skeptics, we should grant that sensory perception is a perfectly good belief forming method. *Ceteris paribus*, if you perceive that *p*, you are rational in concluding that *p.* Do we have good reason to think that perception would lead us to true beliefs about our surroundings? Not if 'good' reason is understood as an appropriately independent reason: for if we set aside all that is in question, we must set aside all beliefs gained by perception. This includes all scientific beliefs, like the belief that evolutionary theory is true.[48]

Either premise 1 of the No Reason Argument is a bad epistemic principle, or a near global skepticism is true. If the former, then the argument fails to debunk moral beliefs. If the latter, then we have a lot more to be worried about than vindicating morality.

---

[47] The most explicit versions of this response can be found in Dworkin (2011) and Setiya (2012), but something like this move appears to be part and parcel of the third factor responses to debunking arguments (e.g. Enoch (2010), Wielenberg (2010)).

[48] Vavova (2014).

Since Vavova's rejection of premise 1 is otherwise epistemologically neutral, it's available to a wide variety of views about non-naturalist epistemology. But even here, I think that a view like *A Posteriori* Ethical Intuitionism has a slight advantage. For the proponent of AEI, moral experiences are considered a special instance of the more general theory of perceptual justification. If moral experiences are properly characterized as perceptual experiences, they will *prima facie* inherit some of the epistemological credentials of other perceptual experiences. In other words, the would-be debunker owes the proponent of AEI a story about why we should be especially suspicious of perceptual experiences of a moral sort despite the fact that perception is, in other cases, a justification-conferring process. This isn't to claim that grounds for suspicion couldn't be given. Rather, it is just to claim that they must, and until they have, the proponent of AEI should be (slightly) less nervous than proponents of *a priori* moral epistemologies which—along with *a priori* justification more generally—have come under fire, both historically and at present.

13. Taking Stock

Epistemological objections to non-naturalism have been around for some time, and they will surely continue to be. The standard approach that non-naturalists have taken to these sorts of concerns has been to argue that the epistemological principles that these arguments rely on are mistaken. The advantage of these approaches is that they are otherwise compatible with a wide variety of positive conceptions of the

mechanisms of moral epistemology. However, I have argued that the two most popular

strategies—partners in guilt and third factor responses—are unsuccessful. Given the

failing of these moral epistemology theory-neutral attempts at responding to debunking

and epistemic access arguments, in this paper I have taken a different tack: I have

suggested that endorsing *A Posteriori* Ethical Intuitionism (AEI) provides the non-

naturalist with the resources to rebut what I take to be the three most powerful versions

of epistemological arguments against anti-skeptical non-naturalism.

According to the first argument, the Impossibility Argument, epistemic access to

the non-natural properties is impossible, since it is inexplicable how we could bear

some positive epistemic relation to them. I claimed that the proponent of AEI has a

surprisingly straightforward answer to this argument: Perceptual experience provides

us with access to the moral properties. It may have seemed as though this response was

a non-starter, given the causal nature of perception and the non-causal nature of non-

natural properties. However, I argued that, once properly understood, the causal

constraint on perception is compatible with the perception of non-natural properties,

even if those properties are causally inefficacious.

Once this response to the Impossibility Argument is secured, responding to the

second argument, the Extreme Unlikelihood Argument, is straightforward. If we

perceive the non-natural properties, then moral perceptual experiences bear a non-

accidental connection to the non-natural properties. This non-accidentality undermines

the claim that moral beliefs formed on the basis of moral experience are extremely unlikely to be true.

At this point, the skeptic balks. At this point in the dialectic, we'd established the following conditional: If we perceive the non-natural properties (and such a thing is at least possible), then there is some non-accidental connection between our moral experiences and the non-natural properties. But the skeptic, we saw, now asks: Why should we think the antecedent of that conditional is met? Bare possibility doesn't establish actuality, after all. The skeptic says that if we have no good reason to trust our moral belief-forming faculty, then any *prima facie* justification we might have had is undercut. This is the third epistemological argument against non-naturalism that I considered, the No Reason Argument. However, as Vavova (2014) points out, the epistemic principle that this argument relies on is faulty, and if true, would entail a near global skepticism. In whatever sense that we don't know that our moral experiences are reliable, we can say the same thing about perceptual experiences more generally—any evidence that they are reliable will be question-begging against the skeptic. The No Reason Argument fails.

A Posteriori Ethical Intuitionism might be right. In other words, we may have moral perceptual experiences, and these experiences may reliably track non-natural properties. If both of these things hold, then we have justified moral beliefs, and, arguably, moral knowledge which is grounded in moral perceptual experience. It

seems, then, that the moral skeptic's attempt to debunk or undercut anti-skeptical non-naturalism has failed. At the same time, non-naturalists should own up to the kernel of truth in the skeptic's thinking. Since we can't assess the status of our moral beliefs from any standpoint wholly independent from our moral belief-forming faculties, we can't confidently claim that we have any moral knowledge at all. In other words, we might know, but we can't know that we know. This is true of many domains of inquiry, but given the importance of morality, non-naturalists should take this epistemic humility to heart.

Works Cited:

Ahlenius, Henrik and Torbjórn Tánnsjó. (2012). "Chinese and Westerners Respond Differently to the Trolley Dilemmas", in *Journal of Cognition and Culture,* 12(3-4):195-201.

Ahluwalia, A. (1978). "An intra-cultural investigation of susceptibility to 'perspective' and 'non-perspective' spatial illusions", in British Journal of Psychology, 69(2):233-241.

Armstrong, D. M. (1980). *Nominalism and Realism: Volume 1: Universals and Scientific Realism.* Cambridge University Press.

Audi, Paul. (2012). "A Clarification and Defense of the Notion of Grounding", in Fabrice Correia & Benjamin Schnieder (eds.), *Metaphysical Grounding: Understanding the Structure of Reality.* Cambridge University Press. Pp.101-121.

Audi, Robert. (2005). *A Theory of the Good and the Right.* Princeton University Press.

Audi, Robert (2008a). "Intuition, Inference, and Rational Disagreement in Ethics", in *Ethical Theory and Moral Practice*, 11(5): 475-492.

Audi, Robert. (2010). Moral Perception and Moral Knowledge. *Aristotelian Society Supplementary Volume*, *84*(1), 79-97.

