

Knowing the Non-natural: Why Morality is as it Appears

By

Daniel Crow

A dissertation submitted in partial fulfillment of

The requirements for the degree of

Doctor of Philosophy

(Philosophy)

at the

University of Wisconsin-Madison

2014

Date of final oral examination: 5/6/14

The dissertation is approved by the following members of the Final Oral Committee:

Russ Shafer-Landau, Professor, Philosophy

Keith Yandell, Professor Emeritus, Philosophy

John Bengson, Assistant Professor, Philosophy

Michael Titelbaum, Assistant Professor, Philosophy

Rob Streiffer, Associate Professor, Medical History and Bioethics

Acknowledgements

I am grateful for the many people who have supported me during this dissertation-writing marathon. My advisor Russ Shafer-Landau has provided extensive commentary on multiple drafts of each of these chapters. He has been a model of excellence within the profession and human flourishing more broadly. Thank you, Russ. John Bengson has also provided essential guidance—both in what to read and in how to write and think. Thank you, John. Finally, many thanks to Keith Yandell for supporting me at every stage of my graduate school career.

And thank you to everyone else.

Introduction

These three essays are about the epistemology of ethical non-naturalism (henceforth, simply “non-naturalism”). Broadly, my aim is to vindicate non-naturalism against epistemological criticism. Before previewing each of my essays, I will clearly characterize non-naturalism and ascertain its unique epistemological challenges.

a. Non-Naturalist moral metaphysics

Non-naturalism is a species of moral realism. And moral realism is a thesis in moral metaphysics that contrasts most directly with moral constructivism (henceforth just “realism” and “constructivism”). To introduce realism, it will be helpful to compare its account of a paradigmatically wrong action with a naïve constructivist account.

Some hoodlums set a cat on fire for fun. Realists and constructivists can agree that this action is wrong. Furthermore, they can agree that the action is wrong “in virtue of” some of its non-moral base properties. Non-moral properties of this action that are explanatorily relevant to its wrongness include mental properties such as the hoodlums’ cruel intentions and the cat’s agony and possibly some physical properties such as the chemical properties of the gasoline that the hoodlums use to ignite the cat. Realists and constructivists agree that the action is wrong “in virtue of” base properties such as these.

I use the phrase “in virtue of” to denote an asymmetrical explanatory relation while remaining neutral on exactly what kind of explanatory relation it is. To avoid awkward “in virtue of” locutions, though, I will refer to the relevant explanatory relation as a grounding relation. Thus, I will say that realists and constructivists agree that the moral properties of an action are grounded in its non-moral properties. Since some of the grounding properties are mental properties (e.g., the hoodlums’ intentions, the cat’s pain),

realists and constructivists agree that moral properties are mind-dependent in the sense that they are grounded in a set of base properties that includes mental properties.

While realists and constructivists agree that moral properties are mind-dependent, in this sense, they disagree over whether they are “stance-independent.”¹ According to realism, moral properties are stance-independent in the sense that whether or not a certain kind of action grounds a moral property does not depend on any (actual or hypothetical) intentional attitudes taken towards that action. On this view, the fact that setting cats on fire for fun is wrong does not depend on anyone’s intentional attitudes towards harming cats. According to constructivism, by contrast, the wrongness of this action does depend on some such intentional attitude. According to the simplest kind of constructivism, the moral properties of this action depend on a social group’s approval or disapproval. That the social group disapproves harming cats is what makes the action ground wrongness.

When I say that moral properties “depend on” intentional attitudes, in characterizing constructivism, the phrase “depends on,” like the earlier “in virtue of” locution, denotes an asymmetrical explanatory relation.² This “depends on” relation is plausibly understood as a playing an enabling role to the grounding relation mentioned above; the attitudes enable the non-moral properties to ground the moral property. According to moral realism, by contrast, the relevant grounding relation does not require enabling by any intentional attitudes. Instead, it is a brute metaphysical fact that certain sets of non-moral properties ground certain moral properties.

Non-naturalism qualifies the realist’s thesis that moral properties are stance-independent with the negative proposition that they are not natural. To fully understand

¹ Russ Shafer-Landau, *Moral Realism* (Oxford: Oxford University Press), p. 15.

² I am counting enabling relations as kinds of explanatory relations.

this qualification, we would need to motivate a general distinction between natural and non-natural properties. One popular way of drawing this distinction has it that a property is natural just in case it is part of the proper subject matter of the natural and social sciences. A property is non-natural, then, just in case it falls outside of the jurisdiction of these sciences. In this schema, a bunch of sciences are lumped together under the heading of “natural” while the subject matter of ethics is set apart as “non-natural.” The division is unsatisfying unless supplemented with an account of the intrinsic differences between the subject matter of ethics, and the subject matter of everything else, that justifies excluding the former from a genus that comprises all of the latter.

There are two ways that moral properties are alleged to be importantly different from all of these other subject matters. The first difference I can only gesture at by saying that moral properties have a “normative dimension” that the subject matters of the other sciences apparently do not. Imagine a psychologist who proves that a common parenting style discourages the development of empathy in children. What the psychologist has discovered is a certain causal relationship. Parenting style X causes (or fails to cause) development of trait Y. This discovery has obvious normative implications. It would be *prima facie* wrong to raise one’s children in this way. Even though the discovery has clear normative implications, though, it is plausibly the causal relationship and not its normative implications that the psychologist studies. By contrast, the proper subject matter of ethics is specifically this kind of normative property.

While realists of all stripes agree that moral properties are normative, most non-naturalists claim that they are special in one other way: namely, they are causally impotent. According to many naturalist realists, moral properties are part of the best

causal explanation of a wide array of non-moral events. According to the naturalist Nicholas Sturgeon, for example, Hitler's depravity is part of the explanation of the Holocaust. By contrast, non-naturalists deny that moral properties play this explanatory role. They claim that another way that moral properties differ from the subject matter of the other sciences is that they do not figure in the best explanations of non-moral events. If "depravity" denotes a moral property that is grounded in Hitler's psychological properties (such as his hatred), it is these subvening psychological properties, and not the supervening moral property, that actually do the explanatory work.

To summarize: According to non-naturalism moral properties are stance-independent and are not natural properties. They are not natural in the sense that they are not part of the subject matter of the natural or social sciences. And they are essentially different from these other subject matters in that they have a special normative dimension that the other subject matters lack and lack causal powers the other subject matters have.

b. Non-naturalist moral epistemology: the background

The metaphysical commitments of non-naturalism constrain the range of promising models for non-naturalist moral epistemology. If moral properties depended on the intentional attitudes of social groups (as is claimed in naïve constructivism), then we should be able to learn about the instantiations of moral properties by discerning the relevant attitudes of the members of the group. Moral epistemology would be as simple as mindreading. The commitment of non-naturalism to stance-independence, however, implies that mindreading is an unpromising model for moral epistemology. We do not gain moral knowledge by discerning others' intentional attitudes to actions but rather (if at all) by discerning grounding relations that obtain independently of these attitudes.

If moral properties were causally potent, then there would be other epistemological possibilities. If moral properties figured in the best explanations of a wide array of non-moral events—as the naturalist claims—then we could come to know about the instantiation of a moral property via an inference to the best explanation. For example, we could know that Hitler was depraved via an inference to the best explanation of the occurrence of the Holocaust. Since non-naturalists deny that moral properties play this explanatory role, though, they must reject this scientific model for moral epistemology.

The non-naturalist's commitment to the causal impotence of moral facts also seems to exclude perception as a viable epistemological model. Certainly many of the physical properties that we perceive are stance-independent in the same sense that non-natural moral facts are; just as our attitudes do not determine what is right and wrong (according to moral realists), so they do not determine the physical property of (e.g.) being a tree. Given that physical properties and moral properties are both stance-independent, one might have thought that we perceive moral properties in the same way that we perceive physical properties.

Perception, however, features a causal connection between the perceiver and the perceived. You can see trees only because their surface properties reflect light rays that strike your sense organs. While there have been a flurry of recent attempts to develop accounts of non-naturalist moral perception—much more on these below—it is worth noting now that all such accounts assume a heavy burden: to explain how we perceive moral properties given that they do not causally impact any of our sense organs.

The commitments of non-naturalist moral metaphysics, then, limit the options for non-naturalist moral epistemology. They apparently exclude mindreading, inference to the best explanation, and perception as viable epistemological models.

One epistemological model that apparently is not excluded by the metaphysical commitments of non-naturalism is what I call *Moral Rationalism*. According to Moral Rationalism, all of our moral knowledge has a source of a priori justification. As a first approximation: a belief is justified a priori if it is justified on the basis of pure thought. The justified belief that hoodlums are hurting a cat must be a posteriori (via perceptual experience). According to Moral Rationalism, though, our justification for believing that this action instantiates the property of wrongness is based on pure thought.

Since many of the moral propositions we know are synthetic, Moral Rationalism implies that much moral knowledge is a kind of synthetic a priori knowledge. Ever since Kant clarified the category of the synthetic a priori, empiricists have claimed that such knowledge is objectionably mysterious. Because it seems to exclude every epistemological model except for Moral Rationalism, non-naturalist moral metaphysics seems to require an empiricist-unfriendly moral epistemology.

In my first paper, “The Mystery of Moral Perception,” I venture an interpretation of why empiricists are uncomfortable with synthetic a priori knowledge. This interpretation also plays an important role in my second paper, “The Causal Objection to Ethical Non-naturalism,” and thus is worth foregrounding here.

According to my interpretation—which is inspired by the work of Paul Benacerraf and Hartry Field—synthetic a priori knowledge is objectionably mysterious in the sense that (allegedly) there is no adequate explanation for the striking fact that we are

as reliable as we are in forming substantive a priori beliefs. With perceptual knowledge, we have a detailed, scientifically credible, causal explanation for the reliability of the perceptual beliefs that compose this knowledge.³ With synthetic a priori knowledge, however, there is nothing of the sort. According to empiricists, there is no adequate non-causal explanation either. Thus, the reliability of our a priori beliefs is inexplicable.

c. Non-naturalist moral epistemology: my contribution

In the “Mystery of Moral Perception,” I argue that many recent accounts of moral perception, which have been advertised as radical empiricist-friendly alternatives to Moral Rationalism, are in fact not substantially different from it. First of all, I argue, each of these accounts implies a source of a priori justification (or at best, a source of justification that is not clearly a posteriori). And, secondly, I argue that each implies that perceptual moral knowledge is as mysterious as synthetic a priori knowledge.

Accounts of non-naturalist moral perception must collapse into Moral Rationalism, in this way, in order to avoid the charge of triviality. Trivially, we perceive some of the non-natural properties that ground moral properties. To avoid the charge of triviality, then, an account of moral perception must claim more than this. Either the account needs to claim that moral perception involves an unconscious inference, the distinct representation of the moral property in perceptual experience, or a moment of perceptual recognition. I argue that, whichever of these sophisticated features is chosen to save the account of moral perception from triviality, that feature smuggles in a priori justification and also renders the entire act of moral perception mysterious.

³ To avoid awkward locutions, I refer to the reliability of beliefs even though reliability is primarily a property of belief-forming processes.

Having eliminated accounts of moral perception as an epistemological model that is distinct from Moral Rationalism, I turn to develop Moral Rationalism in the context of responding to the Causal Objection. In the second essay, “The Causal Objection to Ethical Non-Naturalism,” I begin by distinguishing the Causal Objection from the closely related Evolutionary Objection. Each of these critiques claims that some feature of the genealogy of our moral beliefs leads to some kind of epistemological crisis if non-naturalism is true. The alleged crisis is something in the ballpark of the proposition that we have no moral knowledge.

To flesh out the Causal Objection, I draw again from the work of Benacerraf and Field. In the previous paper, I unpacked the charge that synthetic a priori knowledge is “mysterious” as an inexplicability problem: there is no adequate explanation for the reliability of our synthetic priori beliefs. In this paper, I understand the Causal Objection in similar terms, although I do try to strengthen the relevant explanandum in one way. If non-naturalists have trouble explaining the *reliability* of our moral beliefs, they have even more trouble explaining the striking fact that many of our moral beliefs are *non-accidentally true*—where “non-accidental” here refers to the fourth condition that, together with belief and justification and truth, constitutes knowledge.

If moral properties cannot be part of the cause of any moral belief—and this is guaranteed by non-naturalist moral metaphysics—then there can be no adequate explanation for the striking fact that many of our moral beliefs are non-accidentally true. Any moral metaphysics that offers no adequate explanation of this striking fact is implausible. Therefore, non-naturalism is implausible. This is how I reconstruct the Causal Objection to Ethical Non-naturalism.

