

NON-NATURALISM AND THE ‘THIRD FACTOR’ GAMBIT

Aaron Elliott and David Faraci

Abstract

Normative realists face a fundamental epistemological challenge to show how we can have epistemic access to realistically construed normative truth. Different forms of realism have different responses to this challenge. But the general consensus seems to be that for non-naturalists, there is only one option: so-called third-factor explanations. Indeed, the current dialectic over the non-naturalist’s (in)ability to meet this challenge has more or less reduced to the question of whether or when third-factor explanations can answer it. In this paper, we argue that once the natures of the relevant challenge and of third-factor explanations are better understood, it becomes clear that the latter offers the non-naturalist no hope in meeting the former. If non-naturalists really are limited to third-factor explanations, their view is inconsistent with the existence of normative knowledge.

Normative realists face a fundamental epistemological challenge to provide “an account of how it is that we can have epistemic access to the normative truths about which they are realists” (Enoch 2011, 151). This is sometimes called the reliability challenge, or the Benacerraf-Field Challenge (in accordance with its roots in the philosophy of mathematics), and many have taken it to be the core of genealogical challenges such as Street’s (2006) ‘evolutionary debunking argment’. We just call it the Challenge. The basic thought is that normative judgements lack the right sort of connection to realistically construed normative truth for those judgements to track that truth.

Different forms of realism have different responses to the Challenge. But the general consensus is that for non-naturalists, there is only one option: so-called third-factor explanations. Indeed, the current dialectic over the non-naturalist’s (in)ability to meet genealogical challenges has more or less reduced to the question of whether or when third-factor explanations can answer the Challenge. In this paper, we argue that once the nature of the Challenge and of third-factor explanations are better understood, it becomes clear that the latter offers the non-naturalist no hope in meeting the former. If non-naturalists really are limited to third-factor explanations, their view is inconsistent with the existence of normative knowledge.

1. The Challenge

David Enoch's (2011) framing of the Challenge, quoted above, and his third-factor proposal, are the most developed and widely discussed.¹ In Enoch's view, the best way to cash out the Challenge is as a demand to *explain correlations*:

[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 judgment *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 [non-naturalist] view of normativity, it can't be that our normative judgments are causally or constitutively responsible for the normative truths, because the normative truths are supposed to be independent of our normative judgments. And given that (at least basic) normative truths are causally inert, they are not causally responsible for our normative beliefs. Nor does there seem to be some third-factor explanation available to the robust realist. And so the robust realist is committed to an unexplained striking correlation, and this may just be too much to believe. (Enoch 2011, 159)

Understanding the Challenge this way, Enoch says, is useful because it involves “no hidden assumption about the nature of knowledge, or of epistemic justification, or anything of the sort. There is just a striking correlation, the need to explain it, and the apparent unavailability of any explanation to the challenged view” (Enoch 2011, 159).

We are sympathetic with Enoch's framing of the Challenge in terms of explaining correlations, but think it is a mistake not to relate this back to some epistemological notion, such as knowledge. Indeed, without such a connection, it is hard to see how this is an *epistemological* challenge at all. Here is Enoch himself making the same point about *access*:

In the following subsections, I discuss versions of the epistemological worry put in terms of justification, reliability, and knowledge. It is possible, of course, that my arguments there fail. But if they do not, what remaining epistemological worry could talk of epistemic access introduce? If in the next subsections I manage to convince you that there are no special problems with the justification of normative beliefs, with the reliability of normative beliefs, or with normative knowledge, it seems to me you should be epistemologically satisfied. I do not see how talk of epistemic access should make you worried again. (Enoch 2011, 152–53)

¹ For similar proposals see, e.g., Skarsaune (2011) and Wielenberg (2010).

That sounds right. But if this line of reasoning is right about access, it's hard to see how it would be any less right about *explaining correlations*. There can still *be* a challenge to explain correlations (just as there can still be a challenge to show that we have 'access' in some sense), but it is not an *epistemological* challenge unless an explicable correlation is a necessary condition on justification, reliability, or knowledge. Otherwise, by Enoch's own lights, it poses no threat to the epistemological status of any of our normative beliefs.

In contrast to Enoch, we maintain that the Challenge is first and foremost a challenge to *knowledge*. It is not a challenge to the *truth* of our beliefs. Indeed, we follow the many who argue that in meeting the Challenge it is acceptable to assume the truth of our beliefs.² Enoch's framing in terms of explaining correlations makes this obvious: it is impossible to explain the correlation between A and B without assuming that both A and B obtain.³ We also take the Challenge to be about justification only derivatively: good evidence that our beliefs don't constitute knowledge is quite arguably a defeater for those beliefs, even if those beliefs would otherwise be *prima facie* justified.

Enoch initially shies away from characterizing the Challenge in terms of reliability. He takes a set of beliefs to be reliable "if and only if a sufficiently large portion of it is true" (Enoch 2011, 155). Given this understanding of reliability, as long as there is normative truth, normative beliefs can be reliable, so the Challenge can't be to the possibility that our beliefs are reliable without evidence that there is no normative truth. Nevertheless, it should be clear that there is little difference between "a sufficiently large portion" of a set of beliefs' being true and there being a correlation between those beliefs and the truth. And, indeed, Enoch later claims that the Challenge *is* "primarily about reliability"

² See, e.g., Berker (2015), Clarke-Doane (2016) and Vavova (2014).

