

1. Introduction

The Genealogical Debunking Challenge for non-naturalism argues that there's something epistemically problematic about the origins of our normative beliefs.¹ There are two main versions of this objection. The first suggests that the conjunction of non-naturalist metaphysics and (most frequently) the evolutionary influences on our moral beliefs entails that our beliefs fail a modal condition on knowledge. The other version locates the problem in the lack of explanatory connection between the facts and our beliefs. This cannot be resolved by modal considerations. As I'll argue below, this version should be expressed in terms of epistemic luck. If the explanation for the belief (and/or justification) is disconnected from the truth, it's lucky that the belief is true. Specifically, it's lucky in a way that drives standard Gettier cases (like how the explanations for your belief that there's a sheep in the field and for the corresponding fact diverge). According to this articulation, the luck itself is the problem, not failure of a modal condition on knowledge.

Non-naturalism seems to preclude the sort of explanatory connection that beliefs and facts must exhibit for the beliefs to avoid problematic luck: either direct or third-factor explanatory connections. If these don't hold, it doesn't matter how probable it is that we'd accidentally get things right if there's no account of an epistemically relevant explanatory connection. Third-factor explanations, where the fact and the belief have a shared explanation, have been popular attempts. But, if the worry is about Gettier-sufficient luck, we will see that non-naturalism is unable to provide the kind of third-factor explanation that solves the problem.

I will sketch an account of moral perception that provides non-naturalism the explanatory

¹ I'll be discussing what I take to be the current standard understanding of non-naturalism: the natural/non-natural distinction is metaphysical rather than epistemological, and non-natural normative properties are sui generis and causally inert. For a more complete discussion of the type of view, see McPherson's (2015) discussion of "Moorean Non-naturalism."

connections needed. If we can see instantiated moral properties, then the moral facts can explain our beliefs about them, and we aren't subject to Gettier-sufficient luck. But for moral perception to be compatible with non-naturalism, moral properties cannot cause us to perceive them. However, non-naturalism can hold that cognitively-mediated perceptual learning can provide us perceptual access to moral properties, without the implausible commitment that wrongness, for example, causally affects our retinas. The most widely familiar version is cognitive permeation.² The idea here is that elements of our cognitive architecture sometimes affect our perception in a way that changes the contents of our perceptual experience. If moral concepts permeate our visual perception, we can see moral properties. This picture is compatible with non-naturalism, and puts moral perception on par with perception of other higher-order properties that many already accept. So, given the dialectically appropriate assumptions that non-naturalism is true and that we have justified true moral beliefs, this proposal can convert an epistemically insufficient third-factor explanation into a successful standard facts-explain-beliefs structure.

The paper will proceed as follows. In section 2, I will explain the nature of the challenge, and why third-factor responses cannot succeed. The version or thread of the challenge that I'll be considering is one that hinges on explanatory connections between beliefs and the facts they are about. Third factor explanations can sometimes be sufficient, but I will show that to be so they must be buttressed by other beliefs that are explained by the facts they are about. In section 3, I will give the first sketch of the perceptualist solution available to non-naturalists. If non-naturalism is compatible with perceptual access to moral facts, then the facts would explain perceptually-based beliefs, ensuring the required explanatory connection to meet the challenge. In section 4, I will consider a number of objections that are raised by my sketch.

² This has historically been called "cognitive penetration," but the literature is in the middle of a transition to "permeation" instead.

2. Bad Luck for Non-Naturalism

The Genealogical Debunking Challenge identifies a range of epistemological problems for Non-naturalism.³ In its most common framing, The Evolutionary Debunking Challenge suggests that the combination of non-naturalist metaphysics with the fact there have been evolutionary influences on the moral beliefs we hold entails that it would be a massive coincidence if we are getting it right. Such a strong possibility of error is a defeater for whatever antecedent justification we might have had.⁴ Call this the Probability of Error Problem. The standard response, which I call the “Expectable Coincidence Strategy,” is to argue that the coincidence isn’t so big (and so not so problematic) after all.⁵ And so, our beliefs can meet any plausible modal condition on knowledge, like Safety or Sensitivity.⁶ The argument starts with the claim that the relevant moral truths are metaphysically necessary, and so not modally variable.⁷ Then, they argue that the actual (relevant) explanation for our beliefs, such as the evolutionary pressures that actually shaped us, is strongly modally stable. And so, given *these* truths, and given *these* pressures, we would expect our beliefs to be true enough on average for us to be well placed to make further progress and improvements over time. Sure, there is some coincidence, but it’s exactly the kind of coincidence we should expect

³ The Debunking arguments that I’ll be discussing are also pressed against stance-independent realism, and minimal realism. But I’ll focus on Non-naturalism, since its commitments make the challenge most difficult, and any solution available to it should be available to stance-independent and stance-dependent naturalisms.

⁴ See Joyce (2006) and Street (2006). It doesn’t matter whether justification is defeated directly, or whether knowledge is undermined and this in turn undermines justification (when we know that we don’t know).

⁵ See for example, Enoch (2011, ch. 7). See Baras (2022) for thorough discussion of when a coincidence is problematic and calls for explanation.

⁶ See for example, Clark-Doane and Baras (2021).

⁷ One of the presuppositions that is required for the Debunking Challenge to be an interesting and distinct epistemological challenge is that at least a significant portion of our beliefs about morality are justified and true. See Copp (2018), Enoch (2011), Korman and Locke (2020), Skarsaune (2011) and Vavova (2015) for related discussion. If they aren’t granted to be true, then we don’t have knowledge anyway, and it doesn’t matter whether the debunking challenge succeeds (would succeed if the beliefs were true). A similar concern holds for the justification of our moral beliefs. If justification isn’t granted, then the genealogical considerations are not introducing any special worry (they don’t issue a defeater for justification truth of the relevant beliefs). This by no means gives the non-naturalist a free epistemological ride. As I’ll discuss below, there is a point where the dialectical assumptions have to be made good on.

(given the assumptions), and not one that comes with an epistemologically concerning probability of error.

The other framing of the Challenge presses an adjacent epistemological problem that is not resolved by demonstrating that our beliefs meet certain modal conditions on knowledge. This problem is that non-naturalism seems unable to provide the right sort of explanatory connections between the normative facts and our true beliefs about them. I take the strongest version of this problem to involve a non-accidentality condition on knowledge. Call this the Accidentality Problem. Sometimes “accidentality” is cashed out in terms of modal conditions, such that meeting Safety secures non-accidentality.⁸ But that’s not the sense of “accidentality” that’s at issue here. Instead, I’m concerned about the sense in which coincidences are always accidental, no matter how expectable they might be.

This Accidentality Problem can be traced to two places in mid-20th century epistemology. The connection to the Benacerraf Problem has been widely discussed elsewhere.⁹ The connection to the Gettier Problem can be seen through considering what types of epistemic luck produce Gettier cases.¹⁰ There is “Benign Luck,” for instance when it’s unlikely that you would find something out, but you did.¹¹ If lightning didn’t strike at just the right moment, you wouldn’t have witnessed the murder, but since it did you know who committed it. This kind of luck has no negative epistemic repercussions. Then, there is “Environmental Luck,” where the way that it was unlikely for you to form the true belief undermines knowledge because would have been too easy for you to form a justified false belief instead.¹² This is the kind of luck you have in Fake Barn County. Whether or not environmental luck cases are Gettier cases, there is a distinct kind of luck

⁸ Enoch (2011, ch. 7, p. 174).