I respond to this argument in two different ways. First of all, I argue that it generalizes over all synthetic a priori knowledge. For reasons that should be clear by now, I think that the same kind of inexplicability worry that is driving this epistemological critique of non-naturalism is also driving the more general empiricist resistance to synthetic a priori knowledge. And if the former worry is merely a narrower application of the latter, then we should expect that an adequate explanation of synthetic a priori knowledge would imply an adequate explanation of moral knowledge as well.

The second way I respond to this argument is by searching for the adequate explanation that is needed to save both rationalism and non-naturalism. I consider a few putative explanations that fail to do so. From these failures, I realize that an adequate explanation of the relevant knowledge must propose an explanatory relation (or series of explanatory relations) that ultimately relates the subject of the knowledge with the relevant facts. I thus endorse the *Knower-Fact Constraint*. In the context of non-naturalist moral epistemology, the Knower-Fact Constraint implies that an explanation (for the striking fact that many of our moral beliefs are non-accidentally true) is adequate only if it relates the moral knower to some relevant moral fact.

Having identified this explanatory constraint, I proceed to consider only putative explanations that satisfy it. Evolutionary explanations and John Bengson's constitutive explanations all satisfy the constraint. Nevertheless, I argue that each of these putative explanations is inadequate. Finally, I develop an account of Intellectual Perception that, I argue, does ground an adequate explanation of the striking fact that our synthetic a priori beliefs (including our moral beliefs) are non-accidentally true.

According to my account of Intellectual Perception, appropriate reflection enables relevant abstracta to cause our intuitions; beliefs based on these intuitions, I argue, would constitute knowledge. To provide *prima facie* motivation for the radical thesis that abstract entities cause our mental states, I show how my account makes sense of an otherwise puzzling (non-epistemic) feature of human psychology.

In my third essay, “A Plantingian Pickle for a Darwinian Dilemma,” I commence the delayed discussion of the Evolutionary Objection. Sharon Street formulates the Evolutionary Objection as a dilemma for normative realism. Either the Darwinian forces that shaped our normative beliefs tracked the truth of realistically construed normative facts or they did not. According to Street, the former horn of the dilemma leads to bad science and the latter horn to a crippling epistemological problem.

In response to Street, I devise a dilemma for the Darwinian Dilemma for normative realism. (To distinguish the two dilemmas, I refer to the second-order one as a “pickle.”) To set up the pickle, I uncover deep similarities between Street’s Darwinian Dilemma for Normative Realism and Alvin Plantinga’s Evolutionary Argument Against Atheism. In his evolutionary argument, Alvin Plantinga appeals to similar evolutionary considerations as Street to support the conclusion that reflective belief in atheism is irrational. The similarities between Street’s and Plantinga’s evolutionary arguments certainly limit the objections to Plantinga’s argument that Street can safely endorse. By proposing a pickle, I argue that any basis for rejecting Plantinga’s evolutionary argument implies some damage to Street’s structurally similar argument.

The pickle is this. In rejecting Plantinga’s argument, Street has to make a claim that is either self-defeating—in the sense that it obliges Street to reject a parallel claim of

her own argument—or elitist—in that it implies a state of diminished rationality for reflective novices. Either implication counts against Street's argument.

Since Street cannot reject Plantinga's argument without damaging her own, I conclude that the following conditional is plausible: if the Darwinian Dilemma for Normative Realism is sound, then so is the Evolutionary Argument Against Atheism. Thus, the Darwinian Dilemma ultimately supports theism. But this consequence is doubly problematic for Street. Atheism is a suppressed premise of the Darwinian Dilemma as well as a commitment of almost all normative anti-realists. If the argument supports theism, it is internally incoherent and should be abandoned by almost everyone. Thus, this influential formulation of the Evolutionary Objection fails.

By defending non-naturalism against two powerful epistemological critiques, and exposing accounts of non-naturalist moral perception as disguised versions of Moral Rationalism, and developing a version of Moral Rationalism that clearly is hospitable to non-naturalism, I contribute to the grand project of non-naturalist moral epistemology.

Chapter 1: The Mystery of Moral Perception

1. Introduction

We make many moral judgments spontaneously and naturally use perceptual terms to describe them. If you were to encounter a group of hoodlums torturing a cat, you would be able to just “see” that the action is wrong. In the last decade, some moral realists have defended a kind of moral epistemology that takes this perceptual language literally. They argue that we literally perceive moral facts.¹

It may seem surprising that these accounts of moral perception have been hospitable to non-naturalism.² A species of moral realism, non-naturalism claims that moral facts are not part of the fabric of the natural world. They have normative powers that natural facts lack—such as giving categorical reasons—and lack powers that most natural facts have—for example, they are causally impotent.³ Because of this commitment to a *sui generis* realm of moral facts, non-naturalism is often criticized on the grounds that it requires an objectionably unscientific or mysterious epistemology.⁴

¹See Robert Audi, *Moral Perception* (Princeton: Princeton University Press, 2013); Robert Cowan, “Perceptual Intuitionism,” *Philosophy and Phenomenological Research* 86 (2013), pp. 1-30; Andrew Cullison, “Moral Perception,” *European Journal of Philosophy* 2 (2010), pp. 159-175; Terence Cuneo, “Reidian Moral Perception,” *Canadian Journal of Philosophy* 33 (2002), pp. 229-258; Justin McBrayer, “Moral Perception and the Causal Objection,” *Ratio* 23 (2010), pp. 159-175, and “A Limited Defense of Moral Perception,” *Philosophical Studies* 149 (2010), pp. 305-320; Sarah McGrath, “Moral Knowledge by Perception,” *Philosophical Perspectives* 18 (2004), pp. 209-28; and Jennifer Wright, “The Role of Moral Perception in Mature Agency,” in J. Winewski (ed.), *Moral Perception* (Cambridge, MA: Cambridge Scholars Publishing, 2008), pp. 1-24.

² For example, Cuneo, “Reidian Moral Perception,” pp. 256-257; Audi, *Moral Perception*, pp. 44, 55; and McBrayer, “Moral Perception and the Causal Objection,” pp. 300-306, each suggests that his account of moral perception can accommodate non-naturalist moral metaphysics.

³ Thus, I assume that non-natural moral properties are causally impotent. I do not here endorse the claim that causal impotence is a sufficient condition for being non-natural.

⁴ A.J. Ayer, *Language, Truth, and Logic* (New York: Dover Publications, 1952), p. 106; J.L. Mackie, *Ethics: Inventing Right and Wrong*, (Harmondsworth: Penguin, 1977), pp. 38-42; and Michael Smith, *The*

Traditionally, non-naturalists have sought refuge from these allegations in the rationalist epistemological tradition. They have tried to show that a general epistemological rationalism—roughly, the thesis that we have substantive a priori knowledge—can be extended, without any problem, into even a non-natural moral domain.⁵ The moral epistemology that results is a version of moral rationalism.⁶

This strategy for developing a non-naturalist moral epistemology is promising to the extent that the underlying rationalism is plausible. But rationalism is controversial.⁷ Many philosophers believe that rationalism—in its general as well as particular applications—is objectionably unscientific or mysterious, too. If rationalism is implausible, then we cannot vindicate any moral epistemology by showing it to be a straightforward extension of rationalism into the moral domain.

Against this background, accounts of moral perception have been advertised as an empiricist-friendly alternative to moral rationalism.⁸ They promise to introduce moral

Moral Problem (Oxford: Blackwell Publishing, 1994), pp. 18-25, emphasize the mysteriousness of non-naturalist moral epistemology. Sharon Street, “A Darwinian Dilemma for Realist Theories of Value,” *Philosophical Studies* 127 (2006), 109-166, at p. 112, and Matthew Bedke, “Intuitive Non-Naturalism Meets Cosmic Coincidence,” *Pacific Philosophical Quarterly* 90 (2009), 188-209, at p. 190, rely on allegedly scientifically established premises in their epistemological critiques of moral realism and non-naturalism. While Street supposes that a particular account of the evolution of our normative judgments is true, Bedke includes, in his Master Argument, the premise that the physical is causally closed—which, he asserts, “is supported daily by advances in causal explanations of the various sciences.”

⁵ Michael Huemer, *Ethical Intuitionism* (New York: Palgrave Macmillan, 2008), pp. 122-127.

⁶ In contemporary metaethics, the term “moral rationalism” usually refers to various theses regarding practical moral reasons. Henceforth, however, I use “moral rationalism” to refer exclusively to the thesis that there is substantive a priori (moral) knowledge.

⁷ For an authoritative defense of rationalism, see Laurence Bonjour, *In Defense of Pure Reason* (Cambridge: Cambridge University Press 1998). For a collection of contemporary essays defending empiricism, see Michael Shaffer and Michael L. Veber, *What Place for the A Priori?* (Chicago: Open Court, 2011).

⁸ See, for examples, McBrayer, “A Limited Defense of Moral Perception,” p. 306, and Cowan, “Perceptual Intuitionism,” p. 1. Cp. Audi, *Moral Perception*, pp. 2-3.

realism—even non-naturalism—back into the empiricist mainstream. Perception has a pristine reputation in empiricist and rationalist circles. Indeed, it is the paradigm of a scientifically credible cognitive process. If non-naturalist moral epistemology can wear the reputable cloak of perception, then it can be absolved of the longstanding mysteriousness charge and break free from its questionable association with an epistemological tradition that flaunts substantive a priori knowledge.

In this essay, I argue that accounts of non-naturalist moral perception have failed to live up to their promise to provide an interesting alternative to moral rationalism. Necessarily, these accounts conceal a core commitment of rationalism and embody its most objectionable feature. In particular, accounts of non-naturalist moral perception conceal a commitment to substantive a priori justification and make moral perception “mysterious” in the sense that they provide no explanation of its reliability.

I proceed as follows. In Part II, I clarify the disagreement between rationalists and empiricists. I parse the dispute into separate disagreements about justification and knowledge. Following Benacerraf and Field, I develop the charge that substantive a priori knowledge—and by extension, rationalism—is objectionably “mysterious.” In Part III, I distinguish three accounts of moral perception and argue that each smuggles in a priori justification. In Part IV, I argue that non-naturalist moral perception is on par with substantive a priori knowledge with respect to mysteriousness.

II. Rationalism and Empiricism

I understand the difference between rationalism and (moderate⁹) empiricism as a disagreement about the scope of a priori knowledge. According to empiricists, all a priori knowledge is, in some sense, trivial. In contrast, rationalists claim that much a priori knowledge is substantive. While I will rely on our loose sense of the distinction between trivial and substantive propositions, I will need to examine more closely the distinction between a posteriori and a priori justification.

As a first approximation: a belief is justified a posteriori if it is justified on the basis of experience; it is justified a priori if it is justified independently (i.e., not on the basis) of experience. This way of drawing the distinction is asymmetrical, however, because it identifies the source of a posteriori justification but tells us only what the source of a priori justification is not. An adequate account of a priori justification should supplement this negative account with something positive. For now, I will specify minimally that the source of a priori justification is “pure thought.”¹⁰ According to rationalism, then, we can be justified in believing substantive propositions on the basis of pure thought.

Rationalists claim not only that many beliefs in substantive propositions are justified a priori, but also that many of these beliefs qualify as knowledge. In reply, many empiricists object that substantive a priori knowledge is objectionably

⁹ I adopt the distinction between moderate empiricism and radical empiricism from Bonjour, *In Defense of Pure Reason*, pp. 18-19. According to moderate empiricism, some beliefs are justified a priori—but all of these are analytic. According to radical empiricism, no beliefs are justified a priori. Radical empiricism may also be understood to include a commitment to the thesis that true analytic propositions are not made true by any feature of the mind-independent world. To avoid ambiguities in the analytic/synthetic distinction, I favor a distinction between the trivial and substantive.

¹⁰ For rigorous attempts to refine the a priori/a posteriori distinction, see Bonjour, *In Defense of Pure Reason*, Ch. 1, and Albert Casullo, *A Priori Justification* (Oxford: Oxford University Press, 2003), Ch. 1-2.

“mysterious.” But in what sense is such knowledge mysterious? Following Benacerraf and Field, I will unpack the mysteriousness charge in terms of inexplicability.¹¹ In particular, I will understand the charge as follows: substantive a priori knowledge is objectionably mysterious in the sense that (allegedly) there is no adequate explanation for the striking fact that we are as reliable as we are in forming substantive a priori beliefs.¹²

The spirit of the Benacerraf-Field challenge is best conveyed by contrasting substantive a priori knowledge with perceptual knowledge. For many of our perceptual beliefs, we have a detailed, scientifically credible, mechanistic explanation of their reliability. The general shape of the explanation is familiar to all of us: human beings have a complex cognitive system that converts causal stimuli into perceptual representations. For example, when you grab a tomato, the roundness of the tomato causally impacts the nerves of your hand and—following a very complicated cognitive process—is then represented in your tactile experience. This causal connection between the perceiver and the perceived explains the fact that our perceptual beliefs are reliable.¹³

With respect to the beliefs we form on the basis of pure thought, however, there is no similar causal explanation for the striking fact of their reliability. This is most obvious when we consider a causally impotent domain such as mathematics (assuming

¹¹ See Paul Benacerraf, “Mathematical Truth,” *The Journal of Philosophy* 70 (1973), pp. 661-679, and Hartry Field, *Realism, Mathematics, and Modality* (Oxford: Basil Blackwell, 1989).