³ This is not to say that there are no limits on *how* the assumption of truth can be used in this context. See [redacted] for discussion.

(Enoch 2011, 160). It's just that the demand isn't to directly demonstrate the possibility of reliability, but rather to *explain* how reliability is possible *given realism*, on the grounds that inexplicable reliability is highly improbable.

We agree with all of this, with the exception of the final clause. *If* we understand reliability in terms of a correlation between beliefs and truth at the actual world,⁴ then the Challenge is to explain reliability. But the reason such an explanation is required, we maintain, is not merely that brute correlations are improbable. Rather, it is because inexplicable correlations are *accidental*, and accidental correlations between beliefs and truth never constitute knowledge. Such beliefs are Gettiered.

Enoch anticipates a Gettier version of the Epistemological Challenge, but thinks that it can be easily dismissed, or ignored:

[L]et me assume (for now) that knowledge is justified true belief of a special kind (the qualification needed in order to deal with Gettier cases). If this is so, all that the realist has to do in order to accommodate normative knowledge is to account for normative truths, normative beliefs, justification of normative beliefs, and that extra bit needed to deal with Gettier. Very well, then, let's see: for the reason mentioned in the previous subsection, in the context of an attempt to articulate an independent challenge to Robust Realism, we are to assume that there are normative truths, and furthermore that at least some normative judgments express beliefs . . . All that remains, then, is to accommodate with regard to normative beliefs the necessary anti-Gettier clause, whatever it is. But then surely the realist is entitled at least to a wait-and-see attitude. "Given that I've given you already so much, what reason do we have to believe that I won't be able to give you that too? When you epistemologists have settled on an anti-Gettier clause, let me know, and we'll take it from there." (Enoch 2011, 156)

Despite what this suggests, however, raising a Gettier challenge does not require settling on an anti-Gettier clause—i.e., does not require knowing what is sufficient for avoiding the Gettier problem. It only requires knowing one thing that is *necessary for avoiding* the Gettier problem and then

⁴ This last bit about the actual world is crucial. Some take the reliability that needs explaining to be a *modal* correlation between beliefs and truth across some suitable set of possible worlds, *given* such a correlation at the actual world. In part because of such lack of agreement on the term, we largely avoid talk of reliability in the remainder of the paper.

arguing that your target fails to meet that minimally necessary condition for knowledge. As suggested above, one common way of understanding what makes a belief Gettiered is in terms of some kind of *accidentality*.⁵ So long as we have some minimal condition for *non-accidentality* on hand, we could frame the Challenge as the worry that, on non-naturalism, all normative beliefs fail to meet that condition and are therefore Gettiered at best.⁶ We maintain that Enoch's version of the Challenge lends itself to *precisely this framing*. If there is no explanation for a correlation, it is natural to say that that correlation is a mere accident or coincidence. Thus, the Challenge is to explain correlations, and it is an *epistemological* challenge because if there is no such explanation, normative beliefs are Gettiered at best, and therefore never constitute knowledge.

One point of clarification: Not just *any* explanation is sufficient for non-accidentality. To see this, consider that a correlation can always be explained by citing the conjunction of the origins of its correlates. This is clearly not the sort of explanation we are looking for, especially given that such explanations are available for *any* true belief. If they could be used to meet the Challenge, we could simply argue as follows:

P1 Our ethical beliefs are by-and-large true.

P2 If our beliefs are by-and-large true, this can be explained by the conjunction of the origins of those beliefs and the origins of those truths (or the truths themselves, if they are fundamental).

C Therefore, there is an explanation for the correlation between our beliefs and the truth in ethics.

Our rejection of this argument seems to stem from the *merely conjunctive* nature of the explanation offered in P2. The conjunction of the origins of our beliefs and the origins of the truth

⁵ E.g., not long after the original Gettier (1963) paper, Peter Unger argues that “[f]or any sentential value of *p*, a man's belief that *p* is an instance of knowledge only if it is not an accident that the man's belief is true” (Unger 1967, 172). For a more recent proposal of a non-accidentality condition and further references, see Jenkins (2006).

⁶ At best since they still might be false, or there might be independent reasons to think them justified.

isn't the sort of explanation we're after precisely because it fails to illuminate a *connection* (as opposed to a mere correlation) between our beliefs and the truth.

This language is still largely suggestive; we haven't said what it is for there to be such a connection, except that it requires more than mere correlation.⁷ But for our purposes we only need a minimal non-accidentality condition that rules out trivial explanations like those just discussed:

Connection Beliefs are non-accidentally true only if there is an explanation for their correlation with the truth at the actual world beyond a trivial explanation such as that provided by the mere conjunction of the origin of those beliefs and the (origins of the) truth.

In what follows, we will say that there is a CONNECTION between beliefs and truth insofar as there is a non-trivial explanation for their correlation at the actual world. The absence of a CONNECTION is one way in which true, justified beliefs can fail to be knowledge. Our view is that this is the best way to understand the Challenge: if a CONNECTION between ethical beliefs and ethical truth is in principle impossible, then even if our ethical beliefs are true, they are so only accidentally, and are therefore Gettiered at best.