⁹ Faraci (2019(b)).

¹⁰ See Ichikawa and Steup (2017) for discussion and further references on the Gettier problem and epistemic luck.

¹¹ Prichard (2005, ch. 5) discusses several types of Benign Luck, the distinctions between which are not relevant. All that matters is there is a kind (natural or disjunctive) of luck that is not epistemically problematic

¹² Prichard (2005, ch. 6).

at work in many Gettier cases: “Intervening Luck.”¹³ This is when the explanations for the facts and for the beliefs diverge. The explanations for Smith believing the person who gets the job has 10 coins and for the corresponding fact diverge. The explanations for you believing there’s a sheep in the field and for the corresponding fact diverge.¹⁴ Because the explanation for the belief (and/or justification) is disconnected from the truth, it’s lucky that the belief is true. We can build this kind of case easily: take a false belief that’s justified enough for knowledge (had the belief been true) and make it true “some other way.”¹⁵

So, the Accidentality Problem for non-naturalism is that there is a sort of explanatory connection that beliefs and facts must exhibit for the beliefs to be in epistemically good standing, and non-naturalism seems to preclude such connections.¹⁶ There are three structures that explanatory connections might take. The Standard structure is that the facts explain the beliefs; you believe it because it’s true. For instance, the fact that you’re reading this explains your belief that you’re reading. The Reverse structure is that the beliefs explain the facts. This is a special kind of mind dependence, where the content of the belief grounds the fact that the belief is about.¹⁷ This could involve constitutive explanation (like if etiquette beliefs constitute etiquette facts) or grounding explanation (like with self-referential or self-entailing beliefs, such as “I have at least one self-referential belief”).

¹³ It also doesn’t matter if there is a third family of Gettier cases. All that matters is that there is an “Intervening Luck” family of Gettier cases, and that the epistemological problem at hand falls into that category.

¹⁴ Of course, these last two kinds of luck often overlap in Gettier cases. But what’s important is that there can be pure versions where only one kind of luck holds in the case.

¹⁵ Zagzebski (1994).

¹⁶ This is something like Enoch’s (2011, ch. 7) version of the problem, but with some important differences. His discussion begins with mention of an “epistemic access” framing of non-naturalism’s epistemological problem, but puts this aside because, first it isn’t clear what access means, and second because if he can defend non-naturalism’s ability to secure reliability, justification, and knowledge, what further epistemic problem could lack of “access” give us? But if “access” is about an explanatory condition on knowledge, Enoch cannot secure knowledge without addressing access. And Enoch punts on Gettier from a cautious optimism: if he’s been able to secure so much, why not think he can do so for whatever we eventually decide is the anti-Gettier condition. But if having the right kind of explanatory connection is (part of) the anti-Gettier condition, then the non-naturalist shouldn’t be quite so optimistic.

¹⁷ There are also some causal “beliefs about the future explain the facts about the future” self-fulfilling prophesy type cases, but these aren’t applicable to normative beliefs/facts.

Finally, the Third-Factor structure is when the fact and the belief have a shared explanation. Beliefs about the future on the basis of the present or the past take this kind of structure.¹⁸ We can have knowledge about tomorrow's weather, for example, on the basis of today's weather. We can know where the ball is going to land based on its initial trajectory. We can know who is going to win the election based on early results reporting. In each of these cases, there is a fact A that explains both fact B and fact C (that we believe that B).¹⁹

But none seem available to the non-naturalist. The Standard structure is unavailable because normative properties or facts are taken to be causally inert on non-naturalism, and there seems no other way for those facts to explain our beliefs.²⁰ The Reverse structure is unavailable because non-naturalism rejects mind-dependence. The reasons that the Third-Factor structure is unavailable are complex, but boil down to this: Gettier-preventing third-factor explanations (how A explains C) will involve induction or abduction, and agents cannot get into a position where such an inference would be epistemically licensed without having Standard explanation beliefs (which are apparently unavailable for non-naturalism).²¹

The Third-Factor structure has been taken to be a promising route.²² I think this is to an extent due to a conflation of the ways the Challenge can be pressed using Environmental Luck and Intervening Luck. Third-factor explanations are indeed promising for showing why we should expect our normative beliefs to line up (well enough) with the facts (given the dialectical

¹⁸ Korman and Locke (forthcoming) revise their (2020) view about domain-explanation in order to correctly capture such cases. I don't think their level of detail is required for my project. What I have in mind are cases of inductive or abductive inference, or their sub-personal quasi-inferential counterparts (e.g. the trajectory-based belief), which are the best cases for third-factor explanations (critiques of explanationism from induction and abduction notwithstanding).

¹⁹ Or, there is a slightly more complicated abductive structure where fact C is good evidence for fact B because A is the best explanation for C, and A would also bring B about.

²⁰ Bengson (2015) gives a kind of constitutive explanation similar to something I discuss in Section 3 as a contrast to my perceptualist solution. This kind of explanation is insufficient for roughly the same reasons.

²¹ This is one of the main contributions of Elliott and Faraci (MS). What I say here draws from that work. See also Faraci (2019(b)).

²² For two early examples, see Enoch (2011) and Wielenberg (2010).

assumptions), at a suitable degree of modal stability, to circumvent the Probability of Error Problem.²³ But for the Accidentality Problem, we need to be careful about whether the explanatory connection has the features required for resolving Gettier cases.

Take the following as a dummy Third-Factor story for normative beliefs: The fact that this event causes pleasure (A) explains the fact that this event is good (B), and it also explains why I believe the event is good (C).²⁴ Unfortunately, for this structure to help with Gettier cases, the normative belief must be formed in light of the explanatory relevance of the natural third factor, and not just in light of the natural third factor.²⁵ Compare to a non-normative case. You look out into the field, and because of your perceptual experience, you believe “there’s a sheep in the field.” As it turns out, what you saw was a decoy sheep. But there’s also a roving sheep who really likes this decoy, and so doesn’t stray too far. In this case, your belief and the fact are both explained by the fact that there’s a decoy in the field. But because your belief was formed under ignorance of the explanatory relevance, this belief isn’t knowledge. If someone who knows that a real sheep is attracted to the decoy sees the decoy and infers that the real sheep is in the field too, their belief is knowledge.

Further, the explanatory relevance belief must itself be in good epistemic standing. Compare three variations on the same inference, “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” (and grant for the illustration that it is sound). First, imagine that you believe that sheep decoys attract sheep through some specious armchair evolutionary psychological reasoning: sheep evolved to travel in

²³ This would be the case for both “ersatz” third-factor explanations, like Enoch’s (2011) pre-established harmony account, as well as true third-factor explanations, like Killoren’s (2021) and Wielenberg’s (2010).

