

The Reliability Challenge in Moral Epistemology

Matt Lutz

ABSTRACT: The Reliability Challenge to moral non-naturalism has received substantial attention recently in the literature on moral epistemology. While the popularity of this particular challenge is a recent development, the challenge has a long history, as the form of this challenge can be traced back to a skeptical challenge in the philosophy of mathematics raised by Paul Benacerraf. The current Reliability Challenge is widely regarded as the most sophisticated way to develop this skeptical line of thinking, making the Reliability Challenge the strongest epistemic challenge to normative non-naturalism. In this paper, I argue that the innovations that have occurred since Benacerraf's statement of the challenge are misconceived and confused in a number of ways. The Reliability Challenge is not the most potent epistemic challenge to moral non-naturalism. The most potent challenge comes from the fact that there is a causal condition on knowledge – or, more precisely, a *because* condition – that non-natural moral facts cannot satisfy.

KEYWORDS: Explanation, Benacerraf/Field Challenge, Reliability, Modal Security, Supervenience, Moral Epistemology, Causal Theory of Knowledge

1. The Reliability Challenge

In *Taking Morality Seriously*, David Enoch writes:

[V]ery often, when we accept a normative judgment *j*, it is indeed true that *j*; and very often when we do not accept a normative judgments *j* (or at least when we reject it), it is indeed false that *j*. So there is a correlation between (what the realist takes to be) normative truths and our normative judgments. What explains this correlation? On a robustly realist view of normativity, it can't be that our normative

judgments are causally or constitutively responsible for the normative truths... And given that (at least basic) normative truths are causally inert, they are not causally responsible for our normative beliefs... And so the realist is committed to an unexplained striking correlation, and this may just be too much to believe (Enoch 2011, p. 159).

Call this the Reliability Challenge to moral non-naturalism. Moral non-naturalism is the view that moral facts are not natural facts. I define a natural fact as the kind of fact that has the power to explain our observations, and can thereby be studied by empirical methods. Moral non-naturalism, therefore, is the view that moral facts cannot (and, thus, do not) explain any of our observations. The Reliability Challenge, says Enoch, is the “*strongest* version of the epistemological challenge [to non-naturalist moral realism]” (p. 163, italics in original). Many moral theorists have followed him in this assessment.¹ In recent years, the Reliability Challenge has become one of the most popular topics in moral epistemology.

Enoch's Reliability Challenge became influential for two reasons. First, Enoch proposed the Reliability Challenge as an interpretation of Sharon Street's “Darwinian Dilemma.” Street's debunking argument was published three years before Enoch's first presentation of the Reliability Challenge, and by the time of Enoch's writing, it was beginning to draw substantial interest. But Street's Darwinian Dilemma is somewhat vague – Street stresses that, if moral realism is true, it would be only a *coincidence* if our moral beliefs are all, or even mostly, true.² But what is a “coincidence,” in the relevant sense, and why should such a coincidence be epistemically troublesome?³ Enoch provided an answer: A coincidence is a reliable correlation that can't be explained. In Enoch's estimation, the Darwinian Dilemma uses Darwin to draw our attention to the *real* problem: that there's no explanation

¹ Schechter (2017).

² See also Bedke (2009).

³ cf. White (2010), Setiya (2012, Ch. 2).

of the reliability of our beliefs.

Enoch's Reliability Challenge became popular for another reason as well: it's based on a familiar skeptical challenge for mathematical realism, commonly known as the "Benacerraf/Field Challenge." Exactly what the Benacerraf/Field Challenge says is a complicated matter. But as a first pass: the Benacerraf/Field Challenge is Enoch's Reliability Challenge, applied to mathematical beliefs rather than normative beliefs; if mathematical realism is true, then there's no way to explain the reliability of our mathematical beliefs.

Thus, Enoch argued that a skeptical challenge in mathematical epistemology could generalize to create a skeptical challenge in moral epistemology.⁴ That skeptical challenge has a strong pedigree and seems to be the best way to make sense of Street's Darwinian Dilemma.⁵ Since then, the literature on moral epistemology has been consumed with the question of whether the reliability of our moral beliefs can be explained. Enoch himself argued that they can be explained.⁶ Skeptics argue that these explanations are all obviously question-begging.⁷ Anti-skeptics respond by attempting to offer non-question-begging explanations of the reliability of our moral beliefs,⁸ or else by arguing that there's nothing wrong with offering a question-begging explanation in this context.⁹ Skeptics are unmoved by any of the explanations that anti-skeptics offer up. All of the anti-skeptic's explanations of reliability seem to rely on some kind of trick;¹⁰ surely the reliability of our moral beliefs can't be explained *in the right kind of way*.¹¹ (The trick is to say what the right kind of way is.)

This focus on the Reliability Challenge in recent years is a testament to Enoch's claim that the Reliability Challenge is the strongest epistemic challenge for normative non-naturalism. It certainly

⁴ Huemer (2005, p. 123-127) was the first to notice that the Benacerraf Problem could generalize in this way.

⁵ Although it is probably *not* the best way to interpret Joyce (2006)'s evolutionary debunking argument.

⁶ Enoch (2011, Ch. 7). See also Skarsaune (2011), Wielenberg (2010),

⁷ Street (2016)

⁸ Behrends (2013)

⁹ Vavova (2014).

¹⁰ Bedke (2014).

¹¹ See Setiya (2012, Ch. 3), Lutz (2015, Ch. 3), Lutz (2017), Faraci (forthcoming). Compare Lutz and Ross (2017).

looks like *one of* the most potent epistemic challenges for normative non-naturalism. But it is not.

In what follows, I will argue that, on a straightforward reading of the Reliability Challenge, this challenge is a kind of metaphysical challenge rather than an epistemic challenge (Section 2). Of course, there are ways to understand the Reliability Challenge as presenting an epistemic challenge. But, as I argue in Section 3, every way of thinking of the Reliability Challenge as an epistemic challenge either involves endorsing a clearly incorrect account of knowledge, or else endorsing a causal condition on knowledge. And if skeptics who press the Reliability Challenge want to endorse a causal condition on knowledge, then the focus of our inquiry needs to shift: the most important question in moral epistemology is not whether the reliability of our moral beliefs can be explained. It is whether there is a causal condition on knowledge and, if there is, what form such a condition takes. In Section 4, I argue that there is such a condition on knowledge – although it is a *becausal* condition rather than a *causal* condition. *Contra* Enoch, it is the non-naturalist’s inability to satisfy this condition that poses the strongest challenge to moral non-naturalism.

2. Skeptical Dilemmas and Explanatory Deficits

Epistemic challenges to a philosophical view – as opposed to metaphysical challenges, linguistic challenges, etc. – attempt to show that a view has certain unpalatable epistemic commitments (as opposed to metaphysical commitments, linguistic commitments, etc.). One common kind of epistemic challenge (which will be of particular interest to us here) is a skeptical dilemma.