Audi, Robert. (2013). *Moral Perception (Soochow University Lectures in Philosophy).* Princeton University Press.

Audi, Robert. (2008b). "Skepticism about *A Priori* Justification: Self-Evidence, Defeasibility, and Cogito Propositions", in *The Oxford Handbook of Skepticism,* John Greco (ed.). Oxford University Press.

Balaguer, Mark. (1998). *Platonism and Anti-Platonism in Mathematics.* Oxford University Press.

Bayne, T. (2009). Perception and the Reach of Phenomenal Content. *The Philosophical Quarterly*, *59*(236), 385-404.

Bealer, George. (1998). "A theory of concepts and concept possession", in *Philosophical Issues*,9:261-301.

Bedke, Matthew S. (2009). "Intuitive Non-Naturalism Meets Cosmic Coincidence", in *Pacific Philosophical Quarterly*, 90(2):188-209.

Bedke, Matthew S. (2014) "No Coincidence?". *Oxford Studies in Metaethics, 9.* Available at http://faculty.arts.ubc.ca/mbedke/Research.html

Benacerraf, Paul. (1973). "Mathematical Truth", in *Journal of Philosophy*, 70(19):661-679.

Bengson, John. (forthcoming). "Grasping the Third Realm", in *Oxford Studies in Epistemology*, 5. Available at: https://sites.google.com/site/johnbengson/Grasping_the_Third_Realm.pdf?attred irects=0

Ben-Ze'ev, Aaron (2000). "Emotions, Responsibility and Morality", in A. van den Beld (ed.), Moral Responsibility and Ontology. Kluwer.

Berry, J.W. (1971). "Muller-Lyer susceptibility: Culture, ecology, or race?", in International Journal of Psychology, 6:193-197.

Blair, J., Sellars, C., Strickland, I., Clark, F., Williams, A., Smith, M., et al. (1996). Theory of Mind in the Psychopath. *The Journal of Forensic Psychiatry*, 7(1), 15-25.

Blair, R. (1995). A Cognitive Developmental Approach to Morality: Investigating the Psychopath. *Cognition*, *57*, 1-29.

Blair, James, R., Jones, L., Clark, F., & Smith, M. (1997). The Psychopathic Individual: A Lack of Response to Distress Cues?. *Psychophysiology*, *34*(2), 192-198.

Blair, James R. (2007). "Empathic dysfunction in psychopathic individuals", in *Empathy in Mental Illness.* Cambridge University Press.

Boghossian, Paul. (2000). "Knowledge of Logic", in Paul Boghossian & Christopher Peacocke (eds.), *New Essays on the A Priori.* Oxford University Press.

Bonjour, Laurence. (1985). *The Structure of Empirical Knowledge.* Harvard University Press.

Bonjour, Laurence. (1998). *In Defense of Pure Reason.* Cambridge University Press.

BonJour, Laurence. (2000). "The Elements of Coherentism", in Sven Bernecker & Fred I. Dretske (eds.). *Knowledge: Readings in Contemporary Epistemology*. Oxford University Press.

BonJour, Laurence. (2007). "Epistemological Problems of Perception", in Stanford Encyclopedia of Philosophy. http://plato.stanford.edu/entries/perception-episprob/

Boyd, Richard. (1988). "How to Be a Moral Realist", in *Essays on Moral Realism,* G. Sayre-McCord (ed.), 181-228.

Brady, Michael. (2014). *Emotional Insight*. Oxford University Press.

Brink, David O. (1984). "Moral Realism and the Sceptical Arguments from Disagreement and Queerness", in *Australasian Journal of Philosophy,* 62(2):111-125.

Brink, David Owen. (1986). "Externalist Moral Realism", in *Southern Journal of Philosophy*, *24*(S1), 23-41.

Brink, David Owen. (1989). *Moral Realism and the Foundations of Ethics.* Cambridge University Press.

Brogaard, B. (forthcoming). Do We Perceive Natural Kind Properties? *Philosophical Studies*.

Burge, Tyler. (2003). "Perceptual Entitlement", in *Philosophy and Phenomenological Research,* 67(3):503-548.

Byrne, A. (2009). Experience and Content. *The Philosophical Quarterly*, *59*(236), 429-451.

Carnap, Rudolf. (1950). "Empiricism, Semantics, and Ontology", in *Revue Internationale De Philosophie,* 4(2):20-40.

Casati, Roberto & Achille C. Varzi. (1994). *Holes and Other Superficialities.* Cambridge: MIT Press.

Cecchi, Ariel S. (2014). "Cognitive Penetration, Perceptual Learning and Neural Plasticity", in Dialectica, 68(1):63-95.

Chalmers, David. (2004). "Epistemic Two-Dimensional Semantics", in *Philosophical Studies,* 118(1-2):153-226.

Chalmers, David. (2004). "The Representational Character of Experience", in *The Future for Philosophy*, B. Leiter (ed.), Oxford: Oxford University Press, 153-181.

Chalmers, David. (2012). *Constructing the World.* Oxford University Press.

Chalmers, David, and Frank Jackson. (2001). "Conceptual Analysis and Reductive Explanation", in *Philosophical Review,* 110(3):315-361.

Chappell, Timothy. (2008). Moral Perception. *Philosophy*, *83*(4), 421-437.

Chemero, Anthony. (2003). "An Outline of a Theory of Affordances", in *Ecological Psychology,* 15(2):181-195.

Cheyne, Colin. (1998). "Existence Claims and Causality", in *Australasian Journal of Philosophy,* 76:34-47.

Cheyne, Colin (2001). *Knowledge, Cause, and Abstract Objects.* Kluwer Academic Publishers.

Cholbi, Michael. (2007). "Moral expertise and the credentials problem", in *Ethical Theory and Moral Practice,* 10(4):323:334.

Chudnoff, Elijah. (2011). "The Nature of Intuitive Justification", in *Philosophical Studies,* 153(2):313-333.

Chudnoff, Elijah. (forthcoming). "Awareness of Abstract Objects", forthcoming in *Nous.*

Church, Jennifer. (2013). *Possibilities of Perception.* Oxford University Press.

Clarke-Doane, Justin. (forthcoming). "What is the Benacerraf Problem?" in Fabrice Pataut (ed.), *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity.* Springer.