²⁴ A case like this doesn’t violate non-naturalism’s commitment to the normative being *sui generis*, as long as A is only a *partial* explanation of B. For instance, perhaps there are explanatory moral principles, normative grounding laws, or irreducible normative essences, that also play some role.

²⁵ This can be construed in different ways. It could be that the explanatory relevance fact is a co-basis of the belief with the third factor, as with explicit inference. Or it could explain the transition from the third factor to the belief. An example of this would be subpersonal quasi-induction, as we might find with Pavlovian conditioning, for example.

flocks, so if a sheep sees a decoy sheep, the real sheep will attempt to start a flock with it. Lacking justification, this belief is unable to transmit or assist with the transmission of good epistemological standing through inference, and the belief that there is a sheep in the field is not knowledge.

Second, imagine that you believe that sheep decoys attract sheep because you read that same evolutionary psychology story in the journal *Nature*. In this case the explanatory relevance belief is in good standing, and defuses the original Gettier case—the belief that there is a sheep in the field is knowledge. Third, imagine that the issue of *Nature* was the inaugural April Fools’ issue of joke science, and you read it before the editors announce the trick. Despite being justified and (for the sake of the illustration) true, the explanatory relevance belief is not in good standing in this case. Because it results from Intervening Luck, it too is a Gettier-belief. But such a Gettier-belief cannot defuse the original Gettier case, and the belief that there is a sheep in the field is not knowledge.

Returning to the normative case, recall that natural fact A is the common explanation for the normative fact B and the belief C (that B is true). The fact “natural A is explanatorily relevant for normative B” is itself a normative fact.²⁶ Of course, A might be relevant for different reasons. There might be an explanatory principle like “events causing pleasure are good.” There might be a non-naturalism compatible essence that explains why pleasure is relevant for goodness. But no matter what underwrites the explanatory relevance fact, it is a fact about the nature of the normative, and is subject to all our epistemological concerns about non-natural normative facts. So, if we rely on the explanatory relevance fact, we have to make sure *our belief in that fact* is in good standing. But we can only do that if we have solved the Accidentality Problem. If we have not, then the explanatory relevance belief is due to Intervening Luck, and we are in a case like the April Fools’ case above. Thus, the third-factor story can’t play a leading role in solving the Accidentality Problem.

²⁶ Heathwood (2012).

If this was all correct, the Accidentality Problem suggests that non-naturalism entails that we're a Gettier position with regard to our normative beliefs. But, you might respond, I can't know that I'm in a Gettier position, because that entails that my belief is true. And if I know that my belief is true, then my belief is knowledge. This is a mistake. The Challenge dialectically grants, or assumes for the sake of the argument, that our beliefs are true—otherwise there's no truth to explain. But like in a natural deduction from assumptions rather than axioms, we have to discharge the assumptions to draw the real conclusion. So, *if* our beliefs are true, they're Gettiered. This fact, that we're "Gettiered-at-best," is a defeater for our normative beliefs, and doesn't provide independent evidence or justification for those beliefs.

3. Looking, for the Solution

I will sketch an account of moral perception that a non-naturalist can use to secure the required explanatory connection. If we can see instantiated moral properties, then the moral facts can explain our beliefs about them, allowing a Standard structure explanation. But for moral perception to be compatible with non-naturalism, moral properties cannot *cause* us to perceive them. So, the strategy will be to use cognitively-mediated perceptual learning to allow us perceptual access to moral properties, without the implausible commitment that wrongness, for example, causally affects our retinas. The idea is that elements of our cognitive architecture sometimes affect our perceptual architecture in such a way that changes the contents of our perceptual experience. So, if moral concepts (for instance) permeate our visual perception, we would be able to see moral properties.²⁷ My contention is that the resulting picture is acceptable to non-naturalists, and puts moral perception on par with perception of other higher-order properties that many theorists want

²⁷ Nothing in principle limits this to visual sensory experience; the kind of perceptual learning at issue would be able to operate on any sensory modality, especially if it is a person's primary modality.

to accept anyways.

The Accidentality Problem suggests that non-naturalist metaphysics would entail that all our normative beliefs are Gettiered-at-best. The third-factor strategy used to establish an expectable coincidence in response to the Probability of Error Problem cannot resolve this other worry.²⁸ So, the non-naturalist must change tactics if she is to resolve the Challenge. She needs to find a way to have the normative facts explain our normative beliefs, and do so despite being causally impotent. There are many senses of explanation in the metaphysician's toolbox, and while it's clear that we aren't here limited to causation, it should also be clear that the available range isn't much wider. Rather than sorting through the various kinds, I hope a gesture at the right and wrong sorts will suffice.

While existential beliefs about the objects in our near surroundings might be formed causally, causation as such isn't what's important. Instead, it's that the beliefs are formed in response to the facts that the objects exist. If a potted plant fell off my shelf, and through brain trauma caused my belief that there was a broken potted plant on my floor, it would not be formed in response to the fact that there is. If externalism about content is true, and my beliefs are partially composed of the things they are about, this would give a kind of *constitutive* explanation of my beliefs by the facts. But, as in the previous example, such externally constituted beliefs aren't necessarily formed *in response to* the facts.²⁹ So, while the right sort of explanation might involve causation, grounding, or some other relation, it's only the right sort if the belief is formed in response to the fact.

Moral perceptualism delivers responsiveness. Further, (unlike Intuitionism) there's no need

²⁸ The kind of third-factor explanations for the correlation between our beliefs and the facts are not the kind that turn resolve standard Gettier cases.

²⁹ Consider how the content-constitutive explanation would misclassify the explanatory structure of beliefs about the future as "the fact explains belief," if we use it to provide a Standard structure explanation for normative beliefs. Further, it seems that there could be no Gettier cases if the content-constitutive explanation resolved Gettier cases, since the beliefs are stipulated to be true.

to motivate a new and special faculty, because our current accepted faculties will do all the work. The first trick is that while perceptual explanation for beliefs might be thought to operate through causation, the non-naturalists' explanation cannot require the instantiation of normative properties to affect what light reaches our retinas, or what sound waves reach our ears, and so on. While that might sound like we've already reached a dead-end, there is precedent for the view. There is a range of views that hold that natural kind properties (or social, or historical properties) can be represented in the contents of our visual experiences, despite not causally affecting our sense organs. Non-naturalists can appeal to those that explain this through cognitive influences on our perception. I've already mentioned the cognitive permeation view, where our beliefs play a role in the determination of our perceptual content. Other perceptual learning views that hold that our beliefs can allow us to develop new perceptual abilities will serve just as well.³⁰

Siegel argues for a "rich content" view of this kind, that the contents of our perceptual experience can include higher-level properties, by using the "phenomenal contrast" method.³¹ This method employs the principle that if there is a contrast between the perceptual phenomenologies of two perceptual states with the same sensory inputs, the states have different content. Then, if we can identify a pair of perceptual states for which the best explanation for such a content difference is that one represents a given higher-level property that the other does not, we have compelling reason to believe that such higher-level properties are represented in perceptual experience.³² Siegel's stock examples include natural kind properties, such as "pine tree," the semantic properties of a new language and script, and your neighborhood when you moved to a new city. These look one way to you when you first see them and don't know what they are.³³ But once you become familiar with

³⁰ For instance, see Burston (forthcoming).