Skeptical Dilemma: An argument that purports to show that a certain theory of Fs entails skepticism about Fs.

Skeptical dilemmas force a choice – either accept skepticism about Fs, or abandon the theory of Fs that gives rise to the dilemma. A general argument for global skepticism or moral skepticism would not constitute an epistemic challenge *for moral non-naturalism*. But an argument that non-naturalism’s

truth would render moral knowledge impossible is an epistemic challenge for non-naturalism. If moral non-naturalism entails moral skepticism while other metaethical views do not, then any reason we have not to be skeptics is a reason to reject moral non-naturalism.

2.1 Benacerraf's Challenge

Benacerraf's original version of the Benacerraf/Field Challenge is a skeptical dilemma; it argues for the conclusion that mathematical knowledge is impossible if mathematical realism is true. Benacerraf was writing at a time when the causal theory of knowledge was widely accepted. According to the causal theory of knowledge, S knows that P just in case there is an appropriate causal connection between the fact that P and S's belief that P.¹² But if mathematical realism is true, our mathematical language refers to abstract (and thus causally inert) objects: *numbers*. So how, worried Benacerraf, could we know any mathematical facts – facts *about numbers* – when there are no causal connections between facts about numbers and our beliefs about numbers?

Benacerraf's Challenge

1. S knows that P if and only if there is an appropriate causal connection between S's belief that P and the fact that P. (The Causal Theory of Knowledge)
2. If mathematical realism is true, there are, necessarily, no causal connections between mathematical facts and anything else.
3. Therefore, if mathematical realism is true, mathematical knowledge is impossible.

This is a problem for mathematical realism and it is, specifically, a skeptical dilemma. Benacerraf's Challenge says that there is a particular necessary condition on knowledge that cannot be satisfied – *but only if realism is true!* Thus, mathematical knowledge is possible only if mathematical realism is false. (It should be noted that Benacerraf did not endorse this conclusion; he argues only that the challenge is an important puzzle for realists like him.)

¹² Goldman (1967).

Premise 2 of Benacerraf's Challenge is hard to resist – numbers are abstract, and thus causally inert – so that means that everything rests on Premise 1. Fortunately for mathematical realists everywhere, the causal theory of knowledge as presented in Premise 1 is a very strong claim, and is now almost universally rejected.

There are many objections to the causal theory of knowledge. But it is a causal connection as a *necessary* condition on knowledge that is a premise in Benacerraf's Challenge, so we will look only to reasons to think that there can be knowledge without a causal connection between belief and truth. Two are particularly relevant for our purposes here. The first concerns knowledge of the future. Current facts cause future facts, not the other way around. So it's impossible for facts about the future to cause our current beliefs, and thus impossible to have knowledge of the future if the causal theory of knowledge is true. But knowledge of the future is possible.¹³ The second concerns *a priori* knowledge. While it may be plausible that there is a causal condition on empirical knowledge, it seems implausible to say that there is a causal condition on *a priori* knowledge.¹⁴

These two objections are obvious enough that Alvin Goldman, in formulating the causal theory of knowledge, anticipated these objections and had responses prepared. In response to the knowledge of the future objection, Goldman proposed that two causal chains that both proceed forward from some common cause suffice to satisfy the causal condition. Thus, we can have knowledge of the future if there is some common cause that explains both my belief that P and the fact that P. (Some terminology: Call chains of causal connections that proceed from directly from one fact to another **direct causal connections**. When there is no direct causal chain between two facts, but some third fact serves as a common cause of those two facts, we will say that they are linked by an **indirect causal connection**.) Second, in response to the *a priori* knowledge objection, Goldman sounds a full retreat, arguing that the causal theory is only appropriate for empirical knowledge. *A priori* knowledge, Goldman claims, is

¹³ White (2010).

¹⁴ Linnebo (2006).

nothing more than JTB. Let's call Goldman's responses to these two objections *the traditional responses*.

The traditional responses might serve to answer at least these two objections to the causal theory of knowledge, but they provide no comfort to the skeptic about math or morals for, in accepting the traditional responses, the skeptic's case is undermined. The traditional response to the *a priori* knowledge objection concedes that there is no causal condition on *a priori* knowledge. Mathematical and moral knowledge are paradigmatic kinds of *a priori* knowledge (if they exist). So it's no mark against them if there is no causal connection between belief and fact. Similarly, the traditional response to the knowledge of the future objection gives the moral non-naturalist (if not the mathematical realist) another way to evade the skeptical challenge. Even if non-natural moral facts *cause* nothing, if non-natural moral facts *are caused* there may exist *indirect* causal connections between moral facts and our moral beliefs. And if indirect causal connections suffice to satisfy the causal condition, then moral facts may satisfy the causal condition.¹⁵ Thus, the two traditional responses save the causal theory at the cost of its skeptical bite.

2.2 Field's Challenge

Recognizing the problems with the causal theory, Field famously attempted to update Benacerraf's Challenge by replacing the causal theory of knowledge with a less controversial epistemic framework. "The way to understand Benacerraf's challenge, I think," writes Field, "is not as a challenge to our ability to *justify* our mathematical beliefs, but as a challenge to our ability to *explain the reliability* of these beliefs" (Field 1989, p. 25, emphasis in original). "If *it appears in principle impossible to explain this*, then that tends to *undermine* the belief in mathematical entities, *despite* whatever reason we might have for believing them" (ibid, p. 26). Philosophers who attempt to explain this reliability tend to invoke a reliable faculty of mathematical intuition in order to carry out this

¹⁵ This is a crucial flaw with the skeptical challenge posed by Setiya (2012, Ch. 3).

explanation. “But special 'reliability relations' between the mathematical realm and the belief states of mathematicians seem altogether too much to swallow. It is rather as if someone claimed that his or her belief states about a remote village in Nepal were nearly all disquotationally true, despite the absence of any mechanism to explain the correlation between those belief states and the happenings in the village. Surely we should accept this only as a last resort” (ibid, p. 26-7).

Field emphasizes that his argument does not rely on the causal theory of knowledge. There is no causal condition on knowledge, says Field, but we should still be worried about the lack of causal connections between facts and beliefs. This is because causal connections are our only way of explaining the reliability of our mathematical beliefs. Mathematical realists cannot explain the crucial explanandum.

Unlike Benacerraf's Challenge, Field's Challenge is not a skeptical dilemma. Benacerraf attempted to identify a necessary condition on knowledge such that our mathematical beliefs cannot satisfy that condition if realism is true. Field, on the other hand, notes that a certain feature of mathematical facts – specifically, their causal impotence – makes it difficult for mathematical realists to explain a certain kind of correlation. This is a theoretical deficit for mathematical realism.