Coons, C. (forthcoming). "How to Prove that Some Acts Are Wrong (Without Using Substantive Moral Premises)", in *Philosophical Studies.*

Copp, David. (2000). "Milk, Honey, and the Good Life on Moral Twin Earth", in *Synthese, 124*(1-2), 113-137.

Copp, David. (2005). "Moral Naturalism and Three Grades of Normativity", in Peter Schaber, ed., *Normativity and Naturalism.* Frankfurt: Ontos-Verlag.

Cowan, Robert. (2013a). Clarifying Ethical Intuitionism. *European Journal of Philosophy,* OnlineFirst: http://onlinelibrary.wiley.com/doi/10.1111/ejop.12031/abstract;jsessionid=C7C599A54EFFB2C4218BEA0880A36526.d03t04.

Cowan, Robert. (2013b). Perceptual Intuitionism. *Philosophy and Phenomenological Research,* OnlineFirst: http://onlinelibrary.wiley.com/doi/10.1111/phpr.12023/abstract

Cowan, Robert. (forthcoming). "Cognitive Penetrability and Ethical Perception". http://www.robertcowan.net/#!research/cm8a

Cullison, Andrew. (2010a). "Moral Perception", in *European Journal of Philosophy*, 18(2):159-175.

Cullison, Andrew. (2010b). "What are Seemings?" in *Ratio*, 23(3), 260-274.

Cuneo, Terence, and Russ Shafer-Landau. (2014). "The Moral Fixed Points: New Directions for Moral Nonnaturalism", in *Philosophical Studies*, Online First, 1-45.

Dancy, Jonathan. (2010). "Moral Perception", in *Aristotelian Society Supplementary Volume*, *84*(1), 99-117.

Daniels, Norman. (1996). *Justice and Justification: Reflective Equilibrium in Theory and Practice.* Cambridge University Press.

Davies, Martin. (1997). "Externalism and Experience", in N. Block and O.J. Flanagan (eds.), *The Nature of Consciousness: Philosophical Debates.* MIT Press.

Deonna, Julien A. & Teroni, Fabrice. (2012). *The Emotions: A Philosophical Introduction.* Routledge.

DeRose, Keith. (2004). "Sosa, safety, sensitivity, and skeptical hypotheses", in J. Greco (ed.), *Ernest Sosa and His Critics.*

De Sousa, Ronald. (1984). *The Rationality of Emotion.* The MIT Press.

De Sousa, Ronald. (2011). *Emotional Truth.* Oxford University Press.

De Sousa, Ronald. (2013). "Emotion", in *Stanford Encyclopedia of Philosophy.* *http://plato.stanford.edu/entries/emotion/*

Döring, Sabine. (2003). "Explaining Action by Emotion", in *Philosophical Quarterly*, 53(211):214-230.

Döring, Sabine. (2009). "The Logic of Emotional Experience: Noninferentiality and the Problem of Conflict Without Contradiction", in *Emotion Review*, 1(3):240-247.

Dowe, Phil. (2009). "Absences, Possible Causation, and the Problem of Non-Locality", in *The Monist*, 92(1):23-40.

Dretske, Fred. (1995). *Naturalizing the Mind.* Cambridge: MIT Press.

Driver, Julia. (2006). "Autonomy and the asymmetry problem for moral expertise", in *Philosophical Studies*, 128(3):619-644.

Dworkin, Ronald. (2011). *Justice for Hedgehogs.* Belknap Press of Harvard University Press.

Earlenbaugh, J. and B. Molyneux. (2009). "Intuitions are Inclinations to Believe", in *Philosophical Studies*, 145(1), 89-109.

Enoch, David. (2011). *Taking Morality Seriously: A Defense of Robust Realism.* Oxford University Press.

Evans, Gareth. (1982). *The Varieties of Reference.* Oxford University Press.

Farennikova, Anna. (forthcoming). "Perception of Absence and Penetration from Expectation", in *Review of Philosophy and Psychology.*

Faucher, Luc and Christine Tappolet. (eds.) (2002). *The Modularity of Emotions.* University of Calgary Press.

Field, Hartry. (1989). *Realism, Mathematics & Modality*. Basil Blackwell.

Field, Hartry. (2005). "Recent Debates about the A Priori", in *Oxford Studies in Epistemology*, Vol. 1. Oxford University Press.

FitzPatrick, William J. (2008). "Robust Ethical Realism, Non-Naturalism and Normativity", in *Oxford Studies in Metaethics, Vol. 3.* Oxford University Press.

Fodor, Jerry A. (1983). *The Modularity of Mind: An Essay on Faculty Psychology*. MIT Press.

Foot, Philippa (1972). "Morality as a System of Hypothetical Imperatives", in *Philosophical Review*, 81(3):305-316.

Foot, Philippa. (2001). *Natural Goodness.* Oxford University Press.

Fox, Elaine. (2008). *Emotion Science: Cognitive and Neuroscientific Approaches to Understanding Human Emotions.* Palgrave Macmillan.

Fraser, Benjamin James. (2014). "Evolutionary Debunking Arguments and the Reliability of Moral Cognitiion", in *Philosophical Studies*, 168(2):457-473.

Fricker, Miranda. (2006). Powerlessness and Social Interpretation. *Episteme: A Journal of Social Epistemology*, 3(1):96-108.

Fumerton, Richard and Ali Hasan. (2010). "Foundationalist Theories of Epistemic Justification", in the *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/entries/justep-foundational/

Gibbard, Allan. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment.* Harvard University Press.

Gibbard, Allan. (2003). *Thinking How to Live.* Harvard University Press.

Gibilisco, C.J.K. (ms). "Moral and mathematical realism: partners in epistemological guilt?"

Gibson, James J. (1979). *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin.

Goldie, Peter. (2003). Emotion, Feeling, and Knowledge of the World. *Thinking About Feeling: Contemporary Philosophers on Emotions,* ed. Robert C. Solomon. Oxford University Press.

Goldie, Peter. (2007). "Emotion", in *Philosophy Compass,* 2(6):928-938.

Goldie, Peter. (2009). Getting feelings into emotional experience in the right way. *Emotion Review,* 1(3), 232-239.