³¹ See Siegel (2006), (2007), and (2010).

³² Siegel (2010, pp. 99-115).

³³ I won't rehearse the counter arguments and alternative explanations that Siegel considers, since my goal is largely to show that non-natural normative properties should be on par with these kinds of properties.

pine trees or learn the language or understand the geographical relationships, they look different. Pine trees look like they are pine trees, and in this way stand apart from other plants in the forest. Words and letters used to look like shapes, but now their semantic meanings are visually apparent. Parts of your neighborhood look oriented with regard to other parts.

It's important to recognize the significance of these divergent experiences sharing the same sensory inputs. According to the "thin content" theories that deny that we experience higher-level properties, there are few properties represented by our perceptual experience, and these are generally limited to those that causally affect our visual system. The only properties that could affect our visual system are the ones that are involved in determining the pattern of photons that strike our retinas. These are plausibly exhausted by things like color, shape and motion. But, if the rich content view is true, then some higher-order properties are represented despite not causally affecting our sensory organs. This is great news for the non-naturalist. She can hold that non-natural normative properties can feature in our perceptual experience, without violating her commitment to the normative being causally inert.³⁴

Someone might object that the relations between the property of being a pine tree and the color and shape properties that affect our retinas are different than the relations between being wrong and whatever properties affect our retinas when we see wrongness. This overstates any connection between being a pine tree and whatever color and shape properties any given pine tree may have.³⁵ Pine-tree-ness isn't constituted by color and shape properties, so we can't offer a tight third-factor connection between our perceptions and the pine trees. Such a story would be even less plausible for semantic properties. While there is some "in virtue of" relationship between the shapes

³⁴ Or, we might be inclined to be liberal about causation and say that pine tree properties and semantic properties do cause us to see them. But the sense of causation that it would take for semantic properties to cause us to see them is so weak that it's compatible with non-naturalism.

³⁵ See Buckner (2018) for a discussion of machine learning and problems with recognition of rich content, demonstrating that there is no codifiable pattern of high-level properties (rich content) corresponding to low-level properties (thin content) withing perception.

or sounds and the meanings, it won't be fundamentally different from how normative properties are instantiated in virtue of natural properties.

With this as our model, we can now construct the basic picture. Since this is a response to the Debunking Challenge, we start with the assumption that we have justified true beliefs about the normative.³⁶ So, those beliefs employ concepts about the non-natural normative. If these concepts permeate or otherwise affect our perception, we would perceive particular moral facts. And if we perceive particular normative facts, those facts could explain our beliefs about them. These fact-explained beliefs will enter our greater cognitive architecture, and "being explained by the facts" would propagate across our other normative beliefs.

In this way, we could bootstrap from (perhaps) a problematic Third-Factor epistemology into a non-problematic Standard epistemology. On Enoch's Third-Factor story, there are evolutionary pressures that selected for a range of moral beliefs that put us in a good position for rational revision to get us relatively close to the truth. While this responds to the Probability of Error Problem, it doesn't involve the right sort of explanatory connection to resolve the Accidentality Problem. And because this story leaves us in a Gettier position, our beliefs wouldn't be in good epistemic standing, even if otherwise justified. But, if this gives us perceptual access to particular normative facts through perceptual learning, from that point forward we would have some beliefs explained by normative facts. The explanatory connection will propagate through the system of our normative beliefs when and because the perceptual beliefs enter into reflective equilibrium (or some other coherence process) with the rest of our beliefs. Other normative beliefs will be maintained (in part) because of their fit with the perceptual normative beliefs.³⁷ And as such

³⁶ This assumption may have to be weakened or vindicated at certain points in order for the solution to not feel cheating or question begging, but it is still the starting point in response to the Challenge.

³⁷ This leads to a sort of explanatory overdetermination, even when (if) we separate forming-explanations from maintaining-explanations. But the sufficiency of the one track of required explanatory connection to the facts should not be undermined by the presence of parallel explanatory tracks that do not connect to the facts.

explanatory connection diffuses through the rest of our normative beliefs, the Gettier problem dissolves.

4. More Problems with Moral Perception?

As an optimistic sketch, this should look promising. There are, of course, some obvious (and less obvious) problems that have to be addressed for the picture to fill out in a satisfying way. First are the empirical problems. It is an empirical psychological question whether cognitively-mediated perception ever occurs, let alone with the specific concepts needed here. I will put these aside. This is a strategy paper, not a vindication paper, so I don't have to tie up every loose end. And as a philosopher rather than a cognitive scientist, my skills are better put to use exploring the conceptual space for this strategy than its empirical adequacy. That being said, I think that the current state of the literature gives good reason to be hopeful that the empirical side will be borne out.³⁸ Also, I may not even strictly need cognitively-mediated perception; other accounts of higher-level perceptual content might do just as well. Perceptual beliefs of higher-level properties need to have an explanatory connection to the facts they're about, and such beliefs are generally best-case examples of explanatory connections holding.³⁹ If cognitive-mediation isn't the phenomenon that explains how higher-level properties come to be represented in our perceptual beliefs, then whatever it is that explains this must be enough to establish a sufficient explanatory connection (on pain of skepticism. All that is needed is for moral perception to be on par with more conventional perceptual experience.

³⁸ See Stokes (2013) for a helpful primer on the scientific literature on cognitive permeation, written for philosophers.

³⁹ Korman (2014) however thinks the inference should go the other way: given the nature of sight, there is not robust enough of an explanatory connection between ordinary objects and our beliefs about them, so that those beliefs are in epistemic bad standing. In other words, my partners in innocence argument would be partners in guilt for someone who is willing to give up such beliefs. Of course, we are then subject to a broad skeptical conclusion, and the evolutionary debunking challenge for non-naturalism (or realism more broadly) is not an interesting and distinct challenge. (See Vavova (2014) for discussion of the need for the challenge to be a special one.)

Then, there are problems from the dialectical assumptions of the paper. The genealogical debunking challenge suggests that if non-naturalism is true, then even if we have justified true normative beliefs, there are serious epistemic problems. While a reply to this challenge is appropriately situated inside the context of the attempted conditional proof, at some point the assumptions have to be discharged and independently vindicated. In addition to needing to address justified true belief, this raises a reference problem: it's not immediately obvious how to reconcile non-naturalism with some of the more dominant meta-semantic theories.⁴⁰ If the concepts fail to refer to non-natural normative properties (even if such properties exist), cognitive influence can't make non-natural normative properties represented in our perceptual experience, and so our beliefs wouldn't be explained by them. Of course, non-naturalism has this reference problem independently. And if our concepts, and so our beliefs, don't refer to non-natural normative properties, the fact that this story doesn't work is moot anyway. So, I'm not trading one problem for another; non-naturalism already has both. In light of this, and for concerns about space, I'll put this question to the side as well. But because my picture directly relies on accurate reference, I should at least say something in support of optimism about reference, and so I'll make some remarks in the concluding section.⁴¹

Next, cognitive-mediation and other views bring along their standard epistemological problems, and my appeal to cognitively-mediated perception in this specific way introduces particular challenges. The epistemic problems can be sorted in different ways. First, there are standard worries about cognitive-mediation, special worries about cognitive-mediation in this context, and worries about whether the picture can play the correct dialectical role. On the other

⁴⁰ Werner, (2020).