Thus, rather than presenting a skeptical dilemma, Field's Challenge has an altogether different structure, which was outlined by Baras (2017):

(1) **Non-Skeptical Robust Realism:** The robust realist [about some domain D] has the following commitments:

- (a) **Realism:** There are D truths.
- (b) **Cognitivism:** We have D beliefs.
- (c) **Mind-Independence:** Our D beliefs do not cause or constitute the D truths.
- (d) **No-Causation:** The D truths do not cause or constitute our D beliefs.
- (e) **Non-Skepticism:** There is a correlation between the D truths and our D

beliefs.

(2) **Striking Correlation:** The correlation (e) is striking.

(3) **Unexplainable Correlation:** If Non-Skeptical Robust Realism is correct then the correlation (e) is apparently unexplainable. ((c)&(d) rule out the most obvious kinds of explanation.)

(4) **Epistemic Principle:** If a theory implies that there is an apparently unexplainable striking correlation, that is a strong reason to believe that the theory is false.

Therefore,

Conclusion: We have strong reason to reject Non-Skeptical Robust Realism (at least one of the commitments (a) through (e)).

Baras notes that Field's Challenge raises skeptical worries for *non-causal realism* about any domain of facts, for non-causal realism seems to leave us without a way to explain the reliability of our beliefs. This gives us a reason to reject non-causal realism in any domain. In addition to mathematics and metaethics, Joshua Schechter has shown that we can make a similar argument for skepticism about logic.¹⁶ Modal metaphysics seems suspect as well.

But if this is how we are to understand the Reliability Challenge, it doesn't look like an epistemic challenge at all. None of the premises of Baras's reconstruction of the Reliability Challenge invoke any epistemic notions; Baras's argument is not about justification or non-accidentally true belief. While Baras labels Premise 4 as his "Epistemic Principle," Premise 4 is epistemic only insofar as it states that we have a *strong reason to believe* that a theory is false. But providing a strong reason to believe that a theory is false is not a *distinctive* feature of *epistemic* challenges. It's a *general* feature of *challenges*. Any argument against a theory will, if successful, provide strong reason to think that the

¹⁶ Schechter (2013)

theory is false.

The Reliability Challenge is a *metaphysical challenge*, not an epistemic challenge. To see this, consider another paradigmatic metaphysical challenge to moral non-naturalism: the “Supervenience Challenge.” According to the Supervenience Challenge, moral facts supervene on natural base facts. But non-naturalists can't give any account of why moral facts supervene in this way. This explanatory shortcoming of non-naturalism – its inability to explain a striking phenomenon – is a strong reason to reject non-naturalism. Or so the argument goes.¹⁷ Whether or not this argument succeeds, it is noteworthy that the Supervenience Challenge and the Reliability Challenge have exactly the same structure. There is a striking explanandum; the non-naturalist (apparently) lacks the theoretical resources to explain it; this reflects poorly on non-naturalism.

The similarities between the Supervenience Challenge and the Reliability Challenge run deeper than this. The two challenges are both instances of a more general explanatory problem. Both challenges begin with the same observation: that moral facts are strongly correlated with facts of some other kind. The challenges concern different kinds of correlations: the Supervenience Challenge concerns modal covariation across possible worlds, while the Reliability Challenge concerns reliably true beliefs in the actual world. But in both cases, the fundamental worry is the same: It's *in principle* impossible to explain these correlations for non-causal realists. If moral facts aren't causally connected to anything else and are not identical to anything else, we should expect that they wouldn't be correlated with anything else. (Causation isn't correlation, but the two tend to travel together). So when we encounter the striking correlation, that looks bad for the theory that can't explain it. This is not a specific challenge to the possibility of moral knowledge or justified moral belief. It is, instead, a general theoretical challenge to the explanatory adequacy of the metaphysical theory of moral non-naturalism.

¹⁷ McPherson (2012).

Given the structural similarities between the Supervenience Challenge and the Reliability Challenge, it is unsurprising that non-naturalists attempt to answer both challenges in the same way. Erik Wielenberg (2014) argues that moral non-naturalism can be saved from both of these metaphysical explanatory challenges by positing a normative determination relation that holds between natural facts and normative facts. He calls this a “making” relation, and says that it is a kind of causal relation. Other philosophers have characterized this relation differently,¹⁸ but the common thread is that non-naturalists appeal to this kind of determination relation to illuminate the common-sense idea that moral facts obtain *because* natural facts obtain.¹⁹ And if there is a normative determination relation that holds between the natural facts and the normative facts, then both explanatory challenges can be answered in the same way: the normative determination relation is what explains the relevant correlation. Even though moral facts are causally inert, they are causally connected to natural facts not in virtue of *being causes*, but in virtue of *being caused* (or determined in a non-causal way).²⁰

This normative determination relation explains why, in all possible worlds in which this law holds, there is a correlation between the natural facts and the normative facts. This same determination relation also features in “third factor” explanations of the reliability of our moral beliefs.²¹ If the same natural facts which (normatively) determine the moral facts are also part of what (causally) determine our moral beliefs, then an indirect causal connection exists between our moral beliefs and the moral facts, and this indirect connection explains the reliability of our moral beliefs.

Skeptics can object that the moral realist has given us no reason to think that there are moral laws like this. The non-naturalist may reply that the skeptic has not yet given any reason to think that there are not. After all, the objection that has been offered thus far is not an epistemic objection: it is a metaphysical objection. If the skeptic is asking how our moral beliefs can be justified, the non-

¹⁸ Enoch (2011, p. 133ff), Bader (2017).

¹⁹ Mackie (1977, Ch. 1). Compare Olson (2014, p. 88-100).

²⁰ Compare Harman and Thomson (1996, Ch. 6).

²¹ See, e.g., Enoch (2011), Skarsaune (2011), Wilenberg (2010), Wilenberg (2014), Behrends (2013), Schechter (2017).

naturalist may talk about reflective equilibrium or intuitions or deliberative indispensability or anything else. But the skeptic's charge is not that our beliefs in moral laws are unjustified, it is that the reliability of our moral beliefs cannot be explained (Field is explicit about this). And, the non-naturalist may point out, it's simply not true that the reliability of our moral beliefs can't be explained: by appealing to the normative laws, I just explained the reliability of our moral beliefs. Thus, that reliability *can* be explained.²² QED.

So if we understand the Reliability Challenge as being on par with the Supervenience Challenge, it is a metaphysical challenge that be answered in the same way that the Supervenience Challenge can be answered.

3. Explaining Reliability

While the most natural reading of the Reliability Challenge is as a metaphysical challenge, the Reliability Challenge is intended to be an epistemic challenge. In this section, I'll attempt to reconstruct the Reliability Challenge as the potent epistemic challenge it is supposed to be. We'll look at two possible reconstructions of the Reliability Challenge as an epistemic challenge, and see that neither reconstruction is very strong.