2. Enoch's Third-Factor Proposal

As seen in the quotation at the start of the previous section, Enoch takes there to be three possible forms of explanation for correlations generally: for any correlated factors A and B, either (i) A is (part of) the ultimate explanation for B; (ii) B is (part of) the ultimate explanation for A; or (iii) some third

⁷ We are far from the first to use such language; many have framed the Challenge in terms of finding a 'connection' or 'link' between our beliefs and the truth, including Benacerraf himself: "[I]f I know that Cleveland is between New York and Chicago, it is because there exists a certain relation between the truth conditions for that statement and my present 'subjective' state of belief (whatever may be our accounts of truth and knowledge, they must **connect** with each other in this way). Similarly, in mathematics, it must be possible to **link up** what it is for p to be true with my belief that p " (Benacerraf 1973, 667, bolding ours).

factor C is (part of) the ultimate explanation for both A and B (Enoch 2011, 167). Thus, the non-naturalist must accept that: (1) normative beliefs explain normative facts; (2) normative facts explain normative beliefs; or (3) some third factor explains both.

The fact that non-naturalism is a form of realism straightforwardly rules out (1).⁸ And few non-naturalists will want to claim that non-natural normative properties directly explain our normative beliefs (as per (2)), especially since, as Enoch notes, most hold normative properties to be causally inert).⁹ In what remains, we stipulatively treat this as a commitment of non-naturalism: the normative never directly explains anything non-normative, including our beliefs about it. This all suggests that if the non-naturalist is to identify a CONNECTION between normative beliefs and normative truth, he will have to appeal to a third-factor explanation.

Enoch takes the apparent force of the Challenge, as a threat to non-naturalism, to stem from a failure to acknowledge the possibility of third-factor explanations. True, non-naturalism seems to rule out belief-to-truth and truth-to-belief explanations. And true, there are empirically respectable explanations for our normative beliefs. But there are also third-factor explanations that are both non-naturalist-friendly and compatible with our beliefs' naturalistic genealogy. Indeed, Enoch thinks that the solution to the Challenge begins with the very premise debunkers take to undermine realism: that our normative beliefs are (largely) explained by selective (evolutionary) pressures. Once this is on the table, he insists, evolved normative truth-tracking¹⁰ becomes significantly less mysterious:

⁸ There are possible exceptions, such as if my belief that X is important partly grounds the fact that it is good to help me realize X. But that would have to proceed via further normative facts (e.g., that it's good to help people act in accordance with their evaluative standards) that we would further need to explain our access to.

⁹ Though see, e.g., Oddie (2009) for a view on which such properties are causally efficacious. There is also the possibility that non-natural properties *non-causally* explain our beliefs in them. See, e.g., Bengson (2015). We address this possibility (indirectly) in §4.

¹⁰ Or, at least, "a close relative of truth-tracking," if one takes tracking to entail a causal connection. See Enoch (2011, 166 n.36).

[S]urvival (or whatever) is good; so behaving in ways that promote it is (pro tanto) good; but one efficient way of pushing us in the direction of acting in those ways is by pushing us to believe that it is good to act in those ways. And in fact, as we have just seen, it is good so to act. So the normative beliefs this mechanism pushes us to have will tend to be true. (Enoch 2011, 169)

If we grant all this (and we should all acknowledge that that's a very big 'if'), it seems that Enoch has explained the correlation between normative beliefs and facts by offering a third-factor explanation of the correlation between normative beliefs and normative facts. That third factor is the goodness of survival.

Whatever one's doubts about Enoch's specific proposal, it is highly plausible that third-factor explanations play a role in epistemology. Consider knowledge of the future. Meteorologists can know certain things about tomorrow's weather. Some will be comfortable saying that relevant future facts explain meteorologists' beliefs.¹¹ But the many who deny that future events can serve as explanantia are not lost. For it is natural to explain the correlation between the meteorologist's belief and the truth by citing a *common cause*: current weather conditions and certain facts about how weather develops. The current weather and those facts about weather progression explain (via his awareness of them) the meteorologist's beliefs about tomorrow's weather. And the current weather and those facts about weather progression likewise explain tomorrow's weather itself.¹²

Of course, Enoch's explanation is non-causal, since for the non-naturalist the goodness of survival does not *cause* anything, so we need to be willing to expand the class of legitimate third-factor accounts from 'common causes' to 'common explanantia'. But this does not seem particularly troubling.

¹¹ See, e.g., Jenkins (2006).

¹² The idea of common-cause explanations can be traced back to a proposal from Alvin Goldman (e.g., 1967), which is that all cases where beliefs and truth have a common cause.

But there is something else troubling here. We stated earlier that we are taking it as a commitment of non-naturalism—one Enoch explicitly endorses (at least with respect to *causal* explanation)—that the normative never explains the non-normative. Yet Enoch’s proposed explanans, the goodness of survival, is a normative fact. If his proposal is meant to be one of common explanation, it seems he would be committed to rejecting it!