⁴¹ While this may look like a liability, this is actually why a cognitive-mediation-free rich content perceptualist view is not available for non-naturalists. The perceptual content has to properly be *about* the non-natural normative, and if that is not inherited by accurately referring beliefs it's not clear how non-naturalists can secure it.

sorting, there are justification problems and there are explanationist problems. It's important to recognize that these are separate kinds of challenges, given that my story bootstraps the explanatory connection. We start out in an epistemically suspect position because we lack the right explanatory connection between our normative beliefs and the facts, *not* because we lack initial *prima facie* justification (given the dialectic). Then, the connection becomes established once cognitive-mediation gives us perceptual access to those facts. Given this epistemological role, perception doesn't need to provide foundational or basic justification for our normative beliefs. It only needs to give an explanatory anchor for the system of our normative beliefs, and to be epistemically on par with natural kind perception. So, it's no problem that the perceptual moral beliefs don't provide novel justification. There might be a question of whether this account gives the right sort of explanatory anchor, but after considering an alternative expression of the explanationist requirement, we'll see that the perceptualist picture must be sufficient.

4.1 Theory-Laden Perception and Epistemic Downgrade

A standard epistemic problem for cognitive permeation (which I'll use as the exemplar from here out for simplicity) is that if perception retains its standard epistemic significance, there would be many cases of justification-bootstrapping that we should find problematic. You believe the world is thus-and-so, this makes the world look thus-and-so, and so you strengthen your belief on the basis of the appearances. A common case to illustrate this is Jill fearing that Jim is angry with her, then this fear makes his neutral expression appear angry to her, and so she forms or strengthens the belief that he is angry. We might consider this analogous to a case of circular testimony. I tell you that campus is open during the holiday, then later you tell me that campus is open during the holiday, and then I strengthen my belief that it is so.

But in other cases, cognitive permeation seems benign or even beneficial. An x-ray looks

different to a technician than to me in a way that puts her in a better epistemic position. The text and sounds of a new language look and sound different after you've learned it, and this improves your epistemic position. Because cognitive permeation can tune you in to novel higher-level properties, your perception can become more epistemically powerful by conveying more information.

This indeed looks like a difference between some of the good cases and some of the bad cases. Jill's fear that Jack is angry isn't responsible for her capacity to visually represent the property "being angry" in the way that the x-ray technician's experience allows her to visually represent "benign spot" or "cancerous growth." But this can't be the whole story, because there are other examples where this enabling function of cognitive permeation might be epistemically malicious. Perhaps people who believe in fortunetelling by reading tea leaves (or stars, or palms) have visual experiences of properties that I can't have, such as "auspicious" or "ominous." But the extra content in their visual experience of the same cup are epistemic liabilities rather than advantages. Forming beliefs on the bases of these extra contents worsens their overall epistemic position. So, to evaluate a cognitive permeation based on perceptualist moral epistemology, we must figure out whether perceiving moral properties is more like the x-ray case or more like the tea leaves case.⁴²

The foremost concern at this stage is about circularity, since it's important to know what beliefs such perceptual experiences can support, and how they relate to our broader cognitive architecture.⁴³ Consider the classic case of physicists setting up a container of supersaturated gas in an attempt to test for the existence of alpha particles. Suppose that, when observing the experiment, in addition to contents like "the gas is whitish" and "that's a condensation trail," the physicists'

⁴² After which, we might evaluate *particular* examples on a case-by-case basis to see if they're like Jill's case, such as when we think someone is a villain and so are disposed to see their actions as wrong. But that would no more call into question the good cases than Jill's case calls into question the x-ray tech's case.

⁴³ Siegel (2017) discusses several ways our epistemic position could be worsened in cases like this, which she calls collectively "downgrade." The other two kinds of downgrade don't raise special concerns for my picture of moral perception.

visual experience included the element “there goes an alpha particle.” This results from the visual inputs and cognitive permeation by the contents of the beliefs that constitute her theoretical commitments. What this means is that the downgrade is localized both to this element of experience and to its ability to support *certain* downstream beliefs. Since the perceptual experience did not take the existential belief “that is an alpha particle” as an input, it can still support a belief with *this* content without circularity. But because the perceptual experience took (belief in) the broader theory as an input, neither the experience *nor* the new existential belief can add support to the theory.⁴⁴ Of course, the element of experience “there is a condensation trail” and *its* corresponding belief can add (abductive) support to the theory, so the visual experience at large could still have this epistemic power. It’s only that a specific element is downgraded with regard to specific beliefs.

How does this apply to cases such as when Harman walks around the corner to find some children burning a cat? Supposing that his visual experience includes the content “that is wrong” in part due to cognitive permeation by his general moral beliefs, we should have a similar story. If the permeating belief is well-founded, and the experiential output is supported, that experiential element is well-founded. And so, a belief that what the children are doing is wrong formed on the basis of this element of experience would be well-founded. But neither the experience nor the resulting particular belief has the power to strengthen the belief “burning cats is wrong.” All moral perceptions and perceptual moral beliefs will have this highly localized circularity downgrade due to their content resulting from some kind of inference, even though not every inferential perceptual experience (and corresponding belief) will exhibit this downgrade.

In light of this, it appears that whenever *the ability to perceptually experience a given kind* depends

⁴⁴ That said, *if* the Bayesian calculus says there is an increase the conditional probability for the broader theory on the basis of an experience or corresponding belief that depended on the broader theory, then even circular support may not be epistemically problematic. Thanks to Jonna Vance for their discussion of this point.

on cognitive permeation (or any cognitively-mediated perceptual learning), perceptual experiences of such properties will invariably exhibit this kind of limited circularity downgrade. For my purposes, it's not important to tease out exactly the extent of cognitive-permeation-dependent property perception, and a rough gesture will suffice. This kind of influence is required to perceive any higher-level properties, and so every experience of them will exhibit the same kind of circularity downgrade. There's still some question about the exact range of beliefs the local circularity downgrade prevents the experience from supporting, but what's important is that perception of moral properties are on par with perception of the sorts of properties that are accepted by theorists who don't limit perception to low-level properties.⁴⁵ So, just as moral perception and alpha particle perception will have this local impact from being theory-laden perception, so will pine tree perception and semantic property perception (and so on).⁴⁶

4.2 Does this provide explanatory connections between the right facts and beliefs?

The inability of moral perception to provide theory confirming justification might seem problematic, potentially undermining my attempted solution. Perceptualism is generally an *a posteriori* moral epistemology, but if our perceptions cannot support beliefs in theories the epistemology is impotent.⁴⁷ But the non-naturalist isn't forced to worry about this, because they aren't forced to accept a pure *a posteriori* epistemology. A mixed epistemology where *a priori* reflection plays a significant role is perfectly acceptable, and sufficient to respond to the

⁴⁵ Of course, this will only apply on pictures of rich content that have higher-order property perception depend *in some way* on an inferential structure akin to the one discussed above, where the ability to perceive some kind depends on having the right antecedent cognitive architecture (e.g. beliefs). I believe this will include all perceptual learning views.