3.1 External Conditions on Knowledge

Although Enoch's primary statement of the Reliability Challenge is the metaphysical challenge discussed in the previous section, Enoch notes that epistemic internalists might demand that the Reliability Challenge be phrased as a challenge to justification. And, he claims, this demand can be met. Even though "the challenge does start from an externalist perspective," it can be "internalized," thereby making the Reliability Challenge into a challenge to the justification of our beliefs (Enoch 2011, p. 161). This is via the following principle:

²² Baras (2017).

Undercutting: If it is impossible to explain the reliability of S's beliefs, and S learns this, then this is sufficient to undercut S's justification in those beliefs.

Undercutting seems plausible, but it is *ad hoc*. What makes Undercutting the kind of principle that we should believe in? Enoch's claim that a challenge that starts from an "externalist perspective" can be "internalized" suggests the following principle:

Internalization: If F is a necessary, external condition on knowledge, and S's belief that P lacks F, then, if S learns that her belief lacks F, S is no longer justified in believing P.²³

Internalization is a principle that is motivated by the literature on the Gettier Problem. In so-called Gettier cases, subjects have justified, true beliefs that do not amount to knowledge because their beliefs are true by accident. It is controversial exactly what it means for a belief to be "true by accident" in a way that destroys knowledge, but we can use "non-accidentally true" as a technical term to pick out that feature that all Gettiered beliefs have in common. This non-accidentality feature seems to be, in some sense, *external* to the agent, since it is possible for two subjects to be internal duplicates such that one subject knows P, and the other subject has a belief that is true only by accident (Williamson 2000, Chapter 3).

As "no defeater" accounts of knowledge emphasize, Gettier cases always seem to include a "but unbeknownst to the subject..." twist. That unbeknownst factor is a *defeater*: it is both the feature that makes the belief non-accidentally true, and also something such that, if the subject knew it to be true (if the unbeknownst claim were no longer unbeknownst), the subject would no longer be justified in her beliefs.²⁴ These claims about the nature of defeat and non-accidentality are not uncontroversial, but they provide a plausible theoretical justification for the Internalization principle. The "necessary external condition on knowledge" mentioned in Internalization is the factor(s) that guarantees that the

²³ This formulation of the principle is my own.

²⁴ Lehrer and Paxson (1969).

subject's belief is non-accidentally true. The lack of this necessary external condition on knowledge is, therefore, a defeater. Accordingly, learning about this defeater will thereby defeat one's beliefs. And that's precisely what Internalization says. So Internalization plausibly states a sufficient condition for undercutting defeat.

Internalization is also plausible as a necessary condition on defeat, as well: If a factor F is not required for knowledge, why would learning that you lack F give you a reason to abandon an otherwise justified belief that P?

Thus, if there is some necessary, external condition on knowledge that is unsatisfied for our moral knowledge (if non-naturalism is true), Internalization tells us that learning about this lack will lead to moral skepticism.

3.2 Reliability, Non-accidentality, and Sets of Beliefs

If we are to understand the Reliability Challenge as an epistemic challenge by way of Internalization, then a simple understanding of the Reliability Challenge presents itself.

Explainable Reliability is Necessary (ERN): S knows that P only if the reliability of S's belief that P is explainable.

ERN, together with Internalization, entails Enoch's Undercutting.

In ERN, what does *reliability* consist in? Just this: That most of our moral beliefs are true. Says Enoch: "A class of beliefs is reliable, I take it, if and only if a sufficiently large portion of it is true" (Enoch 2011, p. 155). That is, in the *actual* world, the *set* of our moral beliefs is largely true. Call this Actual Set Reliability (AS-Reliability). The AS-Reliability of the set of our moral beliefs is the surprising fact that Enoch thinks needs an explanation. Enoch is not alone in being concerned about the AS-Reliability of our beliefs. Field's striking explanandum is the AS-Reliability of our mathematical beliefs. Schecter's striking explanandum is the AS-Reliability of our logical beliefs.

But AS-Reliability is not the only conception of reliability. We can also talk about the reliability

of *methods*. A *method* is reliable just in case it generates the correct results *across a variety of counterfactual situations*. Call this Modal Method Reliability (MM-Reliability). For instance, a sharpshooter's method of aiming is (MM-)reliable just in case that method leads the sharpshooter to hit her target in a wide variety of counterfactual situations. And, similarly, a method of *belief formation* is (MM-)reliable just in case that method of belief formation leads to the formation of true beliefs in a wide variety of counterfactual circumstances.

MM-Reliability and AS-Reliability are not the same thing. A sharpshooter can be AS-Reliable without being MM-Reliable if he hits the target on each of his six shots, but each bullseye is purely a matter of luck. And a sharpshooter can be MM-Reliable without being AS-Reliable if she possesses the skill to hit the bulls-eye most of the time, but strange exogenous factors (like random gusts of wind) cause the sharpshooter to miss most of her shots in the actual world.²⁵

Note that MM-Reliability is a property of *methods* and AS-Reliability is a property of *sets*. Neither is a property of *particular beliefs*. Particular beliefs may be true or false. To call a particular belief reliable or unreliable would be a category mistake.

How, then, do we understand ERN? Because a single belief isn't the kind of thing that can be reliable or unreliable, ERN must concern either the AS-reliability of some set of S's beliefs, or else it must concern the MM-reliability of S's method of belief-formation. As noted, the most straightforward way to read those concerned with the Reliability Challenge is as concerned with the AS-reliability of our moral beliefs. This would make the key premise of the Reliability Challenge:

Actual Reliability is Necessary (ARN): If S knows that P, where P is a proposition about subject matter M, then S's beliefs about M are mostly true, and there is an explanation of this AS-reliability.

But ARN is an unacceptable epistemic principle. First, there is the question of how to

²⁵ Compare Bonjour and Sosa (2003, p. 100-107).

individuate sets of beliefs by their subject matter. My belief that President Kennedy was shot in Dallas in 1963 is a belief that I have about Kennedy, but it is also a belief about shootings, about Presidents, about Dallas, and about 1963. Which of these sets must be mostly true if I am to know that Kennedy was shot? But even if this problem can be resolved, the AS-reliability of sets of beliefs (where those sets are individuated by subject matter) is not relevant to knowledge because it's possible to know things about a subject even if one's beliefs about the subject as a whole are unreliable. Someone's believing a huge number of conspiracy theories about the Kennedy assassination does not prevent them from knowing that Kennedy was shot.

3.3 Non-accidentality and MM-Reliability

The lesson of the failure of ARN is that, while the AS-Reliability of our moral beliefs might be an important explanandum in a metaphysical challenge, AS-Reliability is not an epistemically important concept. But MM-Reliability may be. So:

Modal Reliability Challenge (MRC)

1. S knows that P (for some moral proposition P) only if S's belief that P was formed by a MM-reliable belief-forming process.
2. If moral non-naturalism is true, S's belief that P was not formed by a MM-reliable belief-forming process.
3. Therefore, if moral non-naturalism is true, S does not know P (for any moral proposition P).