¹² I follow Joshua Schechter, “The Reliability Challenge and the Epistemology of Logic,” *Philosophical Perspectives* 24 (2010), pp. 437-464, in formulating the Benacerraf-Field challenge in terms of reliability. Note that this is controversial. John Bengson, “Grasping the Third Realm,” forthcoming in T. Szabo Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology, Vol. 5*, argues that Benacerraf’s challenge is to identify a relation between intuitions and abstracta that explains why the former are non-accidentally true. He contends that this explanatory challenge is “prior to” the reliability challenge (see pp. 3, 41-43).

¹³ Even though reliability is primarily a property of representation-forming processes (a genus term that encompasses both belief-formations and the formation of perceptual content), I will sometimes use the term to refer to the beliefs or perceptual representations formed via a reliable process.

mathematical Platonism) or morality (assuming moral non-naturalism). Plausibly, we are able to know many mathematical and moral principles just by thinking about them. But no causal interaction with abstract mathematical entities or non-natural moral facts can explain the reliability of these beliefs.

The Benacerraf-Field challenge extends even further—to our a priori beliefs about domains that are arguably causally potent. For the reliability of these beliefs, there is also no causal explanation.¹⁴ Consider the necessary truth that nothing can be both red and green all over at the same time. We are able to know that this is true just by thinking about it. To be sure, I causally interact with particular green and red things frequently. But these causal connections do not seem to explain the reliability of the process by which I form this kind of belief. In case this is not obvious, note that I could still know this proposition even if I were to learn that I am a brain in a vat—causally isolated, we may assume, from every particular green and red thing.¹⁵ That a brain in a vat could reliably form beliefs about necessary truths suggests that causal connections with concrete particulars is not, even in normal circumstances, part of the explanation of the reliability of these a priori beliefs.¹⁶ Generalized over all substantive a priori knowledge,

¹⁴ Bonjour, *In Defense of Pure Reason*, p. 156, claims that the “objection can be extended to rationalist *a priori* knowledge and justification generally, and [...] has frequently generally been invoked in this broader form by others.” Joshua Schechter, “The Reliability Challenge and the Epistemology of Logic,” on pp. 441-443, discusses the scope of the objection. Justin Clarke-Doane, “What is the Benacerraf Problem?,” in F. Pataut (ed.), *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity* (forthcoming), catalogues different attempts to generalize the objection to non-mathematical domains.

¹⁵ Bengson, “Grasping the Third Realm,” p. 14, offers the similar “case of Trip” to illustrate one kind of non-causal knowledge. In Bengson’s story, Trip has a hallucination in which he seems to see some colors and shapes that he has never encountered before. On the basis of this experience, he is able to know certain propositions about the relationships of these shapes and colors (e.g., that red is more like orange than it is like blue). In fn 27, Bengson discusses envatment in response to the objection that the Case of Trip is metaphysically impossible. Since a properly stimulated brain can experience features to which it bears no causal connections, it is not impossible that Trip could hallucinate new shapes and colors.

then, the Benacerraf-Field challenge, as I will understand it, states that substantive a priori knowledge is “mysterious” in the sense that there is no causal explanation for the reliability of substantive a priori beliefs.

Whether such “mysteriousness” is objectionable will depend on whether there is some other (presumably non-causal) adequate explanation for the reliability of our substantive a priori beliefs. Most rationalists will say that there is. Empiricists will disagree. On this matter, I do not intend to take a stand. For my aim here is not to defend either rationalism or empiricism, but rather to contest the suggestion that non-naturalist moral perception is empiricist-friendly. I will argue that if there is such perception, then there is no causal explanation for its reliability. If such mysteriousness is an objectionable property of a priori beliefs, it is also an objectionable feature of the perception of moral facts (and of the corresponding perceptual beliefs).

III. A Priori Justification in Moral Perception

A philosophically interesting account of moral perception must distinguish itself from the boring thesis that we perceive actions (or at least parts of actions) that in fact have moral properties—what Robert Audi calls “mere perception of a moral phenomenon.”¹⁷ We can perceive hoodlums hurting a cat. And we can perceive a nun helping a homeless person. Certainly, these actions instantiate moral properties. And,

¹⁶In order for a cognitive domain to fall within the scope of the Benacerraf-Field challenge, it is not enough that there is no causal explanation for the reliability of our beliefs about that domain. The domain must also satisfy other conditions. Schechter, “The Reliability Challenge and the Epistemology of Logic,” p. 439, proposes a plausible “objectivity” condition that encompasses (a) meaningfulness (the claim that statements about that domain are truth-apt), (b) independence (the truths do not “depend on us”), and (c) no plenitude (not every consistent set of beliefs or practices is equally correct).

¹⁷ Audi, *Moral Perception*, p. 31. See also McBrayer, “Moral Perception and the Causal Objection,” p. 293, and McBrayer, “A Limited Defense of Moral Perception,” p. 307.

certainly, we can perceive these kinds of actions. For an account of moral perception to be of any interest for moral epistemology, though, it needs to state more than this obvious fact. It should say not only that we perceive actions that are *in fact* wrong, but also *that* these actions are wrong—or perhaps perceive *wrongness* itself.

I see three different ways of distinguishing an account of moral perception from the boring thesis. Moral perception either involves (i) an inference from what is represented in perceptual experience to a corresponding moral belief, (ii) a special representation of a moral property in perceptual experience, or (iii) a recognition of one's perceptual experience as instantiating a moral property. In short: if moral perception is non-trivial, then it is inferential, representational, or recognitional. What I will argue now is that each kind of account of moral perception presupposes a priori justification.

A. The Inferential Account

I will begin by examining the least popular of the three accounts.¹⁸ According to an inferential account, moral perception is partly constituted by an inference from what is immediately represented in perceptual experience (*viz.*, some of the action's non-moral properties) to a conclusion that states more than what is represented (*viz.*, that the action

¹⁸ No one who defends moral perception clearly endorses an inferential account. But a few philosophers endorse the account according to one not unreasonable interpretation of their views. Audi, *Moral Perception*, pp. 52-53, draws a distinction between inferences and “a belief-formation that is a direct response to a recognized pattern”—and apparently allows that moral perception can involve either kind of cognitive process. Pekka Väyryen claims that his account of moral perception is non-inferential, but allows that its “causal etiology involves unconscious inference,” such as occurs in the following situation: “You turn to me at a rock concert and I hear you say ‘Awesome Solo!’ ... only because my brain merges auditory and facial movement signals into a unified experience [...] to repair degraded sounds and resolve ambiguities” (“Some Good News and Some Bad News for Ethical Intuitionism,” *Philosophical Quarterly* 58 (2008), pp. 489-511, at p. 498). Finally, Jennifer Wright distinguishes perceiving that an action is cruel and perceiving that the cruelty is reason-giving; the latter kind of perception is refined, she argues, in mature moral agents. According to one way of understanding this idea, the perception that cruelty is reason-giving involves an inference (Wright, “The Role of Moral Perception in Mature Agency,” at p. 9).

instantiates some moral property). For example, when you perceive that the hoodlums' action is wrong, your perception includes both (i) a representation of certain non-moral properties of the action (e.g., the hoodlums' malignant expressions, the cat's fiery tail), and (ii) an inference to the conclusion that this action is wrong. The inference is not a separate operation we perform on the deliverances of moral perception but is rather constitutive of moral perception itself.

What is an inference? Paradigmatically, an inference involves a series of mental states: antecedent beliefs in premises, an apprehension that the premises stand in some evidential connection, and the formation of a new belief in the conclusion on the basis of the premises.¹⁹ If moral perception involves a valid inference, then one of the premises of the inference must be a bridge-principle that links the non-moral and the moral. The inferential account of moral perception, then, implies that the moral perceiver has an antecedent belief in a bridge-principle. Plausibly, inferential moral perception is a source of epistemic justification for some moral belief only when the antecedent belief in the bridge-principle is also justified. The justification conferred by inferential moral perception draws from the source of justification of this antecedent belief.

If the antecedent belief in a bridge-principle is justified, then it is justified a priori or a posteriori. And it is not justified a posteriori. For if it were, then it would be justified either via perceptual experience or via induction. But it can be justified in neither way.

The antecedent belief is not justified via perceptual experience. To see why, it helps to remember the original motivation for the inferential account. We originally

¹⁹ In inferences, these mental states can occur consciously or unconsciously. To do justice to the spontaneity of our moral judgments, the proponent of inferential moral perception should say that the belief and apprehension that partly constitute "inferring" normally occur unconsciously in moral perception.

posited the perceptual inference to explain how moral perception could be a source of justification for a moral belief, provided that perceptual experience represents only the (boring) natural properties of actions. The inference is supposed to carry us beyond the perceptual experience to more interesting moral beliefs. If we now claim that the bridge-principle is justified via perceptual experience, then we introduce an objectionable kind of circularity into the view. We are invoking the justified belief in the bridge-principle to explain why perceptual experience is a source of epistemic justification for some moral belief, and we are invoking perceptual experience to explain the justification of the belief in the bridge-principle.

The claim that the antecedent moral belief is justified via an inductive generalization also introduces objectionable circularity into the inferential account. In general, the conclusion of an inductive generalization is justified only if the beliefs that constitute its inductive base are justified. For example, if we are justified in believing that torturing babies is wrong via an inductive generalization, then there must be a set of justified beliefs about particular cases of torturing babies—beliefs of the form *that instance of torturing babies was wrong*. But now we have to ask how these beliefs about particular cases are justified. The proponent of the inferential account cannot say: via perceptual experience. For the inductive generalization is supposed to explain the justification of the antecedent belief in the bridge-principle which, in turn, explains why inferential moral perception is a source of epistemic justification. To then claim that perceptual experience justifies the inductive base is, again, circular.

Perhaps the inductive base enjoys a non-perceptual source of justification. But what could this source possibly be? Whatever it is, it is either a source of a priori

justification or a posteriori justification. If it is a source of a priori justification, then we have introduced a priori justification into the inferential account of perception—which is just to concede that this kind of moral perception conceals a core commitment of rationalism. And there is no plausible non-perceptual source of a posteriori justification for the particular beliefs that compose the inductive base. Besides perception and induction, sources of a posteriori justification include introspection and kinesthesia. Clearly, these sources do not play a starring role in moral perception. Since it would be circular to claim that the inductive base is justified via perceptual experience, and there is no plausible non-perceptual source of a posteriori justification, the antecedent belief in the bridge-principle is not justified a posteriori via an inductive generalization.

Let us consider one other kind of inductive argument that is so different from inductive generalization that it is sometimes considered its own category: inference to the best explanation (i.e., abduction). If we are justified in believing moral principles via an inference to the best explanation, then there is some explanandum that the moral principles supposedly explain. Either the explanandum includes moral observations or it is purely non-moral. Suppose that the explanandum includes moral observations: for example, the observations that many different cases of torture are wrong. This explanandum obviously presupposes particular moral beliefs. The inference to some explanans is justified only if these particular beliefs are justified as well. Because the inference to the best explanation requires antecedent justified moral beliefs, the inferential account of moral perception that incorporates this abductive inference is circular in the same way as the kind that involves an inductive generalization.

Perhaps there is something that moral principles best explain that does not imply any moral facts: some purely non-moral explanandum. This possibility, however, is unavailable to the non-naturalist. While many naturalists claim that various moral facts explain different facts of the natural world—famously, Nicholas Sturgeon claims that Hitler’s depravity explains some of the events of the Holocaust—non-naturalism is usually understood (and is assumed here) to exclude this possibility.²⁰

Moreover, even if it were true that moral properties did explain various non-moral features of the world, this fact (or belief in this fact) is not plausibly a source of epistemic justification for all of our spontaneous moral judgments about particular cases. When a normal person just “sees” that the hoodlums are doing wrong, it is unlikely that his epistemic justification for the perceptual belief is derived from any inference to the best explanation of certain non-moral facts. Normal people should not be credited (or charged, as it were) with making such an inference.