Goldman, A.I. (1967). "A Causal Theory of Knowing", in *The Journal of Philosophy*, 64(12):357-372.

Goldman, A.I. (1976). "Discrimination and Perceptual Knowledge", in *The Journal of Philosophy*, 73:771-791.

Goldman, A.I. (2008). "Reliabilism", in *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/reliabilism/

Green, Mitchell. (2010). "Perceiving Emotions", in the *Aristotelian Society Supplementary Volume,* 84(1):45-61.

Greene, Joshua., & Sinnott-Armstrong, Walter. (2007). The Secret Joke of Kant's Soul. *Moral Psychology: Volume 3: The Neuroscience of Morality*, 3, 35-81.

Gross, James J. (2002). "Emotion Regulation: Affective, Cognitive, and Social Consequences", in *Psychophysiology,* 39:281-291.

Gross, James J. (ed.) (2007). *Handbook of emotion regulation.* New York: Guilford Press.

Gross, James J. (2008). "Emotion and Emotion Regulation", in *Handbook of Personality: Theory and Research.* John, Oliver P., Richard W. Robins, & Lawrence A. Pervin (eds.). The Guilford Press.

Gross, James J. & Oliver P. John. (2003). "Individual differences in two emotion regulation processes: Implications for affect, relationships, and well-being", in *Journal of Personality and Social Psychology,* 85:348-362.

Harman, Gilbert. (1977). *The Nature of Morality: An Introduction to Ethics.* Oxford University Press.

Hare, R. (1985). "Scoring Manual for the Psychopathy Checklist". *Unpublished Manuscript*.

Hawthorne, John. (2002). "Deeply Contingent a Priori Knowledge", in *Philosophy and Phenomenological Research,* 65(2):247-269.

Heathwood, Chris. (2015). "Irreducibly Normative Properties", in *Oxford Studies in Metaethics,* 10. Oxford University Press.

Horgan, T., & Timmons, M. (1992). "Troubles on Moral Twin Earth: Moral Queerness Revisited", in *Synthese,* *92*(2), 221-260.

Horgan, T., & Timmons, M. (1991). "New Wave Moral Realism Meets Moral Twin Earth", in *Journal of Philosophical Research*, *16*, 447-465.

Horgan, T., & Timmons, M. (2000). "Copping Out on Moral Twin Earth", in *Synthese*, *124*(1-2), 139-152.

Horgan, T., & Tienson, J. (2002). The Phenomenology of Intentionality and the Intentionality of Phenomenology. *Philosophy of Mind: Classical and Contemporary Readings*, 520-533.

Huemer, Michael. (2000). "Naturalism and the Problem of Moral Knowledge", in *Southern Journal of Philosophy,* 38(4):575-597.

Huemer, Michael. (2001). *Skepticism and the Veil of Perception.* Lanham: Rowman & Littlefield.

Huemer, Michael. (2005). *Ethical Intuitionism.* Palgrave Macmillan.

Huemer, M. (2007). Compassionate Phenomenal Conservatism. *Philosophy and*

Phenomenological Research, 74(1), 30-55.

Huemer, Michael. (2008). "Revisionary Intuitionism", in *Social Philosophy and Policy*, 25(1):368-392.

Huemer, Michael. (2009). "Apology of a Modest Intuitionist", in *Philosophy and Phenomenological Research,* 78(1):222-236.

Huemer, Michael. (ms) "Why I am Not an Objectivist", available at http://spot.colorado.edu/~huemer/rand.htm

Ichikawa, Jonathan Jenkins and Matthias Steup. (2012). "The Analysis of Knowledge", in *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/entries/knowledge-analysis/#CauTheKno

Jackson, Frank. (1998). *From Metaphysics to Ethics: A Defense of Conceptual Analysis.* Clarendon Press.

Jackson, Frank. (2001). "Locke-ing onto content", in *Evolution, Naturalism and Mind,* D. Walsh (ed.). Cambridge University Press. Ch.7.

Jacobson, Daniel (2005). "Seeing by Feeling: Virtues, Skills, and Moral Perception", in *Ethical Theory and Moral Practice,* 8(4):387-409.

James, William. (1884). "What is an Emotion?" in *Mind,* 9(34), 188-205.

Jenkins Ichikawa, Jonathan, & Matthias Steup. (2012). "The Analysis of Knowledge", in *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/entries/knowledge-analysis/

Johnston, Mark. (2001). "The Authority of Affect", in *Philosophy and Phenomenological Research,* 63(1):181-214.

Joyce, Richard. (2001). *The Myth of Morality.* Cambridge: Cambridge University Press.

Joyce, Richard. (2006). *The Evolution of Morality.* MIT Press.

Kaplan, David. (1967). "Quantifying In", in *Synthese,* 19:178-214.

King, Jeff. (2011). "Structured Propositions", in *Stanford Encyclopedia of Philosophy.* Available at http://plato.stanford.edu/entries/propositions-structured/ (Accessed 7/21/2014)

Klein, Peter. (1998). "Review: Foundationalism and the Infinite Regress of Reasons", in *Philosophy and Phenomenological Research,* 58(4):919-925.

Klein, Peter. (2011). "Infinitism", in Sven Bernecker & Duncan Pritchard (eds.), *Routledge Companion to Epistemology.* Routledge.

Kornblith, Hilary. (2007). "Naturalism and Intuitions", in *Grazer Philosophische Studien,* 74(1):27-49.

Korsgaard, Christine. (1996). *The Sources of Normativity.* Cambridge University Press.

Korsgaard, Christine. (2009). *Self-Constitution: Agency, Identity, and Integrity.* Oxford University Press.

Krigolson, Olav E., Lara J. Pierce, Clay B. Holroyd, & James W. Tanaka. (2009). "Learning to Become an Expert: Reinforcement Learning and the Acquisition of Perceptual Expertise", in Journal of Cognitive Neuroscience, 21(9):1833-1840.

Lacewing, Michael. (2015). "Emotion, Perception, and the Self in Moral Epistemology", in *Dialectica,* 69(3):335-355.

Lewis, David & Stephanie Lewis. (1970). "Holes", in *Australasian Journal of Philosophy,* 48(2):206-212.

Lewis, David. (1983). "New Work for a Theory of Universals", in *Australasian Journal of Philosophy,* 64, 25-46.