Note that this interpretation takes us very far from the original statement of the argument, which begins from the observation that our beliefs *are reliable*, and that this fact is striking enough to stand in need of some observation. Now the challenge is that our beliefs *are not reliable at all*, and that this indicates that they cannot be non-accidentally true.

But while we are rather far from the original strategy of talking about striking and unexplained

correlations, it is not hard to trace a line from MRC to the original statement of the Reliability Challenge. It is plausible that the AS-reliability of a set of beliefs can be explained in a satisfactory way only if those beliefs are formed by an MM-reliable method. Thus, the fact that the AS-reliability of our moral beliefs is unexplained counts as evidence that our moral beliefs are not formed by an MM-reliable process. If it is a necessary external condition on knowledge that one's belief be formed by an MM-reliable method, this would, by Internalization, entail that our moral beliefs are unjustified. In this way, MRC promises to capture the thought that underlies Enoch's Undercutting principle.²⁶

Clarke-Doane (2016) endorses a version of MRC as the strongest version of the skeptical challenge.²⁷ Clarke-Doane accepts the epistemic principles of Safety and Sensitivity as illuminating our notion of MM-reliability and thereby non-accidentality.

Safety: If S knows that P by coming to believe P using method M, then in all nearby possible worlds in which S forms a belief using M, that belief is true.

Sensitivity: If S knows that P by coming to believe P using method M, then in the nearest possible world in which P is false, S does not believe P via M.

By Internalization, learning that one's beliefs are either unsafe or insensitive will be a defeater for one's beliefs. And Clarke-Doane endorses exactly this conclusion:

Modal Security: If information E undermines all of our beliefs of a kind F, then it does so by giving us reason to doubt that our F-beliefs are both sensitive and safe (Clarke-Doane 2016, p. 30)

Given this setup, Clarke-Doane then argues that our beliefs are modally secure, and thus that the skeptical challenge can be answered.

There are two major problems with MRC. First, there are numerous counter-examples to both

²⁶ I think that Enoch means to endorse MRC when he writes that "given the purported unexplainability of her [i.e. the non-naturalist realist's] reliability of our normative beliefs, she should come to believe that they are unreliable" (Enoch 2011, p. 162).

²⁷ As does Schechter (2017).

Safety and Sensitivity as necessary, external conditions on knowledge. For instance, the following case is a counter-example to Sensitivity:

Flagpole: Bill frequently slightly overestimates the height of things by several inches. He sees a flagpole that is 40 feet tall, and judges that it is more than 10 feet tall. He clearly knows this. Yet in the nearest possible world where his belief is false (i.e. the flagpole is just under 10 feet tall), he falsely believes that it is more than 10 feet tall.²⁸

And this is a counter-example to Safety:

Bank Hacker: A hacker has hacked into your bank, and is prepared to run a program that will steal all the money from your account, while leaving behind a dummy database so that anyone who checks the bank website won't notice the money is missing (at least not for a few days – enough time to get away). Just before running the program, the hacker wins the lottery, and so decides not to steal the money. You log onto the bank website and check your account balance. You know how much money is in your account. But in all nearby possible worlds you'd believe the same thing by the same method, and you'd be wrong.²⁹

Second, it is not enough for one's belief to reliably track the truth across nearby possible worlds.

It seems that there also must be some sort of *connection* between belief and fact, or else beliefs that are formed by MM-Reliable processes will still seem to be true by accident.³⁰ Consider:

Hypnotist: Connie has been hypnotized by a stage magician to form the belief that methane is a compound of carbon and hydrogen when she hears a bell ring. (The hypnotist chose this claim at random, just to demonstrate his hypnotic powers. He

²⁸ Williamson (2000).

²⁹ This is a modified version of a case from Comesana (2005).

³⁰ Yamada (2012); Setiya (2012, Ch. 3); Lutz and Ross (2017); Faraci (forthcoming).

does not know that it is true; in fact, he thinks it is false.) Since Connie is strongly susceptible to hypnotism, after she awakes, she comes to believe that methane is a compound of carbon and hydrogen when she hears a bell.³¹

The hypnotist has given Connie a disposition to form a certain belief in response to a certain stimulus, and this disposition will reliably result in Connie's forming true beliefs; Connie's method of belief-formation can only result in the formation of a belief in one particular proposition, and that proposition is necessarily true. Thus, Connie's belief in the composition of methane will be both Safe and Sensitive; her method is MM-reliable. Yet it is quite clear that Connie's belief is only accidentally true.³²

One might object that this misdescribes the case; Connie's method should be taken to include the circumstances of her hypnotism, and thus her method is modally unreliable, as the hypnotist could have hypnotized her to believe something different. But we should distinguish between the circumstances that lead to the adoption of a method and the deliverances of the method itself,³³ and Hypnotist is described in order to cleanly separate the method (response to bell-ringing) from the cause of adopting that method (hypnotism). I share the intuition that Connie's belief could easily have been wrong. But, on reflection, it seems to me that I have this intuition because Connie's belief has nothing to do with the fact that carbon really is a compound of carbon and hydrogen. She does not believe this *because* it's true, but for some other reason. Thus, Hypnotist suggests that there is a further necessary condition on non-accidentality that is not modal, but categorical; the belief must *actually be connected* to the truth (e.g., causally).³⁴

³¹ This case is inspired by Yamada (2012)'s "Two Scales" case.

³² One might object that Connie's belief is *unjustified* in this case, and this lack of justification explains her lack of knowledge. If you're worried about this, assume further that the hypnotist's suggestion gives Connie not only the belief, but also the *intuition* that the belief is true. (See Huemer (2005, Ch. 5) for more on intuitions and justification.) Even with such an intuition, it is still clear that she lacks knowledge.

³³ Yamada (2012).

³⁴ See Yamada (2012), Setiya (2014, Ch. 3), Lutz (2015, Ch. 3), Lutz and Ross (2017), Faraci (forthcoming), and Lutz (ms) for more on this point.

Non-naturalists (like Clarke-Doane) argue in response that no such connection is required. After all, we already saw the problems with supposing that we can extract a skeptical conclusion from a causal condition on knowledge. The strongest necessary condition on knowledge that non-naturalists like Clarke-Doane will concede is a modal condition on knowledge.