Helpfully, the problem is not with Enoch’s proposal, but with his taking it to involve common explanation, for not all third-factor explanations involve common explanations. To see this, consider Enoch’s own example of how third-factor explanations function in a non-normative context:

What explains the correlation (if indeed there is one) between giving rise to strongly affectionate feelings and having a poor sense of time? Here’s one possible answer: Young children are cute. Being a young child explains—indeed, perhaps causes—having a poor sense of time. And of course, being cute is closely though perhaps not causally related to giving rise to strongly affectionate feelings. The fact that young children are cute, then, pre-establishes the harmony between giving rise to strongly affectionate feelings and having a poor sense of time. (Enoch 2011, 169)

The correlated factors are (A) giving rise to strongly affectionate feelings and (B) having a poor sense of time. What explains this correlation, according to Enoch, is (C) the fact that young children are cute. Yet, crucially, the fact that young children are cute is itself responsible *for neither A nor B*.

If C is not responsible for A or B, how can it explain their being correlated? The answer is not hard to find. Cute things tend to garner affection. Young children tend to have poor senses of time. And in cases exemplifying the correlation between A and B, certain things’ being cute explains their garnering affection and those same things’ being young children explains their having a poor sense of time. So there is a correlation between A and B not because something is responsible for both, but because what (in some cases) explains some particular thing’s instantiating A is correlated with what (in some of those same cases) explains that same thing’s instantiating B. This further correlation—between being cute and being a young child—has simply been expressed, by Enoch, as the single claim (C) that children are cute.

Here is what we have discovered: The obvious way for a correlation to be explained is by one correlate's being responsible for the other. A less obvious way, as Enoch tells us, is for some third factor to be responsible for each correlate. From here, we refer to these as *common explanans* (CE) explanations. But yet *another* possibility is that the correlation to be explained, between A and B, is explained by a *further* correlation whose correlates are *respectively* responsible for A and B. It is an explanation of this last kind that is at work in Enoch's example above: the correlation between being cute and being a young child explains the correlation between giving rise to strongly affectionate feelings and having a poor sense of time. We refer to these as *higher-order correlation* (HOC) explanations.

Of course, some explanations that appear to be HOC will turn out ultimately to be CE explanations, for there will be some common explanans further up the explanatory chain. Indeed, the case just discussed might be of this sort; there may well be a single explanation for the correlation between being cute and being a young child. In what follows, when we refer to HOC explanations, we always mean *pure* HOC explanations, ones that are HOC "all the way up."

Since the goodness of survival cannot directly explain our beliefs, Enoch's proposal appears to be HOC, not CE. This makes sense: for most non-naturalists, the fact that survival is good is a fact about a non-explanatory relation—supervenience—between two distinct properties: survival and goodness. What's more, this would be true of *any* normative third factor: it must be a correlation between distinct properties.¹³

¹³ In the non-normative case above it was fairly clear what was correlated: the two properties of 'garnering affection' and 'having a poor sense of time'. Here, it is less clear. Is it the case that, as just suggested, the correlation is between the properties of 'promoting survival' and 'being good'? We are genuinely uncertain. But it certainly seems as though *some* correlation is being expressed. The only alternative (on the plausible assumption that relations between distinct properties are always correlations) is that "survival is good" expresses a property identity or constitutive relationship of some kind. But this is clearly not the case, especially given that, as Enoch himself explicitly tell us, survival is only *sometimes* good. And, of course, for a non-naturalist like Enoch, the relevant properties have to be distinct anyway.

The question, then, is whether such explanations merely *appear* to be HOC, or whether there might be some single explanans further up the chain. Given how non-naturalism is typically understood, it seems these must indeed be pure HOC explanations. This is because non-naturalism is often distinguished from other metanormative contenders by virtue of its acceptance of:

Discontinuity The non-normative plays no role in the explanation of the fundamental normative facts.

Given such an understanding of non-naturalism, no non-normative third factor could explain the normative truths our beliefs concern (at least assuming some of those beliefs concern fundamental normative facts). From here, we treat Discontinuity as a defining commitment of non-naturalism. This is stipulative, but not purely so, as the vast majority of those who self-identify as ‘non-naturalists’ accept Discontinuity or something with relevantly similar implications where our arguments are concerned.¹⁴ It follows that non-naturalists are limited to HOC explanations, with a potential exception to be discussed in §4. In §3, we argue that HOC explanations cannot be used to meet the Challenge, for they are *trivial* explanations of the sort introduced at the end of §1, and as such do not constitute CONNECTIONS between normative beliefs and normative truths.

¹⁴ We borrow the term ‘discontinuity’ from McPherson (2012). In addition to Enoch, McPherson cites Dancy (2006), Fitzpatrick (2008), Huemer (2005), Shafer-Landau (2003) and Wedgwood (2007) as examples of non-naturalists thus understood. (It should be noted, though, that Shafer-Landau’s inclusion is debatable, given some of what he says about the relation between the normative and the non-normative.)

It is imperative that we keep in mind here that Discontinuity is about explanations of *fundamental* normative facts. Nothing blocks non-naturalists from holding that particular things’ non-normative properties partially explain their normative properties. But for the non-naturalist, such cases must involve some further fundamental normative fact that is part of the ultimate explanation in these cases. For example, if Bob is bad because Bob is a liar, it seems Bob’s being a liar explains (in the immediate sense) his being bad. But this is true, for the non-naturalist, only because (say) it is an independent normative fact that being a liar makes one bad. Ultimately, Bob’s badness depends not just on his being a liar, but also on that normative fact.