Like inductive generalizations, inferences to the best explanation fail to provide an adequate a posteriori account of the justification of the antecedent beliefs that are essential to moral perception on this view. Since neither perceptual experience nor the most promising inductive sources can explain this justification, it is unlikely that such justification is derived from any posteriori source. If the antecedent belief is justified at all, then it is justified a priori. Therefore, if moral perception involves an inference, it implies a priori justification of an antecedent belief in a bridge principle. Thus, the inferential account conceals a core commitment of rationalism: to the a priori justification of substantive bridge-principles that specify links between the moral and the non-moral.

²⁰ Nicholas Sturgeon, “Moral Explanations,” in David Copp and David Zimmerman (eds.), *Morality, Reason, and Truth* (Totowa, NJ: Rowman & Allanheld, 1985), pp. 49-78.

B. Representational Accounts

Most defenders of moral perception, however, have denied that moral perception involves any inference at all. Proponents of non-inferential accounts of moral perception, then, must find a different way to distinguish their account from the boring thesis that we perceive actions that are in fact wrong. One way they have tried to do so is by building into their account a strong phenomenological constraint.

According to a representational account of moral perception, when we perceive instantiations of moral properties, the phenomenology of our perceptual experience represents not only some of the base properties that ground the moral property, but also the moral property itself. According to this account, perceptual experience represents not only actions that are in fact wrong but, additionally, wrongness itself.

Robert Audi and Robert Cowan have each developed representational accounts of moral perception.²¹ The striking similarities between their accounts—which, I will argue, are constraints on any remotely plausible representational account—raise the worry that these accounts do not supply a source of justification that is clearly a posteriori.

Let us note some of these similarities. In each of their accounts, perceptual experience has two phenomenological layers; the top layer is less “robust” (Cowan’s term) than the base layer, and is phenomenally dependent on it.²² For Audi, the base layer is the “perceptual” and the upper layer is the “perceptible;” for Cowan, the base layer is a “phenomenal presence,” which underlies an upper layer that involves

²¹ Audi, *Moral Perception*. Cowan, “Perceptual Intuitionism.”

²² Cowan, “Perceptual Intuitionism,” p. 7 (italics omitted).

“phenomenal presence as absence.”²³ In each account, properties of color and shape serve as paradigms of the base layer. In contrast, the representation of a moral property happens in the more rarified upper layer.²⁴

According to Audi, the representation of a moral property is “phenomenally integrated with” the “cartographic” representation of some of its perceptual base properties. Cowan claims that representations of moral properties are “added to” a “spatial framework” that is established by phenomenally present properties.²⁵ While Audi does not clearly identify any non-moral examples of the perceptible, one gathers that he would be amenable to the examples Cowan uses to illustrate phenomenological presence as absence. For an example, Cowan mentions the way that you can have an experience “as of the backside of a tomato”—even when all that is phenomenologically present is the side facing you.²⁶ In Audi’s more abstract discussion of the phenomenological upper layer, he suggests that the perceptible representation of the moral property is “intellective” and, when the property of “injustice” is represented, involves a “felt unfittingness between the deed and the context.”²⁷

Unlike Audi, Cowan claims that the representation of the moral property in perceptual experience is the result of “cognitive penetration.”²⁸ Cognitive penetration occurs when one of a person’s extra-perceptual mental states affects perceptual content.

²³ Audi, *Moral Perception*, p. 35. Cowan, “Perceptual Intuitionism,” p. 6 (italics omitted).

²⁴ Audi, *Moral Perception*, p. 35. Cowan, “Perceptual Intuitionism,” pp. 6-7.

²⁵ Audi, *Moral Perception*, pp. 38-39. Cowan, “Perceptual Intuitionism,” p. 10 (italics omitted).

²⁶ Cowan, “Perceptual Intuitionism,” p. 7.

²⁷ Audi, *Moral Perception*, pp. 33, 43.

²⁸ Cowan, “Perceptual Intuitionism,” p. 13.

For an example of cognitive penetration by a moral belief, we can imagine a person who believes *that torture is wrong* so firmly that, when he sees an actual case of it, the background belief causes a representation of wrongness to appear in his visual field.

I agree with Cowan's suggestion that, if the representation of the moral property in perceptual experience is the effect of cognitive penetration, then the justification conferred by this experience must be derived from the penetrating mental state.²⁹ For example, if a perceptual representation of wrongness is the effect of the background belief *that torture is wrong*, then the justification of the associated perceptual belief *this action is wrong* must be derived from the justification of the background belief.

An argument parallel to the one I developed in the previous section establishes that this kind of penetrating background belief, like the antecedent belief in inferential moral perception, must have a source of a priori justification. If this penetrating background belief is justified a posteriori, then it is justified via perceptual experience or induction. It is circular to claim that it is justified via perceptual experience, because the penetrating background belief is supposed to explain why perceptual experience is a source of justification for some moral belief. You can fill in the rest of the parallel argument.

Beliefs are mental states that transmit justification. Elaborating on the above example, we can picture the justification passing from its a priori source, to the background belief *that torture is wrong*, to the representation of the moral property in perceptual experience, and finally to the justified perceptual belief *that this action is wrong*. In the next section, I will consider cognitive penetration by moral memory, which implies a similar transmission of a priori justification. For now, though, we need

²⁹ See Cowan's discussion of "epistemic dependency" ("Perceptual Intuitionism," at pp. 4, 24-29).

to consider accounts of moral perception via cognitive penetration from mental states that do not transmit justification from a more basic source.

If the penetrating mental state does not transmit justification from a deeper source, then that mental state must be the (basic) source of justification for the perceptual moral belief. If the penetrating mental state is a source of justification, then either it is a source of a posteriori justification or a priori justification. But it cannot be a source of a posteriori justification, for reasons that have already been suggested. Sources of a posteriori justification include perceptual experiences, introspective states, and kinesthetic states. Again, these last two states should not be assigned a starring role in accounts of moral perception. And it would be bankrupt to appeal to the penetrating effects of perceptual experience to explain why perceptual experience is a source of justification. Whether the penetrating mental state is a transmitter or a source of justification, then, the representational account that incorporates a commitment to cognitive penetration implies a source of a priori justification.

Cognitive penetration, however, is a distinguishing feature of Cowan's account. Audi's account does not include this commitment. Nevertheless, it must smuggle in a priori justification. But there is a subtle difference in the nature of the a priori involvement. If the representation of the moral property is the result of cognitive penetration, then moral perception *has* a source of a priori justification. But if it is not the result of cognitive penetration, then it *is* a source of a priori justification—or, at best, is a source of justification that is not clearly a posteriori. The representation of the moral property has this confused epistemological identity because it fits the profile of an a priori intuition as well as it does the profile of an “upper layer” a posteriori experience.

To explain what I mean, I will need to problematize my symmetrical characterization of the distinction between a posteriori and a posteriori justification. To recall: a justified belief is a posteriori if it is justified on the basis of experience and a priori if it is justified on the basis of pure thought. This demarcation of the two concepts is inadequate for the fairly obvious reason that “pure thought” is experiential in some broad sense of the term. And some of these experiences arguably play a justifying (as opposed to a merely enabling) role in a priori justification. For example, when you contemplate the proposition $2+2=4$ —and you form a justified belief on the basis of pure thought—you have an intuition that this proposition is true. This intuition is an experience in the sense that it is a conscious state. And it is a justifier. To uphold the distinction between beliefs that are justified on the basis of experience and beliefs that are justified on the basis of pure thought in light of the awareness that pure thought is experiential in some broad sense of the term, we will need to draw a more fundamental distinction among kinds of justifying experiences: those that are sources of a posteriori justification (henceforth, a posteriori experiences) and those that are sources of a priori justification (henceforth, a priori experiences). And this is hard to do.

When we consider only paradigm cases, the distinction between a posteriori and a priori experiences seems obvious, and the request for a principle to justify the distinction might seem pedantic: in so many ways, seeing a red tomato is unlike intuiting that $2+2=4$. But paradigm cases of perceptual experience, which stand in such striking contrast to a priori intuition, are examples of the perceptual (Audi) or the phenomenally present (Cowan), i.e., the phenomenal base layer. While the phenomenal base layer of perception is strikingly different from a priori intuition, the upper layer is not. On purely

phenomenological grounds, it is difficult to distinguish the rarefied representation of phenomenal presence as absence from the subtle phenomenological flavors of an a priori intuition.

To avoid this confusion, a proponent of moral perception might attempt to relocate the representation of the moral property in the phenomenal base layer. But the account, so amended, is implausible. It is no coincidence that both Audi and Cowan locate the representation of the moral property in the upper layer. For it is obvious that the representation of the moral property in perceptual experience—if in fact it is represented at all—is utterly unlike the phenomenology that represents properties of space and color. If the phenomenology were similar, we would not be able to seriously question whether moral properties are represented in perceptual experience at all. (Nobody denies that we are “appeared to redly.”) This characteristic of Audi’s and Cowan’s account, which raises the suspicion that the representation of the moral property is an a priori intuition, is an essential feature of any plausible representational account.

If the thin phenomenology of the representation of the moral property raises the suspicion that it is an a priori intuition, its phenomenological dependence on the representation of base properties further arouses this suspicion. As W.D. Ross has noted, the order of discovery of a necessarily true general principle often involves a prior apprehension of the application of that principle in some contingent event.³⁰ For example, suppose that you were to observe an actual Gettier case: you see Jones form the justified true belief that it is 9 a.m. on the testimony of a broken clock that Jones has every reason to believe works. You would just be able to “see” that Jones does not know that it is 9 a.m. Obviously, your sensory experience justifies your beliefs about many of

³⁰ W.D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930), p. 33.

the contingent facts of this case. But rationalists will say—plausibly—that there is an a priori source to your justified belief that Jones does not have knowledge. When we apprehend necessary truths in particular cases, perceptual experience and a priori intuition are mixed together. Apparently, the a priori intuition exhibits the same kind of phenomenological dependence on the perceptual properties that, according to Audi and Cowan, the perceptual representation of the moral property has on the representation of its base properties.

Again, we find that this shared feature of Audi's and Cowan's account is a constraint on any remotely plausible representational account. It is obvious that the representation of the moral property (if it is represented at all) must exhibit a kind of phenomenological dependence on base properties. After all, we do not experience the representation of wrongness floating free from all other phenomenology—supervening, perhaps, in the corner of the basement next to a lost sock. An account of moral perception is plausible only if it claims that the phenomenological representation of the moral property is both subtler than, and dependent on, the phenomenological representation of some its base properties. But it is just these features of the perceptual experience that make it resemble a priori intuition. In order for an account of moral perception to succeed as an alternative to moral rationalism, its defenders need to do more to show that moral perception is significantly different from a priori intuition.

C. Perceiving-As

So far, we have looked at inferential and representational accounts of moral perception. We have found that the former relies on a priori justification and that the

latter invokes a kind of experience that either *has* or *is* a source of a priori justification. We will now consider one additional attempt to distinguish an account of moral perception from the boring thesis, by claiming that moral perception is a kind of perceiving-as.

In “Moral Perception and the Causal Objection,” Justin McBrayer suggests that any other kind of account of moral perception is trivial:

Relying on Dretske’s (1969) distinction between seeing and seeing as, we can say that all cases of perception are either cases of perception simpliciter or perception-as. “Perception as” requires identification. For example, upon seeing the university president for the first time, I perceived simpliciter the university president but failed to perceive him as the university president. [...] If some form of moral realism is true, it is obvious that we have moral perception simpliciter. We see actions that are, in fact, morally wrong. The contentious claim is that we might also have moral perception-as, e.g. that we might be able to see that an action was wrong, etc. I shall use “moral perception” to mean perception as if some moral property or other is instantiated.³¹

In this passage, McBrayer suggests that accounts of moral perception are distinct from the boring thesis only if they are accounts of perceiving-as. If we draw the distinction between perceiving and perceiving-as so that the latter encompasses both the inferential and representational account, then McBrayer is surely right. But this way of carving up logical space is crude, for there is a kind of perceptual identification that is arguably neither inferential nor representational. I will use “perceiving-as” in this narrower sense to denote perception that involves this kind of identification.

Such “identification,” I take it, is no different from the “kind of recognitional awareness” that Andrew Cullison independently claims is constitutive of moral perception.³² Cullison compares moral perception with a chicken-sexer’s ability to

³¹McBrayer, “Moral Perception and the Causal Objection,” p. 293. See Fred Dretske, *Seeing and Knowing* (London: Routledge and Kegan Paul, 1969), p. 9.

distinguish male from female chicks, with a park ranger's ability to see that a certain tree is a maple rather than a pine, and even with his own ability to identify a friend from a distance.³³ Aside from these suggestive analogies, Cullison does not offer any account of the kind of recognitional awareness that he takes to be essential to moral perception.