We arrive now at a dialectical stalemate. Non-naturalists only propose necessary conditions on knowledge that are weak enough that lots of facts can satisfy them – including non-natural moral facts. The necessary conditions they propose seem far too weak to fully account for non-accidentality, but non-naturalists aren't trying to state sufficient conditions on non-accidentality; to propose necessary and sufficient conditions on non-accidentally true belief is to attempt to solve the Gettier problem, a daunting task that non-naturalists are happy to leave to skeptics. Non-naturalists thereby assume an enviable defensive position, rebutting the over-reaching epistemic principles proposed by skeptics. (Of course, it is controversial whether or not non-naturalists are entitled to assume this defensive position.)³⁵

This dialectical stalemate highlights a second, more fundamental problem for MRC: that there is no obvious way for a skeptic about non-causal realism to avoid endorsing a causal theory of knowledge. A **logical gap** remains between the source of the skeptical concerns – non-causal realism – and the skeptical conclusion – that some necessary condition on knowledge is unsatisfied. We can present this gap as a dilemma. Either the MM-Reliability of one's belief-forming methods can *only* be secured by the presence of a causal connection between belief and truth, or the reliability condition is weaker than this, and can be satisfied in non-causal ways. The first horn of the dilemma amounts to an endorsement of the causal theory of knowledge; but we've already seen the problems that the skeptic encounters in endorsing a causal theory. Those problems have seemed pressing enough that most skeptics have, following Field, embraced the second horn of the dilemma.

³⁵ See, e.g., McGrath (2014) and Woods (2016).

But on the second horn of the dilemma, the skeptic must articulate some weaker condition on knowledge that is still strong enough to generate the requisite skeptical conclusion. But what condition could fulfill both of those desiderata? If it is not the case that knowing P requires a causal connection between the belief that P and the fact that P, why should it be the case that P's causal inertness precludes knowledge of P? On the second horn of the dilemma, Premise 2 of MRC is entirely unmotivated; non-naturalists reject it with a shrug.³⁶

What do we make, then, of the key question of the Reliability Challenge: Can the AS-reliability of our beliefs be explained? The most straightforward way to understand this question is as posing a metaphysical challenge, since AS-Reliability is not an important epistemic concept. And if we understand the challenge along the lines of MRC, the focus of the dialectic shifts to the question of whether a causal connection is a necessary condition on non-accidentally true belief.

4. An Explanatory Condition

The problem for the Reliability Challenge is that its advocates have too readily followed Field in holding that the essence of the skeptical challenge has something to do with reliability. Instead, they should have followed Benaceraff. There is a causal condition on knowledge – or, more accurately, a *becausal* condition – that moral non-naturalists cannot satisfy.

Explanatory Challenge

1. **Explanatory Constraint (EC):** S has non-inferential knowledge that P only if there is a direct explanatory connection between the non-factive mental states on which S bases her belief that P and the fact that P.
2. Non-natural moral facts never directly explain our non-factive mental states.
3. Therefore, non-inferential knowledge of non-natural moral facts is impossible for us.

³⁶ White (2010), Clarke-Doane (2015).

The Explanatory Challenge is a skeptical dilemma, similar to Benacerraf's Challenge. As with Benacerraf's Challenge, the second premise of the Explanatory Challenge is true by definition of 'non-natural.' A fact is non-natural just in case it cannot explain our observations and cannot thereby be explored by empirical methods. Thus, because observations are a kind of non-factive mental state, non-natural facts cannot explain our non-factive mental states. This means that, as with Benacerraf's Challenge, everything turns on the first premise: EC.

EC is similar to the old causal theory of knowledge, only it requires an *explanatory* connection rather than, specifically, a *causal* connection between belief and fact. It also requires that this connection run through some other mental states of the subject. In this section, we'll see why these differences are significant.

4.1 The Case for EC

There are two reasons to accept EC. First, EC states a necessary, external condition on knowledge. It follows from Internalization that, if a subject comes to learn that EC is unsatisfied for one of her beliefs, that belief is unjustified. Thus, if a subject comes to learn that the evidence that she took to support a non-inferential belief that P is not explained by P, that belief is unjustified. This is an independently plausible conception of undercutting defeat. We're no longer justified in believing P on the basis of evidence E if evidence E is *explained away*. In a recent paper, I've defended that conception of undercutting defeat in the form of the principle:

Explaining Away Defeats (EAD): New evidence, D, defeats the support that E provides for S's belief that P *if*: D is evidence in favor of a complete explanation, A, of E, such that S may not infer P from A and S's independent background information.

where background information is "independent" just in case it is not based on either E or D.³⁷ Korman

³⁷ Lutz (2017, p. 8)

and Locke (this volume) defend a very similar principle.

Internalization, EAD, and EC are coherent and mutually supporting. If it's a necessary condition on a non-inferential belief that P that one's evidence for P be explained by P (per EC), then, by Internalization, if you learn that your belief that P is not explained by the fact that P, your belief that P is either unjustified or else an (independently) inferentially-justified belief. And that's just what EAD says. EAD is motivated by an analysis of the notion of *undercutting defeat*, rather than by an analysis of the notion of non-accidentally true belief. Thus, an independently plausible principle (EAD) can be shown to follow from EC and Internalization. And EC follows from Internalization and EAD if Internalization gives necessary as well a sufficient conditions for undercutting defeat, which it plausibly does. This is good reason to think that EC is on the right track.

Second, the need for some sort of explanatory constraint on knowledge was demonstrated in Section 3.3. It's not that a lack of an explanatory connection is an indication of some other epistemic deficit, as Field, Enoch, and Schechter would have it. The explanatory connection *is itself* a necessary condition on knowledge.³⁸ A modal connection between beliefs and facts that is not underwritten by an explanatory connection between belief and fact seems sufficient to make that belief accidentally true, thereby destroying knowledge.

Thus, both the conditions of undercutting defeat and internalization and concerns about non-accidentality serve to motivate EC. Indeed, it is these very concerns about non-accidentality that motivated the original causal theory of knowledge. But as we've seen, the causal theory of knowledge is subject to strong counterexamples, and the traditional responses to those counterexamples are both unsatisfying and deprive the causal theory of any potential skeptical bite. So it will count strongly in favor of EC if it does a better job at answering these traditional objections.

4.2 EC and Future Knowledge

³⁸ Korman and Locke (this volume).

Consider first the knowledge of the future objection. There are no direct explanatory connections between future facts and current beliefs, so if a direct explanatory connection between fact and belief is required for knowledge, then knowledge of the future is impossible. The traditional response is to say that indirect explanatory connections can secure non-accidentality.

But there are two problems with this response. First, indirect connections are omnipresent. All that is required for two facts to be indirectly causally connected is for those two facts to have a common cause. But all facts have a common cause, if we are willing to go back far enough: the Big Bang. Because all facts are (ultimately) caused by the Big Bang, all facts are indirectly causally connected. The Big Bang theory might turn out to be false, but the point is a general one; if we go back far enough in time, we will find that lots of seemingly-completely-unrelated events have distant common causes and are therefore indirectly connected. It follows that a requirement of indirect explanatory connections is satisfied by all beliefs. The problem with this is not that an indirect connections view is too weak to provide sufficient conditions for non-accidentality – it's not my goal in this paper to provide necessary and sufficient conditions for non-accidentality and thereby solve the Gettier problem. It's that the view is too weak to account for the accidental truth of *any* accidentally-true beliefs. (Smith's belief that either Jones owns a Ford or Brown is in Barcelona is indirectly connected to the fact that Brown is in Barcelona, if we're willing to go far enough back in time to look for a common cause.) The requirement that there be an indirect explanatory connection between belief and fact is, effectively, trivial.