Lewis, David. (1986). *On the Plurality of Worlds.* Blackwell Publishers.

Lewis, David. (2004). "Void and Object", in John Collins, Ned Hall & L.A. Paul (eds.), *Causation and Counterfactuals.* MIT Press.

Liggins, David. (2010). "Epistemological Objections to Platonism", in *Philosophy Compass,* 5(1):67-77.

Linsky, Bernard and Edward N. Zalta. (1995). "Naturalized platonism versus platonized naturalism", in *Journal of Philosophy,* 92(10):525-555.

Logue, Heather. (forthcoming). "Visual Experience of Natural Kind Properties: Is There Any Fact of the Matter?" in *Philosophical Studies.*

Loftus, Elizabeth F. and Edith Greene. (1980). "Warning: Even Memory for Faces May be Contagious", in *Law and Human Behavior*, 4(4):323-334.

Loftus, Elizabeth F., David G. Miller and Helen J. Burns. (1978). "Semantic Integration of Verbal Information into Visual Memory", in *Journal of Experimental Psychology,* 4(1):19-31.

Lycan, William. (2001). "The Case for Phenomenal Externalism", in *Philosophical Perspectives.* 15, 17-35.

Lycan, William. (2006). "Representationalism", in *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/entries/consciousness-representational/

Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong.* Penguin.

Macpherson, Fiona. (2012). "Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism", in *Philosophy and Phenomenological Research*, 84(1):24-62.

Manley, David. (2007). "Safety, Content, Apriority, Self-knowledge", in *Journal of Philosophy,* 104(8):403-423.

Maund, B. (2012). Color. *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/entries/color/

McBrayer, Justin P. (2010a). "Moral Perception and the Causal Objection", in *Ratio,* 23(3), 291-307.

McBrayer, Justin P. (2010b). "A Limited Defense of Moral Perception", in *Philosophical Studies,* 149(3):305-320.

McGrath, Sarah. (2004). "Moral Knowledge by Perception", in *Philosophical Perspectives,* 18(1):209-228.

McGrath, Sarah. (2005). "Causation by Omission: A Dilemma", in *Philosophical Studies,* 123(1-2):125-148.

McGrath, Sarah. (2014). "Relax? Don't do it! Why moral realism won't come cheap", in *Oxford Studies in Metaethics,* 9.

McNeill, William E.S. (2012). "On Seeing That Someone is Angry", in *European Journal of Philosophy,* 20(4):575-597.

McPherson, Tristram. (2011). "Against Quietist Normative Realism", in *Philosophical Studies*, 154(2):223-240.

Merrill, G.H. (1980). "The Model-Theoretic Argument Against Realism", in *Philosophy of Science*, 47(1), 69-81.

Merli, David. (2002). "Return to Moral Twin Earth" in *Canadian Journal of Philosophy*, *32*(2), 207-240.

The Muller-Lyer Illusion. Rochester Institute of Technology. Retrieved July 1, 2014, from http://www.rit.edu/cla/gssp400/muller/muller.html

Nanay, Bence. (2012). "Action-Oriented Perception", in *European Journal of Philosophy*, 20(3):430-446.

Nozick, Robert. (1981). *Philosophical Explanations.* Harvard University Press.

Nussbaum, Martha.C. (2001). *Upheavals of Thought: The Intelligence of Emotions.* Cambridge University Press.

O'Callaghan, C. (2011). "Against Hearing Meanings", in *The Philosophical Quarterly*. 61(245), 783-807.

Oddie, Graham. (2005/2009). *Value, Reality, and Desire.* Clarendon Press.

Parfit, Derek. (2011). *On What Matters.* Oxford University Press.

Patrick, Christopher J., Bradley, Margaret M., & Lang, Peter J. (1993). "Emotion in the Criminal Psychopath: Startle Reflect Modulation", in *Journal of Abnormal Psychology*, 102(3), 82-92.

Peacocke, C. (1992). *A study of concepts*. Cambridge, Mass.: MIT Press.

Phillips, Ian. (2013). "Hearing and Hallucinating Silence", in Fiona Macpherson & Dimitris Platchias (eds.), *Hallucination.* MIT Press.

Powers, Peter A., Joyce L. Andriks, and Elizabeth F. Loftus. (1979). "Eyewitness Accounts of Females and Males", in *Journal of Applied Psychology*, 64(3):339-347.

Price, Richard. (2009). "Aspect-Switching and Visual Phenomenal Character", in *The Philosophical Quarterly*, *59*(236), 508-518.

Prinz, Jesse. (2004). *Gut Reactions: A Perceptual Theory of the Emotions.* Oxford University Press.

Prinz, Jesse. (2008). "Is Emotion a Form of Perception?" in Faucher & Tappolet 2008.

Prosser, Simon. (2011). "Affordances and Phenomenal Character in Spatial Perception", in *Philosophical Review,* 120(4):475-513.

Pryor, James. (2004). "An Epistemic Theory of Acquaintance", the UCSB Conference on Content and Concepts. Available at: http://www.jimpryor.net/research/papers/Acquaintance.pdf

Pust, Joel. (2012). "Intuition", in *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/entries/intuition/

Pylyshyn, Zenon W. (1999). "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception", in Behavioral and Brain Sciences, 22(3):341-365.

Railton, Peter. (1986). "Moral Realism", in *Philosophical Review,* 95:163-207.

Reed, Edward S. (1996). *Encountering the world: Toward an ecological psychology.* Oxford University Press.

Richell, R., Mitchell, D., Newman, C., Leonard, A., Baron-Cohen, S., & Blair, R. (2003). Theory of Mind and Psychopathy: Can Psychopathic Individuals Read the 'Language of the Eyes'?. *Neuropsychologia, 41*(5), 523-526.

Ridge, Michael. (2008). "Moral Non-Naturalism", in *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/moral-non-naturalism/

Robb, David & John Heil. (2013). "Mental Causation", in the *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/entries/mental-causation/

Roberts, Robert C. (2003). *Emotion: An Essay in Aid of Moral Psychology.* Cambridge University Press.

Roberts, Tom. (forthcoming). "Emotional Regulation and Responsibility", in *Ethical Theory and Moral Practice.*

Roeser, Sabine. (2006). "A Particularist Epistemology: 'Affectual Intuitionism'", in *Acta Analytica,* 21(1):33-44.