To understand how McBrayer's and Cullison's accounts of moral perception(-as) fit into the conceptual framework I have established, we need to better understand the kind of "recognition" (I will drop synonymous reference to "identification") that is distinctive of perceiving-as. On one plausible account, perceptual recognition involves matching the contents of one's current perceptual experience with some stored representation. I have a belief or memory that represents, in some semi-abstract fashion, certain distinguishing characteristics of the university president or a maple tree. If moral perception is like this, then I have a similar belief or memory that represents distinguishing characteristics of right and wrong actions. When the moment of recognition happens—when I suddenly see the person as the university president or the tree as a maple or the action as wrong—what has happened is that I have matched the contents of my current perceptual experience with the relevant stored representation.

Can a recognitional account of moral perception be reduced to either of the accounts we have considered? It cannot be reduced to the inferential account, although it may be similar in one respect. A fully developed recognitional account should specify the content of the semi-abstract representation against which one's current experience is matched. According to one such specification, the stored representation of wrongness is so abstract as to be a kind of moral principle. In that case, recognitional moral perception

³² Cullison, "Moral Perception," p. 162.

³³ Cullison, "Moral Perception," pp. 160-163.

and inferential moral perception are similar in that they both presuppose representations of moral principles. But this similarity is not enough to ground a reduction of moral perceiving-as to inferential moral perception, because retrieving a moral principle for matching is different from enlisting it in an inference to some conclusion.

Whether recognitional accounts ever reduce to representational accounts depends on whether the mental state retrieved in perceptual recognition can cognitively penetrate perceptual content. Susanna Siegel has argued that experts' perceptual recognition is often representational.³⁴ In one of her examples, she notes how your experience of seeing a pine might undergo a significant phenomenological shift if you spent your entire summer working a job where you had to identify them.³⁵ She contends that the best explanation of the phenomenological difference is that your perceptual experience now includes a representation of the property of being a pine. If repeated recognition begets new representations in perceptual content via this process of cognitive penetration, then recognitional perception is sometimes a species of representational perception.

Even if the two accounts can overlap in this way, though, they should remain distinct during the early stages of matchmaking. To ensure that we cover new ground, I will focus on moral perceiving-as in its novitiate stage, before perceptual recognition has had adequate opportunity to penetrate perceptual content. If perceptual matching in this

³⁴ Susanna Siegel, "Which Properties Are Represented in Perception?" in T. Szabo Gendler and J. Hawthorne (eds.), *Perceptual Experience* (Oxford: Oxford University Press, 2006), pp. 481-503, at pp. 490-491. These kinds of "phenomenological contrast" cases play a starring role in Siegel's Master Argument for a "permissive view" of perceptual content, according to which perceptual experience represents not only simple properties of color, shape, motion, and illumination but also what Siegel calls "K properties"—a diverse class that encompasses natural kind properties along with an indeterminate number of other properties that are not "standardly taken to be represented [by visual experience]" (see p. 482).

³⁵ Cowan, "Perceptual Intuitionism," at p. 14, offers this example of Siegel's as one possible "model" of an account of cognitive penetration.

early stage implies a source of a priori justification—as I argue it must—then this characteristic will be passed on to perceptual content when penetration occurs.

At this early stage, perceptual recognition relies on a background belief or memory. Suppose that it relies on a belief. For example, suppose that an instance of recognizing your perceptual experience as the president relies on a background belief about what the president looks like. In that case, the instance of perceiving-as is a source of epistemic justification of a corresponding perceptual belief only if the background belief is justified as well. And here we re-enter familiar territory.

Either this background belief is justified a priori or a posteriori. If it is justified a posteriori, it is not justified on the basis of perceptual experience or induction, for reasons we have already seen. A parallel version of my argument to the conclusion that inferential accounts of moral perception presuppose a priori justification establishes that perceptual recognition, when it relies on a background belief, must do the same.

But suppose that moral perception-as relies not on a background belief but rather on a (non-doxastic) memory. In that case, defenders of the recognitional account still owe us an explanation of why the matching of some new non-doxastic representation to some old one is a source of epistemic justification for an associated moral belief. To account for this justification, the recognitional account of moral perception will again have to rely on one of the other accounts. Either the matching non-doxastic representations (i.e., the current perceptual experience and the memory) include a distinct perceptual representation of a moral property (e.g., of wrongness) or they do not. If they do include such a representation, then the recognitional account can appeal to the distinct

representation of the moral property to explain why moral perception-as is a source of epistemic justification—but only by smuggling in the representational account.

But suppose that neither of the matching non-doxastic representations includes a distinct perceptual representation of the moral property. Then the justification of the associated perceptual belief must involve an inference from what is represented in the non-doxastic representations to a conclusion that represents more than what is represented. And once again we re-enter familiar territory. Because the recognitional account must draw from one of the other accounts to explain why perceiving-as is a source of epistemic justification, it inherits the problems of one of the other accounts. In all of its variations, it requires an element of a priori justification for the states that play an essential role in justifying moral beliefs.

IV. The Mystery of Moral Perception

A. The Reliability Challenge

So far I have been concerned with the epistemic justification of the beliefs formed on the basis of moral perception. I have argued that this justification is a priori (or at best: not clearly a posteriori). I have not discussed what might be called the “mechanism” of posteriori justification. As was mentioned above, we know quite a lot about the causal mechanism that governs perception. As the Benacerraf-Field challenge emphasizes, however, we do not have a similar understanding of the cognitive mechanism with which we gain substantive a priori knowledge.

Focusing on the mechanism of moral perception, defenders of moral perception might contend that there is an important respect in which their accounts are empiricist-

friendly, even if the justification implied by their accounts is not clearly a posteriori. According to this perspective, these accounts are empiricist-friendly because they assimilate the formation of moral beliefs into a broader, reputable causal mechanism. This causal mechanism provides an adequate explanation for the reliability of moral perception—and thus renders it, unlike substantive a priori knowledge, unmysterious. This putative feature of moral perception should be attractive to empiricists even if, from the standpoint of justification, it has a suspiciously a priori look.

The attempt to prove that moral perception is non-mysterious may be the hidden agenda driving lively discussion of the causal constraint on perception. To my knowledge, everyone who has defended an account of moral perception accepts the causal constraint or something much like it.³⁶ Moreover, their attempts to show that their accounts satisfy the constraint have been highly uniform.

Defenders of moral perception note that, even if moral properties are themselves causally inert, they stand in a very close relationship—at least supervenience—with simpler natural properties that are obviously causally active. When hoodlums torture a cat, the wrongness of the action does not causally affect you. But the cat's fiery tail and the hoodlums' malign expressions obviously do. Even if we do not have a direct causal connection with any moral property, we may still enjoy an indirect causal connection with many moral properties; we are connected to the moral property via our direct causal connection with some of the natural properties that ground it. According to defenders of moral perception, this indirect causal connection satisfies the causal constraint.

³⁶ E.g., Audi, *Moral Perception*, pp. 55-56, Cuneo, "Reidian Moral Perception," pp. 256-257. McBrayer, "Moral Perception and the Causal Objection," p. 303, proposes an "emendation" on the causal constraint that would allow some other relations that ensure "non-accidentality" to satisfy the constraint.

Even if it does, though, it fails to uphold a causal explanation for the reliability of moral perception in any of the three kinds of accounts that we have considered. This is most obvious when we consider inferential accounts of moral perception. Suppose again that moral perception is constituted by both a perceptual representation of natural properties (the hoodlums hurting the cat) and an inference to some moral conclusion (this action is wrong). The causal connection between the perceiver and the natural properties that ground the moral property explains the reliability of the perceiver's beliefs about these natural properties. But it does not seem to explain the reliability of the belief in the bridge-principle that the perceiver enlists in the inference to the moral conclusion. As we saw earlier, the antecedent belief in the bridge-principle cannot be justified on the basis of perceptual experience. We see now that the causal mechanism of perception does not explain the reliability of such a belief.

Since representational accounts of moral perception do not posit antecedent beliefs about bridge-principles, they might seem to avoid the charge I have made against the inferential account. Since they do not require antecedent beliefs, they should be spared the burden of having to explain the reliability of this kind of belief.

While it is true that they are spared this particular explanatory burden, they assume a parallel one. Even though representational accounts of moral perception do not presuppose that we reliably form beliefs about bridge-principles, they do presuppose that we reliably follow bridge-principles—in the way that our perceptual systems convert natural causal stimuli into phenomenological representations of normative properties. The indirectness of the causal connection between the moral perceiver and the moral property becomes relevant here. Unlike perception that involves a direct causal

connection with the perceived property, perception that involves an indirect connection must rely on a bridge-principle. When you (directly) perceive the round surface of a tomato, the part of the tomato that impacts your sense organs is also represented in your perceptual experience. As noted above, there is an extremely complex mechanical process whereby our cognitive faculties convert the causal stimulus into a phenomenological representation. But in cases of direct perception, cognition operates in service of a simple mimetic goal: like a Xerox machine, it simply produces a copy of its input.

When perception is indirect, however, our cognitive machinery must adhere to a more sophisticated algorithm. If moral perception is representational, our machinery does not simply produce a copy of the natural properties that causally impact us; it also adds normative detail. For example, if you see hoodlums torturing a cat, your cognitive faculties would not merely produce a representation of the natural properties of the event that impact you—they would also integrate an original representation of wrongness into the cartography. And the formula by which the natural stimuli are transformed into representations of the non-natural moral property must, again, depend on bridge-principles that specify links between the non-moral and the moral.

Since the recognitional account must invoke one of the other two accounts to explain why moral perception-as is a source of epistemic justification, it seems unlikely that this account will be able to provide, by itself, an adequate explanation of the reliability of moral perception. To discharge this explanatory burden, it will again have to piggyback on one of the other accounts. The retrieval and matching that is distinctive of the recognitional account explains how we relate our present moral experiences to

previous ones, but it cannot explain, on its own two feet, the reliability of either experience (or more exactly the belief-forming processes associated with them). If the inferential and representational accounts do not provide an adequate explanation for the reliability of moral perception, then the parasitic recognitional account does not provide an adequate explanation either.

While none of these accounts explain the way we reliably believe or follow these bridge-principles, they all imply that the cognitive success of moral perception depends crucially on them. Consider the (false) moral principle that torturing babies is morally right. If you believed this false bridge-principle, and enlisted this belief in a perceptual inference, the result would be perceptual error. You would perceive that an action is right, though it is in fact wrong. Similarly, if your perceptual faculties followed this principle when it produced representations of moral properties in response to certain natural stimuli—if, every time you were causally affected by events that constitute the torture of babies, your perceptual experience included a representation of rightness—then again the result would be perceptual error. If there is moral perception, and our cognitive lives do not involve widespread perceptual error, then most of the bridge-principles we believe or follow must be true.

The success of moral perception depends on our ability to reliably believe or follow a very complex set of bridge-principles. Since accounts of moral perception provide no explanation of this reliability, they are no improvement on the rationalist epistemological model.

B. Objections and Replies

One objection to my suggestion that non-naturalist moral perception is mysterious appeals to companions in innocence.³⁷ Other presumably non-mysterious forms of perception presuppose reliable bridge-principles, too. Consider Cowans's paradigm of phenomenal presence as absence. If we can perceive the occluded backside of a tomato, then there must be some bridge-principle that specifies that the present part of the tomato—which most directly makes contact with your sense organs—has a similarly round posterior. As with moral perception, more is represented in your perceptual experience than directly impacts your sense organs.

The realization that many forms of perception involve bridge-principles complicates the mysteriousness charge against rationalism. If substantive a priori knowledge is mysterious because of its reliance on bridge-principles, but perception involves similar bridge-principles, then the mysteriousness charge succeeds only if there is an explanatory disparity between the various bridge-principles. Empiricists who make the mysteriousness charge against rationalism need to say that there is an empiricist-friendly explanation for the bridge-principles that figure in (e.g.) the perception of the backside of a tomato, but not for the bridge-principles that figure in moral perception. While this disparity claim does not strike me as implausible, I will not discharge the empiricist's burden by attempting to defend it here.

I will note instead that defenders of moral perception have not upheld their share of this burden—to provide some explanation of the reliability of the bridge-principles on which the success of moral perception depends. Until accounts of moral perception explain rather than presuppose the reliability of our use of these bridge-principles, they

³⁷ For similar Companions in Innocence arguments, see Cuneo, "Reidian Moral Perception," p. 256, and McBrayer, "Moral Perception and the Causal Objection," pp. 295-298.

do not constitute an interesting alternative to moral rationalism. The plausibility of my thesis does not depend on my endorsement of any particular putative basis of an explanatory disparity that would render the perception of tomatoes unmysterious while upholding the charge against substantive a priori knowledge.