Second, once we realize how common indirect connections are, when we try to find cases in which we have knowledge by indirect connections (rather than direct connections), we find that the only kinds of examples of knowledge by indirect connection are also examples of inferential knowledge. Knowledge of the future, in particular, is knowledge that we arrive at via inductive inference through our knowledge of the past and present.

Keeping these two points in mind, a general picture suggests itself: indirect connections can secure non-accidentality only when the knowledge in question is inferentially justified on the basis of premises which are non-inferentially known. But when the knowledge in question is non-inferential, the connection between belief and fact must be direct. This is precisely what EC says.

4.3 EC and *A Priori* Knowledge

Consider now the *a priori* knowledge objection. Facts that are known *a priori* don't seem to involve causal connections of any kind between fact and belief. As Benacerraf famously pointed out, mathematical facts are causally inert.

This is why EC does not refer to *causal* connections but rather to *explanatory* connections – *because* connections, if you're feeling cheeky. The notion of an explanatory connection is very closely related to Aristotle's notion of a cause. If we reject the notion of a “final cause” as being genuinely explanatory – as is common, post-Darwin – we can identify three different kinds of explanatory relations: formal explanation, material explanation, and causal explanation. Setiya (2012, Ch. 3) groups formal and material explanation together as a *constitution* relation, and I follow that convention. The statue exists *because* the lump exists, in the form of a statue. The window breaks *because* I threw the rock. Those are the two modes of explanatory connection – at least as the phrase “explanatory connection” occurs in the context of EC.

Note also that EC says that the fact that P must be responsible not for S's belief that P, but for S's non-factive mental states on which S bases her belief that P. This requirement that the causal chain pass through the agent's non-factive mental states is intuitive because it explains the role that our experiences play in our empirical knowledge. Consider my knowledge, by visual perception, that there is a cat on that mat. I know that there is a cat on the mat (in part) because I believe that there is a cat on the mat because there really is a cat on the mat. But note the path of the causal chain between fact and belief: The cat is on the mat; light reflects, strikes my retina; I now have a visual experience as of a cat

on a mat. And, on the basis of this visual experience, I come to believe that the cat is on the mat. The causal chain therefore passes through my experiences which serve as my evidence that the cat is on the mat, and this is an important feature of why my belief amounts to knowledge. If the fact that there is a cat on the mat caused my belief, but *not* by causing any of the experiences which are evidence for my belief, the causal chain would be “deviant” and thus my belief would not amount to knowledge.

Yet the added requirement that the causal chain pass through the agent's non-factive mental states is consistent with the idea that empirical knowledge requires a causal connection between fact and belief. EC requires both that P explain some of S's non-factive mental states, and that S's belief that P be based on those very mental states, where basing is a (species of) causal relation. Thus, if causation is a transitive relation (and I'm inclined to think it is), and if the fact that P causes S to have the mental states on which she bases her belief that P, it follows that the fact that P causes S to believe P. So the traditional idea of a causal requirement between belief and fact follows from EC, if we are thinking only of cases of empirical knowledge where the connection between belief and fact is purely causal.

But given EC's invocation of the more general notion of an explanatory (rather than, narrowly, causal) connection, and the requirement that the fact in question must explain an agent's mental states, a fact can satisfy EC in virtue of being part of a constitutive (rather than causal) explanation of EC's non-factive mental states. And what are the constituents of an agent's mental states? At least the following: *concepts* or *ideas* and *the relations between them*. This means that facts concerning the relations between ideas are (partly) constitutive of our non-factive mental states. Thus, EC does not rule out knowledge of facts concerning relations between our ideas. This means that EC does not rule out *a priori* conceptual knowledge. Therefore, EC also avoids the *a priori* knowledge objection.³⁹

So, again, we have an objection to the causal theory of knowledge, and two possible responses: the traditional response (which restricts the causal condition to *a posteriori* knowledge) and EC (which

³⁹ For an account of *a priori* knowledge that is compatible with EC, see Bealer (2000).

posits an explanatory condition on the bases of beliefs). There are two reasons to prefer EC to the traditional response. First, the traditional response implies that it is impossible to have a Gettiered belief in some *a priori* proposition. Yet this is surely possible – cases where a subject believes some logical or mathematical truth because they've taken a belief-pill or been struck on the head or are dupes of evil demons are common examples of accidentally-true *a priori* belief. EC explains why these beliefs count as accidentally true. Second, to account for this possibility, the traditional response will have to offer up a different kind of Gettier condition for *a priori* and *a posteriori* knowledge. EC promises to unify the *a priori* and *a posteriori* under one explanatory condition. Yet it still has the theoretical resources to distinguish the *a priori* and the *a posteriori*, since the former are known via constitutive explanatory connections, and the latter are known via causal explanatory conditions. Once again, EC is superior to the traditional response.

This sketch of an account of *a priori* knowledge suggests a potential way for non-naturalists to evade the Explanatory Challenge. If moral facts are conceptual truths, then they are the kinds of facts of which we can have *a priori* knowledge. And Terence Cuneo and Russ Shafer-Landau have argued that moral facts are non-natural, and that some of them – the “moral fixed points” – are conceptual truths.⁴⁰ So their view seems to have the resources to avoid the Explanatory Challenge. But I have two major worries about Cuneo and Shafer-Landau's position. First, the Open Question Argument seems to give us good reason to think that there are no moral conceptual truths. Someone who holds that killing is good might be *morally* confused, but they do not seem to be *conceptually* confused in the same way that someone who asserts that there are married bachelors is conceptually confused. Second, Cuneo and Shafer-Landau's view explicitly relies on the position that concepts are abstracta that reside in Frege's heaven; Cuneo and Shafer-Landau call this the “traditional view” of concepts. But this conception of concepts and “conceptual truths” is not one on which conceptual truths are constitutive of our mental

⁴⁰ Cuneo and Shafer-Landau (2014).

states, and thus not one on which knowledge of conceptual truths satisfies EC. The traditional view of concepts is a version of non-causal realism about concepts, and as such makes concepts the kinds of things that we can't know about.⁴¹

In sum: EC has the resources to answer both of the traditional objections to the causal theory of knowledge; it gives better answers to these traditional objections than the traditional responses; and EC maintains its skeptical upshot through its role in the Explanatory Challenge.