Roeser, Sabine. (2011). *Moral Emotions and Intuitions.* Palgrave Macmillan.

Rogers, J. and J. Matheson. (2011). "Bergmann's dilemma: exit strategies for internalists", in *Philosophical Studies*, 152, 55-80.

Rosati, Connie. (2006). "Moral Motivation", in the *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/moral-motivation/

Rosen, Gideon. (2010). "Metaphysical Dependence: Grounding and Reduction", in Bob Hale & Aviv Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press.

Russell, Bertrand. (1912). *The Problems of Philosophy*. Barnes & Noble Books.

Russell, Gillian. (2011). *Truth in Virtue of Meaning: A Defence of the Analytic/Synthetic Distinction*. Oxford University Press.

Salmela, Mikko. (2011). "Can Emotion Be Modelled on Perception?" in *Dialectica*, 65(1):1-29.

Scanlon, Thomas. (2014). *Being Realistic about Reasons*. Oxford University Press.

Schechter, Joshua. (2010). "The Reliability Challenge and the Epistemology of Logic", in *Philosophical Perspectives*, 24(1):437-464.

Schroeder, Mark. (2008). "How Does the Good Appear to Us?" in *Social Theory and Practice*, 34(1):119-130.

Schroeter, Laura. (ms) "Can Metasemantics Resolve the Epistemic Challenge to Moral Realism?" Available at https://www.academia.edu/2923803/Can_Metasemantics_Resolve_the_Epistemic_Challenge_to_Moral_Realism (Accessed 7/21/2014)

Schwitzgebel, Eric. (2010). "Introspection", in *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/introspection/

Segall, Marshall H., Donald T. Campbell, & Melville J. Hesrskovits. (1963). "Cultural Differences in the Perception of Geometric Illusions", in Science, 139(3556):767-771.

Setiya, Kieran. (2012). *Knowing Right From Wrong*. Oxford University Press.

Shafer-Landau, Russ. (2003/2005). *Moral Realism: A Defence*. Oxford University Press.

Shafer-Landau, Russ. (2009). "A Defence of Categorical Reasons", in *Proceedings of the Aristotelian Society,* 109(1pt2):189-206.

Sharvy, Richard. (2007). "Who's to say what's right or wrong? People who have Ph.D.s in philosophy, that's who", in *Journal of Libertarian Studies,* 21(3):3-24.

Shoemaker, David W. (2011). Psychopathy, Responsibility, and the Moral/Conventional Distinction. *Southern Journal of Philosophy,* 49(s1), 99-124.

Silins, Nicholas. (2013). "The Significance of High-Level Content", in *Philosophical Studies* 162(1):13-33.

Siegel, Susanna. (2006). Which Properties are Represented in Perception? *Perceptual Experience*, *1*, 481-503.

Siegel, Susanna. (2010). *The contents of visual experience*. New York: Oxford University Press.

Siegel, Susanna. (2011). Cognitive Penetrability and Perceptual Justification. *Nous*, *46*(2), 201-222.

Siegel, Susanna. (2014). "Affordances and the Contents of Perception", in Berit Brogaard (ed.), *Does Perception Have Content?* Oxford University Press.

Siegel, Susanna. (forthcoming). "Rational Evaluability and Perceptual Farce". http://philpapers.org/rec/SIEREA

Silva Jr., Paul. (2013). "Epistemically Self-Defeating Arguments and Skepticism About Intuition", in *Philosophical Studies,* 164(3):579-589.

Simon, Paul & Art Garfunkel. (1964). "The Sound of Silence", on *Wednesday Morning, 3 AM* [CD]. Columbia Records.

Singer, Peter. (2005). Ethics and Intuitions. *Journal of Ethics*, *9*(3-4), 331-352.

Sinnott-Armstrong, Walter. (2006). "Moral intuitionism meets empirical psychology", in *Metaethics After Moore*, Terry Horgan and Mark Timmons (eds.), Oxford University Press, 339-366.

Sinnott-Armstrong, Walter. (2011). "An Empirical Challenge to Moral Intuitionism", in Jill Graper Hernandez (ed.), *The New Intuitionism.* Continuum.

Skarsaune, Knut Olav. (2011). "Darwin and moral realism: Survival of the iffiest", in *Philosophical Studies,* 152(2):229-243.

Skeem, J. L., Polaschek, D. L., Patrick, C. J., & Lilienfeld, S. O. (2011). Psychopathic Personality: Bridging the Gap Between Scientific Evidence and Public Policy. *Psychological Science in the Public Interest*, *12*(3), 95-162.

Skorupski, John. (2013). *The Domain of Reasons.* Oxford University Press.

Smith, Angela. (2005). "Responsibility for Attitudes: Activity and Passivity in Mental Life", in *Ethics*, 115(2):236-271.

Smith, Michael. (1994/1995). *The Moral Problem.* Blackwell.

Snowdon, Paul F. (1979). "Perception, Vision, and Causation", in *Proceedings of the Aristotelian Society*, 81:175-192.

Soloman, Robert C. (1984). *The Passions: The Myth and Nature of Human Emotions.* Doubleday.

Sorensen, Roy A. (2008). *Seeing Dark Things: The Philosophy of Shadows.* Oxford University Press.

Sorensen, Roy A. (2009). "Hearing Silence: The Perception and Introspection of Absences", in Matthew Nudds & Casey O'Callaghan (eds.), *Sounds and Perception.* Oxford University Press.

Sosa, Ernest. (1985). "The Coherence of Virtue and the Virtue of Coherence: Justification in Epistemology", in *Synthese*, 64(1):3-28.

Sosa, Ernest. (2007). Intuitions: Their Nature and Epistemic Efficacy. *Grazer Philosophische Studien*, 74(1), 51-67.

Soteriou, Matthew. (2011). "The Perception of Absence, Space, and Time", in Johannes Roessler, Hemdat Lerman & Naomi Eilan (eds.), *Perception, Causation, and Objectivity.* Oxford University Press.

Stewart, M.V. (1973). "Tests of the 'carpentered world' hypothesis by race and environment in America and Zambia", in International Journal of Psychology, 8:83-94.