3. The Epistemological Insignificance of HOC Explanations

Consider a classic case: someone comes to believe that it is 3:00, at 3:00, because her clock reads '3:00'. In epistemically good versions, there is some CONNECTION between her belief and the truth, say that its being 3:00 explains her belief via explaining why her clock reads '3:00'. By contrast, in the stopped-clock Gettier case, her belief is true because the clock has read '3:00' since it stopped last night; and it happens to be 3:00. This is a trivial explanation.

It should be clear that HOC explanations take precisely this form. Indeed, if we allow for one-off correlations, the above *just is* an HOC explanation: the correlation between when she looks at the clock and (whatever explains) its being 3:00 explains the correlation between her belief and the truth. Even if we take correlations to require multiple iterations, the point stands. Imagine a massive version of the stopped-clock case: a world in which there is a purely coincidental correlation between when clocks stop and when people next look at them. This admits of essentially the same HOC explanation: there is a correlation between what stopped clocks say (the explanation for people's beliefs) and (whatever explains) the time. Yet this is no less a Gettier case.

Of course, as the non-naturalist will be quick to point out, Enoch's proposal adds something to the HOC explanation above: the necessity of the higher-order correlation. But it is unclear why this should make a difference. Trivial explanations for a belief's being true are intuitively no better for being necessary. Suppose Donald believes that $2+2=4$. And even suppose that it is necessary that he believes that $2+2=4$.¹⁵ This does nothing to remove the sense that Donald's being right might be

¹⁵ Fleshing this out a bit: Suppose we accept an essentialist view of identity (à la Kripke) according to which in any world where Donald exists, he is the product of a particular sperm and egg. And suppose it is partly constitutive of that sperm and egg pair that they give rise to a person who believes that $2+2=4$, not as the result of some rational line of reasoning, but simply as a matter of brute fact. In that case, Donald believes that $2+2=4$ at every possible world where Donald exists. Thanks to [redacted] for this framing.

coincidental in the Gettier-relevant sense, that he just *happens* to believe the truth (despite his doing so at every possible world).¹⁶

Now scale up: Donald necessarily believes a whole bunch of necessary truths. Of course, it might seem rather implausible that such a correlation would persist without a vindicating explanation. And so, unlike in the one-off case, this correlation might seem to indicate that there *is* some further non-HOC explanation—that, say, Donald believes these things *because* they are true. But if we *stipulate* that this is not the case, that the necessity of both the beliefs and the truth is the end of the story, this explanation seems no less trivial, no less irrelevant to the Challenge, than the one-off case above.¹⁷ The relevant point is that for the non-naturalist committed to Discontinuity, there *cannot* be a further explanation. The buck stops with a pure HOC explanation: there is some brute, fundamental correlation between the normative and the non-normative, such as the goodness of survival. And any beliefs whose correlation with the truth admits of only an HOC explanation are Gettiered at best.

4. Normative Grounding: A Non-Naturalist CE Account?

In this section, we consider the only possible avenue for offering a non-naturalist-friendly CE account: taking the fundamental normative facts to be *meta-explanatory*. We focus on Ralf Bader's (forthcoming) proposal. Bader proposes that non-naturalists hypothesize a distinctive *normative* form of grounding that runs from non-normative properties to the normative properties they give rise to. Thus, for Bader, the claim "survival is good" is a *normative law* dictating a *normative grounding* relation between survival

¹⁶ This is to implicitly reject a purely modal reading of "coincidence," such as that defended by Clarke-Doane (2016) and assumed by many others. For arguments see [redacted]. See also Bengson (2015).

¹⁷ And of course, metaphysical possibility includes all highly improbable but still possible states of affairs somewhere.

and goodness.¹⁸ Adapting this for Enoch's purposes, survival might serve as a third factor that can explain both our normative beliefs (via evolutionary forces) and the normative facts (via normative grounding). This is consistent with Discontinuity only because the latter explanatory relation is meta-explained by a more fundamental normative fact: the law in question.

Unfortunately, this proposal faces a different (though intimately related) problem to its HOC cousin: in order to meet the Challenge, it is not enough to cite a possible CE explanation; that explanation must be *available* to us (in a sense to be clarified below). But such availability is impossible given non-naturalism.¹⁹

Begin with a standard Gettier case: Abe has the justified true belief *there is a sheep in the field* because he saw a decoy sheep in the field, which happened to be obscuring a real sheep. The fact that Abe holds this belief is coincidental to the fact it concerns. This admits of an HOC explanation: the decoy sheep explains Abe's belief; (say) the presence of delicious grass in the field explains the sheep's

¹⁸ The details here could vary, though this won't matter much for our purposes. This could be a fact about states of affairs (e.g. the extent to which they involve survival as compared to some previous state of affairs *normatively grounds* the extent to which they are good), the survival-promoting features of particulars (e.g. the extent to which X promotes survival *normatively grounds* the extent to which X is good), or something else entirely. Since we will be developing Bader's account as a variation on Enoch's, we will assume something like the second option. We employ this option, since, of the possibilities, this one can perhaps most obviously play the relevant role in an Enoch-style evolutionary story.