⁴⁶ What about if we imagine that *there are* properties such as auspicious and ominous (or some other non-causal supernatural properties), would my view have the implication that we can gain visual access to them in the same way? My inclination is to say "no," but it depends on whether we can have concepts that successfully refer to those properties and justified true beliefs about those properties. While I'm optimistic about the prospects for normative concepts and beliefs, I think the moves to be made there will be hard to duplicate for other sorts of properties. Thanks to Tristram McPherson and Eric Wiland for pressing me on this point.

⁴⁷ C.f. Cowan (2015).

Genealogical Challenge as pressed here. The perceptualist element is only needed for securing an explanatory connection between our moral beliefs and the facts. And so long as our perceptual beliefs feed into our overall system in a way that they affect the overall basing relation for the set of our beliefs, this requirement will be satisfied.⁴⁸

To make this clear, I must first recap a few of the steps we've taken. First, the evolutionary debunking challenge must assume that we have otherwise-justified true beliefs about normativity. Second, it is suggested that the evolutionary pressures that the human species in fact underwent disposed us to have normative beliefs roughly as we in fact have them. Third, third-factor responses that are sufficient for the Probability of Error problem don't provide the right explanatory connection between our beliefs and the normative facts that they are about. But these beliefs could provide us with the ability to perceive particular normative facts, through cognitive permeation. And beliefs formed on the basis of those perceptual experiences will hold a sufficient explanatory connection to the facts they are about. So, this story involves some nativism about (some) normative concepts, related to the evolutionary explanation for some beliefs.⁴⁹ And it involves some unspecified source of justification for those beliefs. The non-naturalist is perfectly free to accept *a priori* justification here.⁵⁰ Of course, a complete non-naturalist epistemology will give a detailed characterization and defense of this element, but in the dialectical context and for the purposes of this paper, this is sufficient. Since the evolutionary debunking objection is supposed to be compatible with all of this, all my picture is required to offer is the right explanatory connection

⁴⁸ Additionally, since the same structural concern should apply to perceptual beliefs about pine trees, so whatever we can say that allows us to gain new justified empirical beliefs about pine trees by perception will likely apply in the moral case as well.

⁴⁹ This is explicitly compatible with Joyce's (2006) presentation of the debunking argument, but it should be compatible with all versions. The space for reluctance or suspicion at this point must be about the successful reference of those concepts. But that is again a separate (non-epistemological) problem that non-naturalism must deal with anyway. I'll briefly address this in the conclusion.

⁵⁰ [still compatible with Debunker's set up, compatible with reference gesture, doesn't reintroduce the Probability of Error problem]

between the facts and our beliefs (without undermining any of the dialectical assumptions). And the perceptualist picture does this by having particular facts explain particular beliefs.

The natural objection at this stage is that the perceptualist picture doesn't provide enough explanatory connection, or connection to the right beliefs.⁵¹ After all, the background assumes that we have some degree of general normative beliefs that are not explained perceptually. And while my suggestion that the perceptual beliefs enter into reflective equilibrium with the antecedent general beliefs might serve to introduce *some* explanatory connections into our overall epistemological position, what is required is for the general beliefs to be explanatorily connected to the general facts. Without this, *those* beliefs are still accidental in the way that's sufficient for Gettier cases.

Framing this slightly differently, it looks like a dilemma arises from my suggested mixed epistemology, where the general beliefs take (foundational) *a priori* justification. As the first horn, it might look like the perceptualist element is insufficient, as it doesn't establish explanatory connections for the right range of beliefs. As the second horn, it might look like the perceptualist element is redundant. This would be the case if *a priori* belief formation provides explanatory connection of the right kind. In such a case, we wouldn't need the perceptualist element. But this would either require the mysterious capacity of Intuitionism, or to include constitutive explanations of the beliefs to count as the right sort of explanatory connection. I have rejected both of these already.

Fortunately, the horn of insufficiency is easy to evade (again, by pulling in an independent problem). Non-naturalists have to offer an account of normative explanation, an account of the metaphysics of particular normative facts and their relation to the natural facts that bring them about.⁵² Most, if not all, of the worked out views of non-naturalist normative explanation are

⁵¹ Thanks to Chris Heathwood for pressing me to clarify this step in my solution.

⁵² We see this most frequently in the Supervenience literature, as the ways supervenience is supposed to be explained on non-naturalism. Of course, those who reject supervenience have to offer accounts of normative explanation too.

generalist. That is, particular normative facts are, in part, explained by general normative facts.⁵³ So, the sketch from above just expands a little. Some particular beliefs are explained by perceptions of the particular facts. Since those particular facts are explained (in part) by general facts, when the corresponding general beliefs are maintained (in part) because of the corresponding particular beliefs (through reflective equilibrium), the general beliefs are transitively explained (in part) by the general facts they are about. Then, other non-perceptual particular beliefs can be formed or maintained in light of those connected general beliefs, and these will exhibit a non-accidental third factor explanatory structure: the general facts explain both the particular facts and corresponding particular beliefs (mediated by an appropriate relevance belief in good standing—the corresponding general belief). This is what I meant above when I said “explanatory connection diffuses through the rest of our normative beliefs.”

4.3 Does this provide *any* explanatory connection?

Some might remain suspicious that perceptual experiences due to cognitive permeation by general beliefs can serve to provide explanatory connections at all, as a partners in guilt argument. We can evaluate whether a mental state that governs the (sub-personal) transition between subdoxastic states establishes a sufficient explanatory connection by looking at the analogous explicit doxastic inference.⁵⁴ The fact that the inference is explicit shouldn't make a difference in the standing of the inferred belief (e.g. whether it's knowledge), so if the transition isn't good enough doxastically, it shouldn't be good enough subdoxastically. This is why third-factor explanations need to have the common factor explain the belief in light of its explanatory relevance for a target fact.⁵⁵

⁵³ I take Bader (2017), Enoch (2019), FitzPatrick (2014?), Leary (2017), and Skarsaune (2015) to all offer generalist explanations. This only makes the view unavailable for particularists.

⁵⁴ Faraci (2019(a)), Elliott and Faraci (MS).

⁵⁵ See note 25 for more characterization of this mediation.

And if we are motivated by an explanationist requirement, we'll probably think that explicitly inferring the moral conclusion from the non-moral perceptual basis and a prior (insufficiently connected) moral belief doesn't establish the right explanatory connection for the new belief. After all, the non-moral input doesn't establish an explanatory connection, and the prior belief lacks an explanatory connection (until demonstrated otherwise).