5. Conclusion

EC is the central epistemic premise of the Explanatory Challenge, which is an external skeptical challenge. My Evolutionary Undercutting Argument,⁴² in which EAD features as the central epistemic premise, is the internalized version of that same challenge. But the notion of *reliability* is no part of this challenge. So, with apologies to Enoch: it is this explanatory challenge (in either its internalized or externalized form) that is the strongest objection to normative non-naturalism.

Two tasks remain. First: The conclusion of the Explanatory Challenge is that non-inferential moral knowledge is impossible for agents like us. But we have not yet ruled out the possibility of inferential moral knowledge. Here, I think Hume's Is-Ought gap looms large: how can we come to have inferential moral knowledge if all of our basic knowledge is non-moral? This is a big question;⁴³ I save it for another day. Second: a complete vindication of EC will require us to first solve the Gettier Problem, thereby providing a satisfying account of knowledge, and show that EC follows from that account. That's too big a task for this paper, but the arguments in this paper suggest a way forward. As EAD suggests, a non-inferential belief is justified just in case it is the best explanation (by the subject's lights) of their evidence. And according to EC, a subject has non-inferential knowledge that P only if P

⁴¹ Copp (2018).

⁴² From Lutz (2017).

⁴³ See Zimmerman (2010, Ch. 5-6); Sinnott-Armstrong (2006, Ch. 7-8); Lutz (2015).

actually explains the subject's evidence. Conjoining these suggests that S has non-inferential knowledge that P if and only if P is both the best explanation of S's evidence and the actual explanation of S's evidence. I explore this account of knowledge in other work.⁴⁴

⁴⁴ Lutz (ms).

Bibliography

- Bader, Ralf (2017). The grounding argument against non-reductive moral realism. *Oxford Studies in Metaethics*, 12.
- Baras, D. (2017). our reliability is in principle explainable. *Episteme*, 14(2), 197-211.
- Bealer, George (2000). A Theory of the A Priori. *Pacific Philosophical Quarterly* 81 (1):1-30.
- Bedke, Matthew S. (2009). Intuitive non-naturalism meets cosmic coincidence. *Pacific Philosophical Quarterly* 90 (2):188-209.
- (2014). No coincidence. *Oxford studies in metaethics*, 9, 102-125.
- Behrends, J. (2013). Meta-normative realism, evolution, and our reasons to survive. *Pacific Philosophical Quarterly*, 94(4), 486-502.
- Benacerraf, Paul (1973). Mathematical truth. *Journal of Philosophy* 70 (19):661-679.
- BonJour, Laurence & Sosa, Ernest (2003). *Epistemic Justification: Internalism Vs. Externalism, Foundations Vs. Virtues*. Wiley-Blackwell.
- Clarke-Doane, Justin (2014). Moral Epistemology: The Mathematics Analogy. *Noûs* 48 (2):238-255.
- (2016). What is the Benacerraf Problem?. *New perspectives on the philosophy of Paul Benacerraf: Truth, objects, infinity*. Dordrecht: Springer.
- Comesaña, Juan (2005). Unsafe Knowledge. *Synthese* 146 (3):395-404.
- Copp, David (2018). Are There Substantive Moral Conceptual Truths? *Moral Skepticism: New Essays*, ed. Diego Machuca. Routledge.
- Cuneo, Terence and Russ Shafer-Landau (2014). The Moral Fixed Points: New Directions for Nonnaturalism. *Philosophical Studies* 171 (3): 399-443.
- Enoch, David (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- Faraci, David (forthcoming). Groundwork for an Explanationist Account of Epistemic Coincidence.

Philosophers' Imprint.

Field, Hartry (1989). *Realism, Mathematics & Modality*. Basil Blackwell.

Goldman, Alvin (1967). "A causal theory of knowing." *The Journal of Philosophy* 64 (12), pp. 357-372.

Harman, Gilbert, and Judith Jarvis Thomson. *Moral relativism and moral objectivity*. Oxford: Blackwell, 1996.

Huemer, M. (2005). *Ethical Intuitionism*. Palgrave Macmillan.

Joyce, Richard (2006). *The Evolution of Morality*. MIT Press.

Korman, D. and Locke, D. (forthcoming). Against Minimalist Responses to Moral Debunking Arguments. *Oxford Studies in Metaethics*.

Lehrer, Keith & Paxson Jr, Thomas (1969). Knowledge: Undefeated justified true belief. *Journal of Philosophy* 66 (8):225-237.

Linnebo, Ø. (2006). Epistemological challenges to mathematical platonism. *Philosophical Studies*, 129(3), 545-574.

Lutz, M. (2015). *The Case for Moral Skepticism*. PhD Dissertation, University of Southern California.

----- (2017). What Makes Evolution a Defeater? *Erkenntnis*, 1-22.

----- (ms). The Matched Explanations Theory of Knowledge.

Lutz, M. and Ross, J. (2017). Moral scepticism. *Routledge Handbook of Metaethics*, ed. McPherson and Plunkett. Routledge.

Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Penguin.

McCain, K. (2014). *Evidentialism and epistemic justification* (Vol. 59). Routledge.

McGrath, Sarah (2014). Relax? Don't do it! Why moral realism won't come cheap. In *Oxford Studies in Metaethics* 9 (2014): 186-214.

McPherson, Tristram (2012). "Ethical non-naturalism and the metaphysics of supervenience." Oxford

studies in metaethics 7 : 205-234.

Olson, Jonas (2014). *Moral Error Theory: History, Critique, Defence*. Oup Oxford.

Schechter, J. (2013). Could Evolution Explain Our Reliability about Logic?. *Oxford studies in epistemology*, 4, 214-239.

----- (2017). Explanatory Challenges in Metaethics. *Routledge Handbook of Metaethics*, ed. McPherson and Plunkett. Routledge.

Setiya, Kieran (2012). *Knowing Right From Wrong*. Oxford University Press.

Sinnott-Armstrong, Walter (2006). *Moral Skepticisms*. Oxford University Press.

Skarsaune, Knut Olav (2011). Darwin and moral realism: Survival of the fittest. *Philosophical Studies* 152 (2):229-243.

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

----- (2016). Objectivity and truth: you'd better rethink it! *Oxford Studies in Metaethics* 11.

Vavova, Katia (2014). Debunking Evolutionary Debunking. *Oxford Studies in Metaethics* 9:76-101.

White, Roger (2010). You just believe that because.... *Philosophical Perspectives* 24 (1):573-615.

Wielenberg, Erik J. (2010). On the evolutionary debunking of morality. *Ethics* 120 (3):441-464.

----- (2014). *Robust ethics: The metaphysics and epistemology of godless normative realism*. OUP Oxford.

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

Woods, Jack (2016). *Mathematics, Morality, and Self-Effacement*. *Noûs*.

Yamada, Masahiro (2011). Getting It Right By Accident. *Philosophy and Phenomenological Research* 83 (1):72-105.

Zimmerman, Aaron (2010). *Moral Epistemology*. Routledge.