¹⁹ In addition to the problem we will discuss below (which is perfectly general), there is a concern that the actual details of this *particular* proposal fail to provide a genuine CE account. The facts that explain our normative beliefs will turn out to be different from the facts that explain the normative facts, even if both are facts about survival. The evolutionary story holds that the survival of our ancestors explains (the relevant set of) our normative beliefs. So, if there is to be a common explanans for the normative fact, the normative fact must be explained by those same facts about the survival of our ancestors. Unfortunately, the only normative facts that the survival of our ancestors could plausibly explain are the normative facts about the survival of our ancestors. But these normative facts won't be the ones that our normative beliefs in question correspond to. Thus, no CE account is available on this version of the story.

presence; and the decoy sheep and the delicious grass are correlated (they're both in the field). But we can easily adjust the case to include a common explanans: the real sheep is there not because of the delicious grass, but because it was attracted to the field by the decoy sheep. In that case, the decoy sheep²⁰ is a common explanans for Abe's belief and the truth. Yet, crucially, it seems no less a Gettier case.

Of course, there is a version of this case where Abe is *not* Gettiered. That's the version where Abe *infers* that there is a sheep in the field, *via* the common explanans. In such a case, Abe reasons along the following lines: "either that is a sheep or it is a decoy sheep; if it is a decoy sheep it will have attracted a real sheep to the field; so, there is a sheep in the field."²¹

For this to be a case of knowledge, however, Abe can't merely *believe* that decoy sheep attract real sheep; that belief needs to have certain epistemic credentials.²² Suppose Abe came to believe that decoy sheep attract real sheep because he read an article to that effect in the inaugural April Fools' issue of *Nature*, before the editors announced the trick.²³ If the story is false, Abe's conclusion is obviously Gettiered. But even if the story is true, Abe's belief is true by coincidence. This leaves Abe Gettiered about both the flocking habits of sheep, and any inferences that rely on this belief.²⁴

²⁰ Or the fact that it looks like there is a sheep in the field. Or the fact that there is a convincing decoy sheep in the field. Take your pick.

²¹ Obviously, Abe would also have to know that there are enough sheep around here that the decoy sheep is going to attract at least one.

²² This is a claim about what's necessary, not what's sufficient. We aren't claiming that Abe couldn't be Gettiered in some other way.

²³ The article explains that sheep evolved to travel in flocks, so if a sheep sees a decoy sheep, the real sheep will attempt to start a flock with it.

²⁴ For those who are skeptical of this claim, compare to a modification of another classic Gettier case. Smith sees Jones driving a Ford, and so justifiably believes that Jones owns a Ford. Smith then makes a valid logical inference to "either Jones owns a Ford, or Brown is in Barcelona." As it turns out, Jones was borrowing (or

Similar points go for the meteorological case we first used to motivate the relevance of third-factor explanations. Recall how we first described that case: “The current weather and those facts about weather progression explain (*via his awareness of them*) the meteorologist’s beliefs about tomorrow’s weather.” The meteorologist has the right set of beliefs, all in good standing: a belief about the current weather, and a belief about the explanatory relevance of the current weather. If the meteorologist lacked one of these beliefs, or if they weren’t all in good standing, then we would react differently to the case.²⁵

Here’s a first pass at the lesson of these cases: CE explanations only alleviate concerns about accidentality when the relevant beliefs are formed by *inference* from *knowledge* of that CE explanation.²⁶

This is likely to strike some as too strong, in two ways. First, it may seem strained to take all relevant cases to involve *inference*. Second, it may seem too strong to require that Abe *know* that decoy sheep attract real sheep. Consider a variation on the weather prediction case. Suppose Jim smells the air and comes to believe that it is going to rain. It seems quite possible that Jim *knows* that it is going to rain. Throughout his life, he has had more than enough experience to indicate a strong correlation between a certain smell in the air and impending rain. We may suppose there is a third-factor

‘borrowing’) the Ford, but Brown *is* in Barcelona. Smith has a justified true belief, but only by coincidence, and so it is a Gettiered belief. Now, suppose that on the basis of this Gettiered belief Smith comes to believe that he has one true disjunctive belief (let him have no other disjunctive beliefs for simplicity). This new belief is justified, since it was formed via a valid abstraction inference from a justified belief. Furthermore, the belief is true. But, finally, this belief is also Gettiered: because the first belief was coincidentally true, the belief about having the true belief is coincidentally true. The Gettier status is transmitted through inference.

²⁵ Suppose that, unbeknownst to our meteorologist, two bolts of lightning strike his weather station; the first erases all of the data, and the second just happens to duplicate that data. The meteorologist collects this data, and uses it to form certain true beliefs about tomorrow’s weather. This is a Gettier case, not because there is no common explanans for his belief and the truth, but because the meteorologist is Gettiered about the data.

²⁶ The point that third-factor explanations matter only for inferential knowledge is due to [redacted].

explanation at work here, that the phenomena Jim is responding to are caused by the same current conditions that will shortly cause it to rain.