Another objection to my claim that moral perception is mysterious appeals to evolution to explain the reliability of our deployment of the bridge-principles. In response to evolutionary objections, moral realists have sketched evolutionary histories (many friendly to non-naturalism) that purport to vindicate our moral beliefs by uncovering evolved tendencies to believe true moral principles.³⁸ Suppose that one of these vindicating evolutionary stories is true (or, if it makes a difference, that we are justified in believing it). Presumably our justification for believing such a story would be a posteriori and non-mysterious. One might suspect that we can appeal to evolutionary forces to provide an empirical vindication of the bridge-principles presupposed by any account of non-naturalist moral perception—thereby rendering it entirely empiricist-friendly.

One problem with this attempt to enlist evolution to vindicate the empiricist-friendly character of moral perception is that empiricists enjoy no special claim to the vindicating evolutionary history. Rationalists are equally entitled to rely on successful evolutionary theses to absolve their own accounts of any charge of mystery.

In case this is not clear, consider how a moral rationalist might appropriate such a story to vindicate moral rationalism. According to rationalists, we are justified a priori in

³⁸ See Erik Wielenberg, “On the Evolutionary Debunking of Morality,” *Ethics* 120 (2010), 441-464; David Enoch, “The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope With it,” *Philosophical Studies* 148 (2010), 413-438; Knut Skarsaune, “Darwin and Moral Realism: Survival of the Fittest,” *Philosophical Studies* 152 (2011), 229-243. For a very different kind of putative evolutionary vindication of our moral beliefs, see Huemer, *Ethical Intuitionism*, p. 216.

believing many moral principles. According to one version of moral rationalism, reflecting on true principles tends to trigger a rational intuition that they are true. If charged with mysteriousness, a moral rationalist could reply that the tendency to have this rational intuition, upon consideration of the true principle, was shaped by evolutionary forces (per the details of some vindicating evolutionary history). If such an evolutionary appeal exonerates moral perception of the mysteriousness charge, then it should exonerate moral rationalism just the same. In their current state of development, vindicating evolutionary accounts of morality do not favor empiricist moral epistemology over moral rationalism. Thus, defenders of moral perception cannot fall back on evolutionary vindications to establish the empiricist-friendliness of their accounts.

V. Conclusion

Accounts of moral perception have been advertised as an a posteriori epistemological alternative to moral rationalism that promises to normalize non-naturalist moral epistemology. I have argued that these accounts do not live up to the hype. If it is not boring, non-naturalist moral perception is inferential, representational, or recognitional. Inferential accounts presuppose a priori justification of bridge-principles. Representational accounts involve a kind of justifying experience that may be a source of a priori justification. And recognitional accounts are parasitic on the inferential or representational account. Thus, each of these accounts conceals a core commitment of rationalism—to substantive a priori justification—or at best does not clearly uphold the empiricist's commitment against it.

We have considered whether these accounts might provide an interesting alternative to rationalism, if not for the justification they provide, then for the way that they explain the reliability of the beliefs so justified. But the accounts do not provide this either. They presuppose rather than explain reliable bridge-principles. Thus, they are on an explanatory par with rationalism. Since accounts of moral perception are not clearly distinct from rationalism, and are objectionable for the same reasons as rationalism, they do not constitute an exciting alternative to rationalism. Our spontaneous moral judgments are, at most, a marginal form of perception that is on par with substantive a priori cognition in the most philosophically interesting respects.

Chapter 2: The Causal Objection to Ethical Non-Naturalism

According to moral realism, moral properties do not depend on the intentional attitudes taken towards their natural base properties. For example, if torture is wrong, the wrongness does not depend on any intentional attitudes taken towards torture. Moral realism claims that moral properties are, in this sense, “stance-independent.”¹ Non-naturalism is a version of moral realism that emphasizes the deep differences between the moral domain and the natural domain.² To mention two of these: moral facts, unlike natural facts, are causally impotent and categorically reason-giving.³

While non-naturalism represents a commonsensical understanding of the nature of morality—one that upholds both its apparent objectivity and its striking discontinuity with the subject matter of the natural sciences—it is, nevertheless, a popular target of epistemological criticism.⁴ Recent epistemological critiques have focused on the genealogy of our moral beliefs.⁵ According to these *genealogical critiques*, certain facts about the formation of these beliefs imply a kind of epistemological disaster if non-

¹ Russ Shafer-Landau, *Moral Realism* (Oxford: Oxford University Press), p. 15.

² For a discussion of different ways of demarcating ethical naturalism and non-naturalism, see Shafer-Landau, *Moral Realism*, chapter 3, esp. pp. 58-65.

³ Thus, I assume that non-natural moral facts are causally impotent.

⁴ For classic critiques of non-naturalism, see A.J. Ayer, *Language, Truth, and Logic* (New York: Dover Publications, 1952), p.106; J.L. Mackie, *Ethics: Inventing Right and Wrong*, (Harmondsworth: Penguin, 1977), pp. 38-42; Michael Smith, *The Moral Problem* (Oxford: Blackwell Publishing, 1994), pp. 18-25.

⁵ See, e.g., Sharon Street, “A Darwinian Dilemma for Realist Theories of Value,” *Philosophical Studies* 127 (2006), pp. 109-166; Matthew Bedke, “Intuitive Non-Naturalism Meets Cosmic Coincidence,” *Pacific Philosophical Quarterly* 90 (2009), pp. 188-209. The target of Street’s dilemma is normative realism, which is wider than ethical non-naturalism along two different dimensions: 1) it encompasses non-moral as well as moral normativity, and 2) it encompasses “value naturalism” as well as non-naturalism (see pp. 135-141).

naturalism is true.⁶ Once you understand how our moral beliefs are formed, so the critique goes, you should see that they do not track non-natural moral facts.

In this essay, I explore one kind of genealogical critique—the Causal Objection—and, ultimately, sketch a solution to it that is provided by my account of Intellectual Perception. In Part I, I distinguish the Causal Objection from the Evolutionary Objection. In Part II, I develop the former objection along the lines suggested by Hartry Field. According to this formulation of the objection, non-naturalism is implausible because it offers no adequate explanation for the striking fact that many of our moral beliefs are non-accidentally true. In Part III, I assess the scope of the objection and conclude that its proper target is much wider than non-naturalism; the causal objection applies to all belief-forming processes about which causal connections with concrete particulars is inessential to the explanation of their reliability. Many of our *a priori* beliefs result from processes that fit this profile. In Part IV, I offer a limited defense of the Causal Objection by criticizing a number of candidate explanations for the non-accidentality of the relevant beliefs. In Part V, I sketch an account of Intellectual Perception that, if true, would ground a general solution to the Causal Objection, and I showcase the account’s surprising explanatory power.

I. Two Genealogical Critiques

We should begin by carefully distinguishing two kinds of genealogical critiques: the Causal Objection and the Evolutionary Objection. We can distinguish them with

⁶ Different critics have spelled out the disaster in slightly different terms. According to Sharon Street, the disaster is that it would be an “incredible coincidence” if many of our normative beliefs were true (“A Darwinian Dilemma for Realist Theories of Value,” p. 124). Matthew Bedke adds that the “need for a cosmic coincidence, once realized, constitutes a defeater” for our moral beliefs (“Intuitive Non-Naturalism Meets Cosmic Coincidence,” p. 190).

reference to the relevant facts that are alleged to generate the epistemological disaster. According to the Causal Objection, the relevant fact is that moral facts never cause our moral beliefs. Since this causal absence is entailed by a metaphysical commitment of non-naturalism—to the proposition that moral facts are causally impotent—the Causal Objection does not depend, like the Evolutionary Objection, on recent empirical work.

The spirit of the Causal Objection is best conveyed by an analogy with perception. According to a widely endorsed account of perception, what makes the difference between perceiving and misperceiving a physical object is an “appropriate causal connection” with it.⁷ If you seem to see your dog Buddy at your feet, but your sensory experience is caused by some hallucinogenic drug acting directly on your brain—rather than the real Buddy affecting your sense organs—then you do not actually perceive Buddy. Because Buddy is not part of the cause of the sensory experience, it would be a coincidence if Buddy actually were, as he appears to be, at your feet.

According to the Causal Objection, moral intuitions are relevantly similar to hallucinations.⁸ Just as Buddy does not cause your sensory experiences, so non-natural moral facts never cause any of your moral intuitions. Thus, if it would be a coincidence if our sensory experience accurately represented Buddy, then it would be a similar kind of coincidence if many of our moral intuitions accurately represented the moral facts.

⁷ In this essay, “appropriate,” when qualifying a causal connection, is meant to exclude deviant causal chains.

⁸ We can formulate genealogical critiques so that moral intuitions or moral beliefs are the primary target. A genealogical critique that targets only our moral intuitions is slightly weaker than the formulation that targets our moral beliefs at large, since moral beliefs that are not based on moral intuition escape this version of the critique. Since moral intuition is certainly an important (and arguably the only primary) source of epistemic justification for our moral beliefs, though, even the weaker critique, if successful, would generate a serious epistemological crisis for ethical non-naturalism (viz., the debunking of moral intuition). I will shift between these two formulations of genealogical critiques depending on what seems most helpful in the particular dialectical context.

While the Causal Objection claims that it is the absence of one kind of causal influence from the genealogy of our moral beliefs that generates the epistemological crisis, the Evolutionary Objection claims that it is the presence of a different kind of causal influence that generates the crisis. In particular, the Evolutionary Objection claims that it is the way that natural selection has shaped our moral beliefs—where “shaped” here denotes a kind of causal influence—that creates the problem. Since the nature of this shaping is an empirical matter, the Evolutionary Objection includes an empirical hypothesis that specifies the nature of this evolutionary influence on our moral beliefs.

In her influential statement of the Evolutionary Objection, Sharon Street delineates this hypothesis. According to Street, evolution selected for tendencies to form some moral beliefs rather than others only because having those tendencies motivated reproductively advantageous behavior in our evolutionary ancestors.⁹ The correspondence of those beliefs with stance-independent moral facts (i.e., the truth of those moral beliefs) does not explain why evolution selected for those tendencies. Given this hypothesis, it would be a massive coincidence if many of our moral beliefs were true.¹⁰

Finding an analogy for the Evolutionary Objection requires more imagination. Suppose that, every morning when you woke up, you seemed to hear a certain kind of songbird singing at your window. Having no reason to be suspicious, you form the belief every morning that there is this kind of songbird outside. Now suppose that evolutionary biologists discover that, in the Pleistocene environment, seeming to hear such a songbird

⁹ While Street advances this hypothesis to explain all kinds of normative judgment, I will consider the hypothesis only as an account of moral judgment.

¹⁰ Street, “A Darwinian Dilemma for Realist Theories of Value,” pp. 113-121, 127-128.

in the morning was selected for only because having this seeming boosted the spirits in a manner that attracted mates. The correspondence of such seemings with real songbirds does not explain why the disposition to have this seeming originally evolved or was passed on to us. Provided the truth of this story, it would be a massive coincidence if our songbird seemings accurately represented the presence of actual songbirds.

According to the Evolutionary Objection, our moral intuitions are relevantly similar to seeming to hear songbirds in this story. Our ancestors were selected to have certain moral intuitions because those intuitions conferred reproductively advantageous psychological effects—they made us take care of ourselves and cooperate with others—and not because those intuitions accurately represented the moral facts.¹¹ If you should not trust your seeming to hear songbirds, then neither should you trust your moral intuitions.

Unlike the Causal Objection, the Evolutionary Objection does not depend on the metaphysical thesis that moral properties are causally impotent. Even if moral properties are causally potent, the crisis is that they have not caused our intuitions in the appropriate way. To extend the analogy, even if such songbirds are real and causally active, the problem is that they do not appropriately cause our apparent experiences of them.

We can, then, state the relationship between the two genealogical critiques as follows. According to the Causal Objection, the absence of the moral fact from the cause

¹¹ This story is not a perfect analogy for every version of the Evolutionary Objection. For example, it does not perfectly fit Richard Joyce's influential statement of the Evolutionary Objection (*The Evolution of Morality* (Cambridge: MIT Press, 2007), Chapter 2). According to Joyce, it is primarily our moral concepts, not our moral intuitions, that evolution has selected for their beneficial psychological effects. According to Joyce, moral concepts imply categorical reasons for action. Evolution selected for concepts such as "right" and "wrong" because the categorical reasons implied by these concepts reinforced pro-social behavior (with all of its attendant benefits to the individual and the group).

of the moral belief is sufficient to generate the epistemological crisis even if the defenders of the Evolutionary Objection are mistaken in their empirical account of the genealogy of our moral beliefs. And according to the Evolutionary Objection, this empirical account is sufficient to generate the crisis even if non-naturalists are mistaken in their metaphysical commitment to the causal impotence of moral facts.