Stratton-Lake, Philip. (2010). "Intuitionism", in John Skorupski (ed.), *The Routledge Companion to Ethics.* Routledge.

Stokes, Dustin. (2012). "Perceiving and Desiring: A New Look at the Cognitive Penetrability of Experience", in Philosophical Studies, 158(3):479-492.

Stokes, Dustin. (2014). "Cognitive Penetration and the Perception of Art", in Dialectica 68(1):1-34.

Stokes, Dustin. (forthcoming). "Toward a Consequentialist Understanding of Cognitive Penetration", in Athanassosios Raftopoulos & John Zeimbekis (eds.) Cognitive Penetrability. http://philpapers.org/rec/STOCPO

Stout, Rowland. (2010). "Seeing the Anger in Someone's Face", in the *Aristotelian Society Supplementary Volume*, 84(1):29-43.

Street, Sharon. (2006). "A Darwinian Dilemma for Realist Theories of Value", *Philosophical Studies,* 127(1): 109-166.

Street, Sharon. (2008). "Constructivism About Reasons", in *Oxford Studies in Metaethics,* 3:207-245.

Street, Sharon. (forthcoming). "Does Anything Really Matter or Did We Just Evolve to Think So?" in A. Byrne, J. Cohen, G. Rosen, and S. Shiffrin (eds.), *The Norton Introduction to Philosophy.* W.W. Norton.

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

Sturgeon, Nicholas. (1986). What Difference Does it Make Whether Moral Realism is True?. *Southern Journal of Philosophy*, *24*(S1), 115-141.

Sturgeon, Nicholas. (1986). Harman on Moral Explanations of Natural Facts. *Southern Journal of Philosophy*, *24*(S1), 69-78.

Sturgeon, Nicholas. (2002). Ethical Intuitionism and Ethical Naturalism. Phillip Stratton-Lake (ed.), *Ethical Intuitionism: Re-Evaluations.* Oxford University Press.

Tait, W.W. (1986). "Truth and Proof: The Platonism of Mathematics", in *Synthese,* 69(3):455-482.

Tanaka, James W. & Lara J. Pierce. (2009). "The neural plasticity of other-race face recognition", in Cognitive, Affective, & Behavior Neuroscience, 9(1):122-131.

Tobia, Kevin, Wesley Buckwalter and Stephen Stich. (2012). "Moral Intuitions: Are Philosophers Experts?" in *Philosophical Psychology*, 1-10.

Tropman, Elizabeth. (2008). "Naturalism and the New Moral Intuitionism", in *Journal of Philosophical Research,* 33:163-184.

Tucker, Chris. (2010). "Why Open-Minded People Should Endorse Dogmatism", in *Philosophical Perspectives*, 24(1), 529-545.

Tucker, Chris. (forthcoming). Seemings and Justification: An Introduction. In C. Tucker (ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism.* Oxford University Press.

Tye, Michael. (1997). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind.* A Bradford Book.

Tye, Michael. (2000). *Consciousness, Color, and Content.* Cambridge: MIT Press.

Unger, Peter. (1996). *Living High and Letting Die: Our Illusion of Innocence.* Oxford University Press.

Van Roojen, Mark. (2006). "Knowing Enough to Disagree: A New Response to the Moral Twin Earth Argument", in *Oxford Studies in Metaethics*, *1*, 161-194.

Van Roojen, Mark. (forthcoming). "Moral Intuitionism, Experiments and Skeptical Arguments", in Anthony Booth & Darrell Rowbottom (eds.) *Intuitions.* Oxford University Press.

Vavova, Katia. (2014). "Debunking Evolutionary Debunking", in *Oxford Studies in Metaethics*, 9:76-101.

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

Väyrynen, Pekka. (2008). "Some Good and Bad News for Ethical Intuitionism", in *Philosophical Quarterly*, 58(232), 489-511.

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

Väyrynen, Pekka. (forthcoming). "Doubts about Moral Perception".

Vetter, Petra & Albert Newen. (2014). "Varieties of cognitive penetration in visual perception", in Consciousness and Cognition, 27:62-75.

Watzl, Sebastian. (2014). "Perceptual Guidance", in *Ratio,* 27(4):414-438.

Werner, Preston J. (forthcoming). "Seemings: Still Dispositions to Believe", in *Synthese.*

Whiting, Demian. (2012). "Are Emotions Perceptual Experiences of Value?", in *Ratio,* 25(1):93-107.

Wielenberg, Erik J. (2010). "On the evolutionary debunking of morality", in *Ethics,* 120(3):441-464.

Wong, Alan C.-N. & Yetta K. Wong. (2014). "Interaction between perceptual and cognitive processing well acknowledged in perceptual expertise research", in *Frontiers in Human Neuroscience*, 8(308):1-3.

Wright, Crispin. (1992). *Truth and Objectivity.* Harvard University Press.

Wright, Crispin. (2003). *Saving the Differences: Essays on Themes from Truth and Objectivity.* Harvard University Press.

Zamulinski, B. E. (2007). *Evolutionary intuitionism: a theory of the origin and nature of moral facts*. Montreal: McGill-Queen's University Press.

Zangwill, N. (2008). "Moral Dependence", in *Oxford Studies in Metaethics, Vol. 3.* Oxford University Press.

Zarinpoush, Fataneh, Martin Cooper, and Stephanie Moylan. (2000). "The Effects of Happiness and Sadness on Moral Reasoning", in *Journal of Moral Education,* 29(4):397-412.

Preston J. Werner
Department of Philosophy, Syracuse University
541 Hall of Languages, Syracuse, NY 13244
pjwerner1@gmail.com, (319)239-5942

Areas of Specialization

Ethics, Philosophy of Mind

Areas of Competence

Epistemology, Metaphysics, Political Philosophy

Education

2016   (anticipated) Syracuse University, Ph.D. (Philosophy).
Dissertation: *Seeing Right from Wrong*
Advisor: David Sobel

2013   University of Nebraska-Lincoln, M.A. (Philosophy).
Advisor: David Sobel

2009   Cleveland State University, M.A. (Philosophy).
Advisor: Diane Steinberg

2007   University of Iowa, B.A.

Publications

"Character (alone) Doesn't Count: Phenomenal Character and Narrow Intentional Content", in *American Philosophical Quarterly,* forthcoming.