Is it plausible to say that Jim *infers* that it will rain? We think so, at least so long as we take a broad enough view of inference. But the terminological point is orthogonal here. All that matters for our purposes is that Jim's belief that it will rain is *parasitic* on his experience of certain current features of the air. That is guaranteed by the fact that the case involves third-factor explanation: the causes of Jim's belief (the current weather conditions) are something other than what those beliefs are about (the rain).²⁷ Jim's belief-forming process must therefore be indirect; there is a *transition* between input and output. We will continue to refer to such transitions as inferences, but again nothing hangs on this terminological choice.

On to the second worry. Air that smells *like this* is a reliable indicator of impending rain. Call this the Rain Fact. Must Jim *know* the Rain Fact in order to infer that it's about to rain? The most obvious objection here is that Jim needn't know the Rain Fact because Jim needn't *believe* the Rain Fact. Jim may have no beliefs whatsoever about air quality and rain; the inferential process may be entirely automatic.²⁸ Given a broad enough account of *implicit* belief, it may be that this objection can be dealt with. But our response does not hinge on this. What matters for our purposes is that, as our discussion of Abe above shows, it is not enough that Jim's system takes him from air quality to future rain. Rather, mirroring the Challenge, there must be a non-trivial explanation for the correlation between Jim's inferential disposition (his sub-personal 'belief' in the Rain Fact) and the Rain Fact itself.

²⁷ Things are somewhat more complicated than this. First, there will be cases where inference is required despite the input and output's being the same property—e.g., I infer that I'm drinking H₂O from a direct experience of water. This is a result of the hyperintensionality of doxastic states. Second, there may be cases where the inference is merely conceptual. I have experiences of certain colors and shapes and I see a bus.

²⁸ This objection is pressed by Werner (forthcoming) in response to similar arguments from Faraci (2015), concerning moral perception. [redacted]

We can show that this is the case by *reductio*. Suppose (like Jim's) Kira's system is disposed to infer that it will rain from certain features of the air, but that the explanation for this disposition has nothing to do with the Rain Fact. It is a mere coincidence that her system is accurate in this way. Suppose (for *reductio*) this does not prevent her from gaining knowledge; all that matters is that her system is accurate. Now compare Kira to Jen, who is exactly like Kira, except that he explicitly *believes* the Rain Fact for the very same (disCONNECTED) reason that Kira's system infers in accordance with the Rain Fact. Since Jen's belief in the Rain Fact is not knowledge, his inferences from it are not knowledge, and therefore Jen does not know that it is going to rain. Thus, according to this view, two agents can differ in their knowledge merely because one of them *explicitly* makes an inference that the other makes automatically. This is implausible. We should therefore conclude that the need for a non-trivial explanation for the correlation between belief and truth extends to the correlation between whatever mental states ground sub-personal inferences and truth.

Putting this all together, we can offer the following necessary condition on CE explanations' meeting the Challenge. Call something an *inference-grounding state* if it is either a belief used in explicit inference, or a mental state that grounds sub-personal inferences. Then:

Further Connection Take any CE explanation EXP—e.g., the claim that E explains P-beliefs and E explains P. EXP addresses the Challenge only if there is a further, possible CONNECTION between agents' E-to-P-inference-grounding states and the fact that E explains P.

Further Connection gives us a better understanding of the good CE cases discussed above. Our driving case is of a meteorologist's belief about tomorrow's weather on the basis of today's weather. She has beliefs in the explanatory relevance of current weather conditions (E) to future weather conditions (P), which serve as the bridging premise from E to P. Presumably abduction on previous observations provides a CONNECTION between her belief in E's explanatory relevance for P

and the fact that E is explanatorily relevant for P.²⁹ Because of this, she meets Further Connection. A similar story holds for Jim, but in his case the inference (explicit or not) is plausibly based on induction from observed correlations. He thereby has a CONNECTION to the fact that E is explanatorily relevant for P. Further Connection tells us that if the Bader-inspired CE explanation—survival both grounds goodness and causes certain normative beliefs—can be used to meet the Challenge, there must be a CONNECTION between the fact that survival grounds goodness and the states that ground our inferences from survival to goodness.

This is untenable for the non-naturalist. Consider Betty, who infers goodness from survival (or something survival-promoting). The non-naturalist needs to explain the CONNECTION between Betty's inference-grounding state—her belief that survival implies goodness or disposition to make inferences in line with this claim—and the relevant normative law. This requires that the law explain her belief, the reverse, or that some third factor explain both.

None of these options are open to the non-naturalist. First, the non-naturalist can't hold that Betty's inference-grounding state explains the law any more than that our beliefs directly explain the normative facts. Second, the non-naturalist can't hold that the law explains Betty's inference-grounding state any more than that goodness itself explains our normative beliefs. The non-naturalist can no more hold that the normative law directly explains our belief in it (or our dispositions to infer in accordance with it) than that any other normative fact explains any mental state. Could the law instead *indirectly* explain our belief in the way that the Rain Fact might explain Jim's belief? No. For the Rain Fact to explain Jim's belief, he must learn the Rain Fact by correlating independent experiences of air quality and rain. Likewise, we would need to learn the normative law by correlating

²⁹ This qualifies as the explanatory relevance fact's explaining the belief in explanatory relevance, since the observed instances that served as the basis for the abduction process were explained (or meta-explained) by the explanatory relevance fact.

independent experiences of survival and goodness. But this would make the CE proposal viciously circular, since it requires the relevant experiences of goodness to be facilitated by knowledge of the law.