Clearly, the Causal Objection and the Evolutionary Objection are not the same objection. While the Evolutionary Objection has received the lion's share of the attention in the metaethical literature, I will focus primarily on the Causal Objection. A thorough examination of the Causal Objection is important, not only because the objection poses a serious challenge to non-naturalism that is logically independent of the Evolutionary Objection, but also because the examination alters the dialectic surrounding the Evolutionary Objection—as we will see when evolutionary considerations reenter the picture in Part IV.

II. The Causal Objection Developed

In the previous section, I have suggested a way to formulate the Causal Objection as an Argument from Analogy. According to this argument, moral intuitions are relevantly similar to sensory experiences that are caused by drugs. Non-naturalists, of course, are likely to reject this analogy. Indeed, they might claim that it reveals what is most wrongheaded about the Causal Objection: namely, that it assumes that any plausible account of moral knowledge must be modeled on perception. To the contrary, non-naturalists can contend, the explanation for moral knowledge is fundamentally different from the explanation for perceptual knowledge. The former does not, like the latter, cite

an appropriate causal connection with some relevant fact. This reply suggests a stronger reformulation of the Causal Objection. Instead of formulating the objection as an Argument from Analogy, we can reformulate it as an Argument from Elimination of putative (non-causal) explanations of moral knowledge.

In formulating the Causal Objection as an Argument from Elimination, I follow Hartry Field. Drawing on the work of Paul Benacerraf, Field formulated a causal objection to mathematical Platonism in the context of defending mathematical anti-realism.¹² According to mathematical Platonism, mathematical objects (such as numbers and sets) are abstract entities that, similar to non-natural moral properties, are stance-independent and causally impotent. Because of these relevant similarities between non-natural moral facts and Platonic mathematical entities, Field's formulation of the objection applies with equal force to (ethical) non-naturalism.

According to Field, there is a "striking fact" that "demands explanation:" namely, that whenever mathematicians believe that P, it is usually true that P.¹³ In other words, mathematicians' beliefs about mathematical propositions are usually true. While Field focuses on experts, I would add that a slightly weaker point is true of mathematical novices. Certainly the accuracy of both experts' and novices' mathematical beliefs is much higher than it would be if we formed our mathematical beliefs entirely at random. And the same is true of moral beliefs. So if this striking fact about our mathematical beliefs demands explanation, then a parallel fact about our moral beliefs demands explanation as well. When normal adult humans beings believe a moral proposition, the

¹² Paul Benacerraf, "Mathematical Truth," *The Journal of Philosophy* 70 (1973), pp. 661-679; Hartry Field, *Realism, Mathematics, and Modality* (Oxford: Basil Blackwell, 1989).

¹³ Field, *Realism, Mathematics, and Modality*, p. 230.

moral proposition is usually true—or at least, much more likely to be true than if we had formed our moral beliefs randomly. Using this last claim as the explanandum, we can reconstruct a causal objection to non-naturalism that parallels Field’s objection to mathematical Platonism:

1. A moral metaphysics (such as non-naturalism) is plausible only if it can offer some plausible explanation of the striking fact that most of our moral beliefs are true.
2. Non-naturalism cannot offer any plausible explanation of this striking fact.
3. Therefore, non-naturalism is implausible.

By slightly enlarging the explanandum, we can make the argument even stronger. If we assume that we have knowledge about some domain (and I will make this assumption about mathematics and morality), then the metaphysics of that domain should be able to explain not only the striking fact that many of our beliefs about that domain are true, but also that many of our true beliefs about that domain qualify as knowledge.

I assume that a belief counts as knowledge if and only if it is justified, true, and satisfies some kind of non-accidentality constraint (whose need is suggested by Gettier¹⁴). Since many of our moral and mathematical beliefs enjoy strong intuitive support and cohere with one another, it will not be difficult to account for their initial justification. The challenge is to explain the striking fact that they are true, and, moreover, that they are non-accidentally true.¹⁵ I will refer to the formulation of the Causal Objection that incorporates this expanded explanandum as the Master Causal Objection (MCO):

¹⁴ Gettier, Edmund. “Is Justified True Belief Knowledge?,” *Analysis*, 23 (1963), pp. 121–123. I make no effort to clarify why Gettier-style examples do not qualify as knowledge.

A1. Non-naturalism is plausible only if it can offer a plausible explanation of the striking fact that many of our moral beliefs are non-accidentally true.

A2. Non-naturalism cannot offer a plausible explanation of the striking fact that many of our moral beliefs are non-accidentally true.

A3. Therefore, non-naturalism is implausible.

In the next section, I will begin a two-stage evaluation of MCO. In the first stage of the evaluation (in Part III), I identify a generalization problem with MCO; this problem suggests that the argument is unsound. In the second and much longer stage (in Part IV), I defend the two premises of the argument. If the argument is unsound, my defense suggests, it is nevertheless difficult to refute. The two stages of the evaluation together motivate the fairly radical proposal I will develop (in Part V).

III. The Generalization Problem

The problem with MCO is that it generalizes across all substantive *a priori* knowledge—moral and non-moral knowledge alike.¹⁶ Since we certainly have much substantive *a priori* knowledge, the argument must be unsound.

To see why the objection generalizes, it will be helpful to consider a paradigm of *a priori* knowledge, such as the belief that nothing can be both green and red all over. Is

¹⁵ I will treat the non-accidentality as a property of a belief. But one could make the case that it is primarily a property of the justification of the belief. Nothing here hinges on this fine point.

¹⁶This is not an original point. For example, Lawrence Bonjour argues that the “objection can be extended to rationalist *a priori* knowledge and justification generally, and [...] has frequently generally been invoked in this broader form by others” (*In Defense of Pure Reason* (Cambridge: Cambridge University Press 1998), p. 156). Justin Clarke-Doane catalogues different attempts to generalize the objection to non-mathematical domains (“What is the Benacerraf Problem?” forthcoming in F. Pataut (ed.), *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity*). Finally, Russ Shafer-Landau argues that Bedke’s causal objection to ethical non-naturalism generalizes over “every case of modal, numerical, and philosophical truths” (“Evolutionary Debunking, Moral Realism, and Moral Knowledge,” *Journal of Ethics and Social Philosophy* 7 (2012), pp. 1-37, at p. 29).

there a causal explanation of why this belief is non-accidentally true? On the surface, it might seem that there is. After all, “red” and “green,” unlike “right” and “wrong,” are respectable natural properties. And particular red and green things affect us regularly. Our causal connection with these things might lead us to suspect that there is an adequate causal explanation for why the belief that nothing is green and red all over is non-accidentally true. And, if there is, then MCO does not generalize over all substantive a priori knowledge.

But the suggestion that there is a causal explanation for this knowledge is mistaken. Certainly, the causal impact of green and red particulars on our sense organs partly explains our ability to know contingent propositions such as that there is a green plant in the room and that the pen is red. But there is no similar causal explanation of our ability to know the necessary truth that nothing can be both red and green all over. In case this is not obvious, consider the epistemological situation of a brain in a vat. We can suppose that the brain is causally isolated from any green or red thing.¹⁷ Obviously, the brain would be at a huge epistemic disadvantage when it comes to knowing about its physical surroundings. It would have no perceptual knowledge. But the brain would be at no similar disadvantage when it comes to knowing necessary truths such as that nothing is red and green all over. Intuitively, the brain is just as well positioned as normal people to know that this is impossible. Since being causally impacted by red and

¹⁷ John Bengson offers the similar “case of Trip” to illustrate one kind of non-causal knowledge (“Grasping the Third Realm,” forthcoming in T. Szabo Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 5, at p. 14). In Bengson’s story, Trip has a hallucination in which he seems to see some colors and shapes that he has never encountered before. On the basis of this experience, he is able to know certain propositions about the relationships of these shapes and colors (e.g., that red is more like orange than it is like blue). In fn 27, Bengson discusses envatment in response to the objection that the Case of Trip is metaphysically impossible. Since a properly stimulated brain can experience features to which it bears no causal connections, it is not impossible that Trip could hallucinate new shapes and colors.

green particulars does not affect one's epistemic position vis-à-vis the proposition that nothing is red and green all over, it is unlikely that the explanation of the non-accidentality of this belief is causal. And since the substantive a priori proposition we have considered is representative of the kind, the same is true of most other substantive propositions that we can know a priori.

The Generalized Causal Objection (GCO) formalizes this reasoning:

B1. We have substantive a priori knowledge only if there is an adequate explanation of the striking fact that many of our a priori beliefs are non-accidentally true.

B2. There is no adequate explanation of the striking fact that many of our substantive a priori beliefs are non-accidentally true.

B3. Therefore, we have no substantive a priori knowledge.

Notice that GCO includes no commitment to non-naturalism about any domain. Even so, there is good reason to think that GCO is on par with MCO (which targets ethical non-naturalism). While non-naturalism about a domain guarantees there is no causal explanation for the relevant knowledge, we find with a priori knowledge that, even without this metaphysical guarantee, its explanation is not causal. Thus, the fate of a priori knowledge in general—and ethical non-naturalism in particular—depends on the adequacy of non-causal explanations (assuming that A1 and B1 are true).

My endorsement of this generalization problem should put my following defense of MCO in perspective. When I defend its most controversial premise—that there is no adequate explanation of the striking fact that most of our moral beliefs are non-accidentally true—my ultimate goal is not to undermine ethical non-naturalism (by

defending the argument) but rather to force a radical reply to the objection (i.e., my account of Intellectual Perception). In this spirit, I now turn to defend MCO.

IV. A Defense of the Master Causal Objection

A. A Defense of Premise A1

Recall MCO:

A1. Non-naturalism is plausible only if it can offer an adequate explanation of the striking fact that many of our moral beliefs are non-accidentally true.

A2. Non-naturalism cannot offer an adequate explanation of the striking fact that many of our moral beliefs are non-accidentally true.

A3. Therefore, non-naturalism is implausible.

This argument is essentially an Argument from Elimination because most of the controversy surrounds A2—the claim that non-naturalism cannot provide the needed explanation—and the only way to defend that premise (as far as I can see) is via the elimination of candidate explanations. Even though the heart of my defense of MCO is my Argument from Elimination for A2, I will begin with a discussion of A1, which will pave the way to our discussion of A2 by raising the issue of explanatory adequacy as it pertains to beliefs about necessary truths.

Field interprets David Lewis as rejecting an analogue of A1. In *On the Plurality of Worlds*, Lewis considers whether a version of the Causal Objection threatens his own modal realism.¹⁸ Since we are causally isolated from the concrete possible worlds that Lewis posits, a parallel version of the Causal Objection would seem to apply to Lewis's account. But Lewis dismisses the objection on the grounds that Benacerraf and Field's

¹⁸ David Lewis, *On the Plurality of Worlds* (Oxford: Blackwell Publishers, 1986), pp. 108-115.

explanatory demand is improper, given that the real modal facts could not have been otherwise. Field summarizes Lewis's reasoning as follows:

“In his recent book, David Lewis has adopted a somewhat similar formulation of the Benacerrafian challenge...but he holds that it does not pose a genuine problem in the mathematical case, because all facts about the realm of mathematical entities hold necessarily. More fully, Lewis' idea is that we do need—and do have, at least in outline—an explanation of the reliable correlation between the facts about electrons and our electron beliefs (...); or, as he puts it, we need and have an account (in this case a causal account) of the way in which ‘electron’ beliefs counterfactually depend on the existence and nature of electrons. But it is only because the existence and nature of electrons is contingent that it makes sense to ask for an explanation of the counterfactual dependence of electron beliefs on the existence and nature of electrons...Consequently, since mathematics consists entirely of necessary truths, there can be no sensible problem of explaining why it is that our mathematical beliefs are a reliable indicator of mathematical facts.”¹⁹

Applied to the mathematical analogue of MCO, there are at least two ways that we can interpret the view that Field attributes to Lewis. First, we can understand Lewis as denying an analogue of Premise A1. On this reading, Lewis is saying that there is no explanation (a fortiori no adequate explanation) for the reliability (and, I would add, non-accidentality) of our mathematical beliefs, but because the contents of these beliefs are necessary truths, this explanatory absence does not count against mathematical realism.

We can also interpret Lewis as objecting to an analogue of A2. According to this reading, Lewis is saying that there is an explanation of the reliability of our mathematical beliefs, and the modality of the mathematical facts (i.e., their necessity) provides this adequate explanation. The difference between these two positions is that the former claims that the necessity of the mathematical facts removes, while the latter claims that it satisfies, the putative explanatory demand.

¹⁹ Field, *Realism, Mathematics, and Modality*, 233. Before dismissing Benacerraf's explanatory demand, Lewis spends a good deal of time puzzling over what kind of explanation is being demanded (*On the Plurality of Worlds*, pp. 113-114).