"Moral Perception and the Contents of Experience", in *Journal of Moral Philosophy,* forthcoming.

"Self-ownership and non-culpable proviso violations", in *Politics, Philosophy & Economics,* Vol.14(1):67-83. February 2015.

"Seemings: still dispositions to believe", in *Synthese*, Vol. 191(8):1761-1774. May 2014.

Presentations

"Normative Asymmetry and Non-Naturalist Epistemology"
Central European University Summer Seminar on Moral Epistemology, 2014

"Comments on 'Utilitarianism, Intuitions, Rationality and Neuroscience', by Rosalind Abdool & Patricia Marino"
American Philosophical Association, 2014 Central Division Meeting

"Does Phenomenal Character Determine Intentional Content?"

American Philosophical Association, 2013 Central Division Meeting
Commentator: Esa Diaz-Leon

"Modal Semantics for Stage Theorists"
2012 UT-Austin Graduate Conference in Philosophy
Commentator: Casey Woolwine

"Nozick's Tracking Theory and the Problem of Neural Malfunction"
5th Annual University of Iowa Graduate Philosophical Society Conference
Commentator: Kristopher Philips

Employment

Teaching Assistant, Syracuse University
2013-Present

Teaching Assistant, University of Nebraska-Lincoln
2009-2013

Teaching Assistant, Cleveland State University
2007-2009

Teaching Experience

As Instructor

Theories of Knowledge and Reality [Introductory Metaphysics and Epistemology]
Summer 2015 (Syracuse University)

Introduction to Philosophy
Summer 2012 (University of Nebraska)

Philosophy and Contemporary Issues [Introduction to Ethics]
Summer 2011 (University of Nebraska)

Philosophy for 8th Graders (Co-Taught)
Spring 2014 (Southside Academy Charter School, Syracuse, NY)

As Teaching Assistant

Logic
Spring 2015 (Syracuse University)

Theories of Knowledge and Reality
Fall 2014 (Syracuse University)

Ethics in the Media Professions

Spring 2014 (Syracuse University)

Human Nature
Fall 2013 (Syracuse University)

Philosophy and Contemporary Issues [Intro to Ethics]
Fall 2009, Spring 2010, Summer 2010, Fall 2010, Spring 2011, Spring 2012
(University of Nebraska)

Introduction to Philosophy
Fall 2011 (University of Nebraska)

Introduction to Logic
Fall 2007, Spring 2008, Fall 2008, Spring 2009 (Cleveland State)

As Guest Lecturer

"Occupy Wall Street and Political Philosophy", Invited Lecture to Adam
Thompson's class on Philosophy and Contemporary Issues.
Spring 2012

Awards

Certificate in Moral Epistemology, Summer Seminar on Moral Epistemology
Summer 2014 (Central European University)

Summer Research Grant
Summer 2014 (Syracuse University)

Departmental Travel Grant
Summer 2014 (Syracuse University)

Hinman Fellowship in Philosophy
Spring 2013 (University of Nebraska)

Departmental Travel Grant
Spring 2012 (University of Nebraska)

Centennial Fellowship
2009-2010 (University of Nebraska)

Service

Referee, *Journal of Moral Philosophy*

Referee, *Synthese*

Submission Referee, Graduate Student Conference
2015, 2016 (Syracuse University)

Coordinator, Syracuse Graduate Working Papers
2013-2014 (Syracuse University)

Graduate Student Liaison, Department of Philosophy
2012-2013 (University of Nebraska)

Coordinator, Graduate Student Research Colloquia
2012-2013 (University of Nebraska)

Graduate Coursework

Independent Study (Evolutionary Debunking Arguments), David Sobel
Spring 2015

Topics in Ethics & Value Theory (Audited), Kate Manne
Spring 2015 (Cornell University)

Concepts, Kevan Edwards
Fall 2014

Intuitions and the A Priori, Andre Gallois
Fall 2014

Independent Study (Death & Emotions), Ben Bradley
Fall 2014

Practical Reason, David Sobel
Spring 2014

Philosophy of Action, Kim Frost
Spring 2014

Independent Study (Moral Epistemology), Hille Paakunainen
Spring 2014

Autonomy, Kenneth Baynes & Hille Paakunainen
Fall 2013

Hyperintensionality, Kris McDaniel
Fall 2013

Independent Study (Science and Philosophy of Perception), Robert van Gulick

Fall 2013

Advanced Metaphysics (Feminist Metaphysics), Jennifer McKitrick
Fall 2012

Ethical Theory (Decision Theory), Aaron Bronfman
Spring 2012

Advanced Epistemology (Epistemology of Perception), David Henderson
Spring 2012

Ethical Theory (Parfit's *On What Matters*), Mark van Roojen
Fall 2011

Philosophy of Language (Singular Thought), Janice Dowell
Fall 2011

Advanced Metaphysics (Individuation and Persistence) Joe Mendola
Fall 2011

Ethical Theory (Libertarianism and Self-Ownership), David Sobel
Spring 2011

Philosophy of Mind, John Gibbons
Spring 2011

Greek Philosophy (Ancient Theories of Properties), Harry Ide
Spring 2011

Philosophy of Language, Janice Dowell
Fall 2010

Advanced Metaphysics (Causation and Dispositions), Jennifer McKitrick
Fall 2010

Metaphysics, Joe Mendola
Fall 2010

Formal Logic, Reina Hayaki
Spring 2010

Ethical Theory, Mark van Roojen
Spring 2010

Ancient Philosophy, Harry Ide
Spring 2010

Advanced Ethics, David Sobel
Fall 2009

Kant, Nelson Potter
Fall 2009

Theory of Knowledge, Al Casullo
Fall 2009

Memberships

American Philosophical Association
American Association of Philosophy Teachers

References

David Sobel, Irwin and Marjorie Guttag Professor, Syracuse University

Russ Shafer-Landau, Professor, University of North Carolina at Chapel Hill

Hille Paakkunainen, Assistant Professor, Syracuse University

W. Ben Bradley, Allan and Anita Sutton Professor of Philosophy, Syracuse University

Letters of Reference may be requested from Lisa Farnsworth, Secretary, Syracuse University Department of Philosophy, lfarnswo@syr.edu.