The final option would be to adopt a third-factor explanation. HOC explanations are non-starters, for the reasons given above. CE solutions raise a dilemma for normative grounding laws: either the law is fundamental, or it isn't. If it is, then it obviously can't share an explanans with our inference-grounding states. If it isn't, then the explanation for the law must be purely normative (on pain of violating Discontinuity). But if the normative laws are grounded in further normative facts, then this requires normative facts to explain our mental states, an option we set aside following the vast majority of self-proclaimed non-naturalists. Unless they are willing to explore such normative explanation of our mental states, non-naturalists simply cannot use third-factor explanations to meet the Challenge.

5. Conclusion

There is a perfectly general challenge to explain how we come to know the truth. In many cases, that challenge can fairly easily be met. Perhaps I come to know that I have an occurrent belief by believing that I have an occurrent belief; my belief explains the truth. Perhaps I come to know the truth about today's weather by looking out the window; the truth explains my belief. Perhaps a meteorologist comes to know the truth about tomorrow's weather by inferring from the same current conditions that will explain tomorrow's weather; the same things that explain her belief explain its truth.

These examples can give the impression that in order to meet the challenge on hand, one need merely identify a possible explanation of one of these kinds for belief-truth correlation. But the truth is more complicated, especially in the case of third-factor explanations. In epistemically good cases, the meteorologist doesn't just *happen* to infer from today's weather to tomorrow's; she has an

epistemically well-founded belief or disposition to treat current weather *as* a common explanation. As such, her being correct isn't a matter of luck or coincidence.

Non-naturalists in metaethics face this same challenge. Many of them acknowledge that they can only respond to it by appealing to third-factor explanations. We have argued that this is a dead end. 'Higher-order correlation' explanations offer no help in meeting the epistemological challenge. And it is not enough to *posit* a 'common explanans' explanation; agents need to be able to *use* the relevant explanation to infer (in the broadest sense) from the natural to the normative. Meteorologists can make relevantly similar inferences because their inferential disposition stems from rigorous observation of instances of weather patterns, and is formed by induction or abduction. Nothing relevantly similar is available to the non-naturalist, for the 'instances' cannot be independently accessed: that would require the very sort of truth-to-belief explanation unavailable to the non-naturalist from the get-go. We conclude that unless the non-naturalist embraces normative explanation of our normative beliefs—something most understandings of his view clearly prohibit—his view is inconsistent with the existence of normative knowledge.

References

- Bader, R. forthcoming. "The Grounding Argument against Non-Reductive Moral Realism." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau.
- Benacerraf, P. 1973. "Mathematical Truth." *The Journal of Philosophy*, 661–679.
- Bengson, J. 2015. "Grasping the Third Realm." In *Oxford Studies in Epistemology*, 5:1–38. Oxford University Press.
- Berker, S. 2015. "Does Evolutionary Psychology Show That Normativity Is Mind-Dependent?" In *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*, edited by J. D'Arms and D. Jacobson, 215–52. Oxford University Press.
- Clarke-Doane, J. 2016. "What Is the Benacerraf Problem?" In *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity*, edited by F. Pataut, 17–43. Springer.
- Dancy, J. 2006. "Nonnaturalism." In *Oxford Handbook of Ethical Theory*, edited by D. Copp, 122–45. Oxford University Press.
- Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- Faraci, D. 2015. "A Hard Look at Moral Perception." *Philosophical Studies*, no. 172: 2055–2072.

- Fitzpatrick, W. 2008. "Robust Ethical Realism, Non-Naturalism, and Normativity." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, 3:159–205. Oxford University Press.
- Gettier, E. L. 1963. "Is Justified True Belief Knowledge?" *Analysis*, 121–123.
- Goldman, A. 1967. "A Causal Theory of Knowing." *The Journal of Philosophy* 64 (12): 357–372.
- Huemer, M. 2005. *Ethical Intuitionism*. Palgrave Macmillan.
- Jenkins, C. S. 2006. "Knowledge and Explanation." *Canadian Journal of Philosophy* 36 (2): 137–163.
- McPherson, T. 2012. "Ethical Non-Naturalism and the Metaphysics of Supervenience." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, 7:205–34. Oxford University Press.
- Oddie, G. 2009. *Value, Reality, and Desire*. Oxford University Press.
- Shafer-Landau, R. 2003. *Moral Realism: A Defence*. Oxford University Press.
- Skarsaune, K.O. 2011. "Darwin and Moral Realism: Survival of the Fittest." *Philosophical Studies* 152 (2): 229–243.
- Street, S. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127 (1): 109–166.
- Unger, P. 1967. "Experience and Factual Knowledge." *Journal of Philosophy* 64 (5): 152–173.
- Vavova, K. 2014. "Debunking Evolutionary Debunking." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, 9:76–101.
- Wedgwood, R. 2007. *The Nature of Normativity*. Oxford University Press.
- Werner, P. J. forthcoming. "Moral Perception Without (Prior) Moral Knowledge." *Journal of Moral Philosophy*. doi:10.1163/17455243-46810075.
- Wielenberg, E. J. 2010. "On the Evolutionary Debunking of Morality." *Ethics* 120 (3): 441–464.