But when we consider other (non-normative) cases, rich content perception is supposed to provide exactly this right sort of explanatory connection, on the same basis. Again, Pine-tree-ness doesn't cause perceptual experiences of pine trees. And neither are there codifiable patterns for moving inferentially from properties admissible as thin content to perceptual beliefs.⁵⁶ And so, for pine tree perceptual experiences to be able to offer the epistemically necessary explanatory connection to the pine tree facts they present, this sort of picture must suffice for explanatory connection in general, on pain of overly-broad skeptical implications. And then there is no reason to exclude perceptual experiences with normative content due to the same subpersonal processes as issuing the required sort of explanatory connection.⁵⁷

4.4 Does this Satisfy the Right Explanationist Requirement?

My cognitive permeation story might not be a successful response to the Genealogical Challenge if the explanationist requirement I used to frame the challenge might not be the right one. Bedke suggests a modal condition on the connection between the explanations of our beliefs and the facts that the cognitive permeation picture might not meet.⁵⁸ But this "Obliviousness" condition

⁵⁶ See Buckner (2018), for a discussion of machine learning and problems with recognition of rich content.

⁵⁷ While this doesn't address the heart of the challenge, showing why the subpersonal transition between thin and rich content is epistemically upstanding while the analogous explicit inference is not, it either suffices as a response (the objection proves too much) or it radically weakens the explanatory requirement to include the explicit inference, offering way more options to the non-naturalist. I prefer the former.

⁵⁸ Bedke's (2014) presentation thus initially seems to resemble the modalist Probably of Error problem. But he constructs it to involve counterfactuals or counter-possibles involving the explanations of our beliefs in order to evade

also threatens to overgeneralize if it rules out the kind of cognitive permeation at work in moral perception.

Bedke offers his Obliviousness condition as a way of getting to the precise possibility of error that should be used in the Probability of Error argument, and in response to the modalist solutions that lean on the necessity of moral truths.⁵⁹ The test involves constructing an “allodoxic” possibility, and using this to assess the connection between the belief and the fact. Allodoxic possibilities are a kind of internalist epistemic possibility, based around the hypothetical falsity of a given belief, which allows us to avoid the complication that there are no metaphysically possible counterfactuals where the fundamental moral truths are different. In the possible world where the belief is false, we look to see if you would have the same belief and the same justification with the same explanations. If so, then your belief is Oblivious to the facts. And being Oblivious is a defeater for whatever justification you might have.

To test my proposal, we don’t evaluate the background beliefs before the perceptually explained beliefs enter the picture and transmit their explanatory basis to the rest.⁶⁰ But at first look, the perceptual beliefs would be just as Oblivious. Like with the concern above, if the moral belief input is Oblivious, and the perceptual input doesn’t have moral content, it seems like the perceptual experience doesn’t offer the allodoxic counterfactual connection to the moral facts that would keep a downstream belief from being Oblivious. If the children burning the cat wasn’t wrong, we would have the same background belief on the same basis, and we would have the same perceptual inputs on the same basis, and so we would have the same moral element of our perceptual experience, and

the sort of Expectable Coincidence strategies described above. Since beliefs have to have the right sort of explanatory connection to the facts to pass the Obliviousness Test, Bedke’s is an explanationist challenge, a version of the Accidentality Problem.

⁵⁹ Bedke (2014, p. 116-119).

⁶⁰ This requires taking a stance on how and when defeaters defeat justification, but hopefully not a very strong one. This must reject the existence condition on defeat, but doesn’t have to accept the possession requirement on defeat. A moderate access requirement is probably enough to block defeat for the background beliefs.

so a belief on the basis of this experience would have the same explanation and doxastic justification as in the actual world, and so this belief is Oblivious.⁶¹

But, if cognitive permeation (or similar subpersonal, quasi-inferential effects) is pervasive in our visual experience, then the Obliviousness test will disqualify too many of our perceptual beliefs (and the beliefs downstream from these) from being knowledge. My assessment is that the results would be massively skeptical, even if not universally skeptical: any inference-based perceptions that neither have the second purely perceptual track to the content, nor inferential inputs (concepts, beliefs, etc.) that are themselves explained by the things they are about would lead to Oblivious beliefs. But to not court too much controversy, I'll use a narrower range of perceptual kinds: those who inherit all their causal powers from their components or constituents. The kinds I have in mind are artifact kinds, social kinds, and (at least some) historical kinds.

Take the artifact kind "chair" as a test case. To test beliefs about chairs for Obliviousness, we need to examine what we would justifiably believe, and on what basis, if the beliefs were false. But we shouldn't look at particular existential beliefs, because the realist about chairs would say "in the allodoxic possibility where this chair wasn't here, I wouldn't believe that that was." This isn't the right comparison for the allodoxic falsity of *non-naturalism*. Instead, we need to examine the belief in realism about chairs, and so consider what we would believe if chairs nominalism were true.⁶² Since the property doesn't itself causally affect our visual systems, if there were no universal "chair," we would still have the same chair-experiences and so have the same chair-beliefs on the same justification and basis. Thus, our chair beliefs (indeed all our artifact-kind beliefs) are Oblivious.

⁶¹ Bedke only concludes that our beliefs would be Oblivious on the condition of non-naturalism, and so this puts pressure on non-naturalists to give up the view, since the conjunction "we know that burning cats is wrong and naturalism is true" is more plausible than the conjunction "non-naturalism is true and we have no moral knowledge." (2014, p 122-124)

⁶² If nominalism isn't the right contrast, we can run the same test on the allodoxic hypothetical "what if we were wrong about which things instantiate chairness," since Bedke's argument against non-naturalism could use either normative nihilism or "what if we were wrong about which things are wrong" as the allodoxic hypothetical. The same too broadly skeptical results follow from this alternate hypothetical.

The same applies, for example, to race and gender properties (on social constructivist realism), to dollars and counterfeits (as historical kinds), and to shadows and holes (as relational kinds). In all cases like these, the explanation and justification for our beliefs do not come from the properties themselves, unless (veridical) cognitively permeated perception provides the minimal explanatory connection required to avoid Obliviousness. And if this is the minimal requirement, our perceptual moral beliefs will satisfy it as well.⁶³

5. Conclusion

Non-naturalist responses to Genealogical Debunking challenges have focused on third-factor accounts. Such responses can successfully show that our normative beliefs are modally stable enough (given our actual histories, and the necessity of the relevant facts) to meet any plausible modal condition on knowledge. Unfortunately, they cannot offer satisfactory responses to the Accidentality worries that drive a standard class of Gettier cases. Even so, non-naturalists can offer a Standard structure, facts-explain-beliefs, explanations that will avoid Gettier accidentality if they adopt an account of moral perception. This has been overlooked, as perceptualist accounts at first seem to require a causal efficacy that is incompatible with standard characterizations of non-naturalism. But a closer examination of the phenomenon of cognitive permeation shows that non-naturalists can support moral perceptual content without the implausible commitment that, for example, wrongness influences how light strikes our retinas. Any rich content view, views on which we perceptually experience higher-order properties, will have to say that we experience properties that don't causally influence our sensory organs. And so, non-naturalist moral perception is on par with standard, common sense, views of ordinary object perception.

⁶³ Again, a debunker (such as Korman, 2014) might be willing to pay this price, but the point here is that it's not unreasonable for non-naturalists to reject a debunking set-up that requires this much.

While this sketch is promising for non-naturalists, at various places I punted on an independent problem for non-naturalism, a solution to which is needed for the perceptualist account to get off the ground. This was the reference problem. Non-naturalists must hold that we have accurately referring normative concepts for there to be cognitively permeation perceptual experiences with veridical normative content. But most standard theories of metasemantics aren't easily compatible with non-naturalism.⁶⁴

Here too, I believe that the non-naturalist has reason for optimism. My brief suggestion is the following. As flagged above, my perceptualist response to the challenge employed a mixed epistemology, where there are antecedent general beliefs that bear *a priori* justification, as well as the subsequent particular beliefs that bear *a posteriori* perceptual justification. This leaves space for the non-naturalist to accept something based on Dickie's epistemic theory of reference, where what justifies a belief fixes the reference of the concepts employed in the content of the belief.⁶⁵ So, the *a priori* justification of those initial general beliefs will serve to fix the reference to the normative.⁶⁶ Further, I believe the appropriate reference-fixing justification is the sort shown in deliberative indispensability arguments.⁶⁷ These are of course controversial, but anyways attractive to non-naturalists. My thought is that the indispensable projects of practical and theoretical deliberation justifies agents in believing in normative reasons, or that there are considerations that *count in favor of* actions and beliefs. Per the epistemic theory, this would fix the reference of the employed concept to the genuine relation of "counting in favor of," and as all other normative properties will bear some relation or other to reasons, this reference foothold will enable successful reference for the

⁶⁴ See Werner (2020) for details.

⁶⁵ Dickie (2016). Werner (2020) offers something in the neighborhood: "the rough idea is that a concept C refers to some property F just in case the mode of justification for beliefs containing C non-luckily converge on F."

⁶⁶ This again involves the dialectical assumption that we have justified true beliefs about the normative. If this assumption proves false, then neither the debunking challenge nor its solution is relevant, because global normative skepticism is true anyway.

⁶⁷ See Enoch (2011, ch. 3).

appropriately related corresponding normative concepts. The refined details of this account unfortunately must wait for another paper.

References

- Bader, R. 2017. "The Grounding Argument against Non-Reductive Moral Realism." *Oxford Studies in Metaethics* 12: 105-134
- Bedke, M. 2014. "No Coincidence?" In *Oxford Studies in Metaethics*. 9: 102-125.
- Bengson, J. 2015. "Grasping the Third Realm." In *Oxford Studies in Epistemology*, 5:1–38. Oxford University Press.
- Burnston, D. Forthcoming. "Perceptual Learning, Categorical Perception, and Cognitive Permeation." *Dialectica*. <https://philpapers.org/rec/BURPLC>
- Buckner, C. 2018. "Empiricism Without Magic: Transformational Abstraction in Deep Convolutional Neural Networks." *Synthese* 195, 5339–5372.
- Clarke-Doane, J. and D. Baras. 2021. "Modal Security." *Philosophy and Phenomenological Research* 102: 162-183.
- Cop, D. 2019. "How to Avoid Begging the Question Against Evolutionary Debunking Arguments." *Ratio* 32: 231-245.
- Cowan, R. 2015. Cognitive penetrability and ethical perception. *Review of Psychology and Psychology*, 6(4), 665– 682.
- Dickie, I. 2016. "The Essential Connection Between Epistemology and the Theory of Reference." *Philosophical Issues*, 26(1): 99-129.
- Elliott, A. and D. Faraci. Unpublished Manuscript. "Non-naturalism and the 'Third-Factor Gambit'."
- Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- . 2019. "How Principles Ground." *Oxford Studies in Metaethics* 14: 1-
- Faraci, D. 2019(a). "Moral Perception and the Reliability Challenge." *Journal of Moral Philosophy* 16 (1): 63-73.
- . 2019(b). "Groundwork for an Explanationist Account of Epistemic Coincidence." *Philosophers' Imprint* 19 (4)
- FitzPatrick, W. J. 2010. "Robust Ethical Realism, Non-naturalism, and Normativity." *Oxford Studies in Metaethics* 3: 159-205.
- Heathwood, C. 2012. Could Morality Have a Source? *Journal of Ethics and Social Philosophy* 6 (2):1-19.
- Ichikawa, J. and M. Steup. 2018. "The Analysis of Knowledge." *Stanford Encyclopedia of Philosophy*. E. N. Zalta (ed.). <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>
- Joyce, R. 2006. *The Evolution of Morality*. MIT Press.
- Korman, D. 2014. Debunking Perceptual Beliefs about Ordinary Objects." *Philosophers' Imprint* 14 (13): 1-21.
- Korman, D. and Locke, D. (2020). "Against Minimalist Responses to Moral Debunking Arguments." In *Oxford Studies in Metaethics*, 15: 309-332

- . Forthcoming. An Explanationist Account of Genealogical Defeat. *Philosophy and Phenomenological Research*. Online first: DOI: 10.1111/phpr.12848.
- Leary, S. 2017. “Non-naturalism and Normative Necessities.” *Oxford Studies in Metaethics* 12: 76-105.
- McPherson, T. 2015. “What Is at Stake in Debates between Normative Realists?” *Nous* 49: 123-146.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press.
- Siegel, S. 2006. Which properties are represented in perception? In T. S. Gendler, & J. Hawthorne (Eds.), *Perceptual experience*. Oxford University Press.
- . 2007. “How can we discover the contents of experience?” *Southern Journal of Philosophy*, 45(S1), 127– 142.
- . 2010. *The Contents of Visual Experience*. Oxford University Press.
- . 2017. *The Rationality of Perception*. Oxford University Press.
- Skarsaune, K.O. 2011. “Darwin and Moral Realism: Survival of the Fittest.” *Philosophical Studies* 152 (2): 229–243.
- . 2015. “How to Be a Moral Platonist” Oxford: *Oxford Studies in Metaethics* 10: 245-272.
- Stokes, D. 2013. “Cognitive Penetrability of Perception.” *Philosophy Compass* 8 (7): 646-663.
- Street, S. 2006. “A Darwinian Dilemma for Realist Theories of Value.” *Philosophical Studies* 127 (1): 109–166.
- Vavova, K. 2014. “Debunking Evolutionary Debunking.” In *Oxford Studies in Metaethics* 9: 76–101.
- . 2015. “Evolutionary Debunking of Moral Realism.” *Philosophy Compass* 10 (2): 104-116.
- Werner, P. J. 2016. “Moral Perception Without (Prior) Moral Knowledge.” *Journal of Moral Philosophy* 13 (3) 294-317.
- . 2017. “A Posteriori Ethical Intuitionism and the Problem of Cognitive Penetrability.” *European Journal of Philosophy* 25 (4): 1791-1809.
- . 2020. “Getting a Moral Thing into a Thought.” In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, (15).
- Wielenberg, E. J. 2010. “On the Evolutionary Debunking of Morality.” *Ethics* 120 (3): 441–464.
- Zagzebski, L. “The Inescapability of Gettier Problems.” *The Philosophical Quarterly* 44 (17): 65-73.