I doubt that there is any substantial difference between these two descriptions of Lewis's argument. Under either description, the substantial point is that, when justified true beliefs are about necessary truths, the modality of the belief entails that the belief is non-accidentally true. Whether this entailment relation satisfies an explanatory demand or prevents the demand from arising in the first place depends on the low-stakes semantic question of what counts as an explanation. The interesting substantive issue is whether this modal property of these justified true beliefs (viz., that their contents are necessary truths) entails that the justification of these beliefs is non-accidental.

Before responding to this substantive issue, it is worth identifying a contested presupposition of Lewis's argument: namely, that mathematics concerns only necessary truths. According to Field, some mathematical facts are "mixed" in the sense that they combine necessary mathematical truths with contingent facts about the world.²⁰ And we cannot always prise apart mixed mathematical facts into their necessary and contingent components. Since many moral facts are similar mixtures of necessarily true moral principles and contingent facts about the world, Field's point generalizes to the moral domain. For mathematics or morality, Lewis's response to the causal objection succeeds only if, and only to the extent that, our beliefs about these domains have as their contents necessarily true propositions.

Even if all of our moral beliefs did concern necessary truths (or can be separated into combinations of necessary truths and explicable contingent truths), still, Lewis's substantive claim is implausible. As I will show, there are examples of justified beliefs about necessary truths that are accidental in the relevant sense. These examples draw

²⁰ Field, *Realism, Mathematics, and Modality*, pp. 233-234.

intuitive support for the claim that this modal property of the content of a justified true belief does not entail that the belief is non-accidental.

The first counterexample is of a belief in a necessary truth that is justified on the basis of specious testimony. Suppose that you believe a complex mathematical equation on the authority of a celebrated professor. Later, however, you learn that the aging professor is losing his mind. In every lecture, he randomly strings together numbers and symbols and teaches the pseudo-equations to his students. Against all odds, it turns out that one of these equations happens to be true. Intuitively, your justified true belief is accidental in the relevant sense; it does not qualify as knowledge.

If issues surrounding testimonial knowledge are distracting, consider a second counterexample. In this one, you believe the same complex equation as in the previous example. But this time, you believe it on the basis of an intuition caused by a brain lesion. The brain lesion causes you to randomly have intuitions about the truth-values of mathematical propositions.²¹ Because of the brain lesion, you believe that this equation is true. Coincidentally, it is true. Once again, intuitively, your justified true belief in this necessary truth exhibits the relevant kind of accidentality.

As these two examples demonstrate, justified true beliefs about necessary truths can be accidental in the relevant sense. Since this is so, a justified true belief in a necessary truth does not entail the relevant kind of non-accidentality. Even if all of our justified true moral beliefs are about necessary truths, it does not follow that many of these beliefs are non-accidentally true. Something else—or something in addition—is needed to explain the distinction between justified true beliefs about morals and mathematics that are merely accidentally true from those that are non-accidentally true.

²¹ This example is adapted from John Bengson's essay ("Grasping the Third Realm," p. 6).

We will now turn to consider putative explanations of why some of these justified true beliefs are non-accidentally true, in the context of assessing Premise A2 of MCO.

B. A Defense of Premise A2

To recall, A2 states:

A2. Non-naturalism can offer no adequate explanation for the striking fact that many of our moral beliefs are non-accidentally true.

In this very long section of the essay, I will offer a limited defense of this premise by criticizing a number of candidate explanations for this striking fact. In the process of eliminating these, I will begin to identify constraints on an adequate explanation. I will begin by looking at purely psychological and purely factive explanations. I claim that both of these are inadequate because they do not relate the knower to a relevant fact.

i. Explanations that fail the Knower/Fact Constraint

In the same essay in which he discusses Lewis's response to a version of the Causal Objection, Field considers whether the systematization of mathematics might be able to explain the striking fact that he is concerned with:

“For as mathematics has become more and more deductively systematized, the truth of mathematics has become reduced to a smaller and smaller set of axioms; so that we could explain the fact that the mathematicians' beliefs tend to be true by the fact that they could be logically deduced from axioms, if we could just explain the fact that what they take as axioms tend to be true.”²²

We can devise a parallel explanation for moral knowledge. After all, fundamental moral principles are relevantly similar to mathematical axioms. Just as it is possible to deduce many non-axiomatic mathematical propositions from mathematical axioms, it is possible to derive other moral principles or (if relevant non-moral facts are known) facts about the instantiations of moral properties from fundamental moral principles. If we can explain

²² Field, *Realism, Mathematics, and Modality*, pp. 231-233.

the non-accidentality of our non-axiomatic mathematical beliefs by appeal to this kind of deduction, then we should be able to do the same for our derived moral beliefs.

The putative explanation just described is actually ambiguous between two putative explanations: one psychological and the other factive. According to the psychological explanation, it is our beliefs (a psychological phenomenon) about mathematical axioms that explain our non-axiomatic beliefs. The former beliefs explain the latter beliefs in the sense that we “deduce” the latter from the former. (“Deduce” may resist reduction to a purely psychological phenomenon, but this complication is not important to what follows.) According to the purely factive explanation, moral principles (understood as extra-mental entities—facts in some broad sense) explain non-axiomatic moral facts via a kind of grounding relation. Which explanation, then, does Field dismiss? I understand him as rejecting the psychological explanation. But in order to be thorough, we will have to consider them both.

One problem with the psychological explanation is that it is unclear whether many people have beliefs with mathematical axioms or fundamental moral principles as their contents. If many people do not have axiomatic beliefs about some domain, then we cannot appeal to these beliefs to explain (via deduction) the striking fact that many of their non-axiomatic beliefs about that domain are non-accidentally true.

In the passage quoted above, however, Field suggests an even deeper problem with this putative explanation—namely, that it offers no explanation for the equally striking fact that many of our beliefs about mathematical (or, I would add, moral) axioms are non-accidentally true. If these axiomatic beliefs are themselves only accidentally

true, then whatever we deduce from them will be at most accidentally true as well.²³

While this putative explanans serves to shrink the explanandum (it leaves us with fewer beliefs to explain), the smaller explanandum is just as inexplicable as the original one, and for exactly the same reason. Unlike the causal explanation we accept for perceptual knowledge, this psychological explanation for moral knowledge seems inadequate.

And we have some idea why it seems inadequate. Note that the explanandum incorporated into MCO (i.e., the striking fact that many of our moral beliefs are non-accidentally true) presupposes the existence of two entities. Just as perception presupposes a perceiver and a perceived, the explanandum presupposes both a moral knower and a moral fact that is known. One obvious difference between the causal explanation of perceptual knowledge and this candidate explanation of our mathematical or moral knowledge is that the former relates (i) some psychological property of the believer (i.e., her sensory experiences) and (ii) the fact that the belief is about (i.e., the physical fact), while the explanation we have just considered relates only the believer's psychological properties to one another. In particular, it relates some of her beliefs to some more of her beliefs. The absence of the relevant fact from the explanation is what appears to be driving the intuition that it is inadequate.

This point of inadequacy suggests an explanatory constraint. An explanation for this striking fact is plausible only if it relates the moral knower (or certain properties of the moral knower) to the moral facts (or certain properties of the moral facts). I will refer to this explanatory constraint as the "Knower/Fact Constraint."

This Knower/Fact Constraint receives some confirmation from the obvious inadequacy of other explanations that fail it. We have looked at a putative explanation

²³ Field, *Realism, Mathematics, and Modality*, p. 232.

that relates different psychological properties of the moral knower to each other. In this explanation, the relevant psychological properties are both beliefs, and they are related via an unanalyzed “deduction” relation. But we can imagine a different psychological explanation that involves mental states besides beliefs and explanatory relations different from deduction.

For example, we can devise a psychological explanation that cites intuitions—where intuitions are understood entirely in psychological terms. Certainly, we often base our beliefs about moral principles on our intuitions. Thus it makes sense to cite our intuitions to explain our beliefs via some kind of basing relation. But if intuitions are a purely psychological phenomenon, then this psychological explanation will be inadequate for the same reason that the psychological explanation that involves only beliefs is inadequate: namely, it doesn’t explain the striking fact that our intuitions are non-accidentally true. To explain this striking fact, the psychological properties of the knower need to bear some relation to the relevant facts that make those beliefs true.

Just as we cannot explain the relevant striking fact by relating the moral knower’s various psychological properties to each other, so we cannot explain it by relating her psychological properties to her physical properties. We could imagine a neurological explanation of the reliability of our moral intuitions or beliefs that relates these psychological states to underlying neurological states via some kind of grounding relation that relates the neurological and the mental. For this explanation to be adequate, however, it obviously must be supplemented with some other explanation that links one of these relata to the facts (i.e., an explanatory connection that relates neurological or psychological facts and moral facts).

We have looked at putative explanations that violate the Knower/Fact Constraint because they relate only various properties of the knower to each other and thereby leave out the facts. But an explanation can also fail the constraint by relating only fact to fact and thus excluding the moral knower. This brings us back to the second interpretation of the position that Field criticizes. According to this interpretation, our axiomatic beliefs explain our non-axiomatic beliefs via a grounding relation. This grounding relation explains why our beliefs qualify as knowledge. Applied to the moral domain, the grounding relation is between fundamental moral principles and derived moral principles (or perhaps instantiations of principles in concrete states of affairs). It should be clear, however, that this explanation of how certain moral facts entail or ground other facts does not explain any striking fact about our moral beliefs. To explain why our moral beliefs are non-accidentally true, these facts must bear some relation to us.

To explain the striking fact that many of our moral beliefs are non-accidentally true, it is not enough to explain how certain properties of the moral knower are related to each other or how certain moral facts are related to each other. So I will turn from these putative explanations that fail the Knower/Fact Constraint to consider some that satisfy it.

ii. Evolutionary Explanations

In their current state of development, evolutionary explanations of our moral faculties are sketchy; they usually tell us why the developed faculty (itself the output of a long evolutionary process) was adaptive, rather than explaining what was adaptive about each of the mutations that went into its design. Even in their sketchy states, though, they do purport to relate the moral knower to the facts—either via a third factor or via a

broader cognitive faculty with which we are able to believe propositions because they are true.

In the beginning of the essay, I distinguished the Evolutionary Objection from the Causal Objection. In response to the Evolutionary Objection, some defenders of moral realism have argued that evolutionary considerations, far from undermining our moral beliefs, actually serve to vindicate them. They propose *vindicating evolutionary explanations* of our moral faculties. I am thus lifting these explanations from their original dialectic context—as responses to the Evolutionary Objection—to see whether they ground an adequate response to the Causal Objection, in particular to MCO.

Responses to the Evolutionary Objection can be divided into two categories: those that concede and those that reject the empirical premise of the objection. For each of these categories of response, there is a corresponding vindicating evolutionary explanation that we will need to consider.

Third-factor explanations fall into the first category.²⁴ They concede to the debunker that evolution selected certain moral belief contents purely for their psychological effects and not because they corresponded with the moral facts. (Recall the songbird analogy.) But they claim that, far from leading to an epistemological crisis for moral realism, the truth of this evolutionary story would provide the needed explanation of the striking fact that our moral beliefs are non-accidentally true.²⁵

²⁴ See Erik Wielenberg, “On the Evolutionary Debunking of Morality,” *Ethics* 120 (2010), pp. 441-464; David Enoch, “The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope With it,” *Philosophical Studies* 148 (2010), pp. 413-438; Knut Skarsaune, “Darwin and Moral Realism: Survival of the Iffiest,” *Philosophical Studies* 152 (2011), pp. 229-243. While I tailor my critique to Enoch’s third-factor explanation, parallel critiques apply to the other two.

²⁵ While proponents of third factor explanations may purport to explain only the weaker explanandum clarified above (the striking fact that our moral beliefs are true) and not the stronger explanandum (which

Key to this vindication is the eponymous third factor that explains both our moral beliefs and the moral facts. In David Enoch’s proposal, the third factor is the “fact that survival is good.”²⁶ This fact partly explains our (pro tanto) moral obligation to promote survival. And survival is, as Enoch explains, part of the “aim” of evolution. Because evolution aims at survival, and the fact that survival is good explains many moral obligations, there is a pre-established harmony between what evolution would cause us to believe about morality and what is, in fact, true.²⁷ Because of this pre-established harmony, it is no accident that many of our moral beliefs are true.

If third-factor explanations succeed, then the empirical premise of the evolutionary objection actually—somewhat ironically—provides the explanation that undermines the Causal Objection; what was thought the source of one epistemological crisis actually provides the solution to another. Unfortunately, there are two strong reasons to think that third-factor explanations are unable to provide such vindication.

The first reason is that they involve an illicit appeal to defeated first-order normative beliefs.²⁸ Enoch assumes the first-order normative belief that survival is good in his explanation of why evolution would push us towards the moral facts. But if the evolutionary story conceded by third-factor explanations is correct, then most of our normative beliefs—including this one—have been selected only for their psychological effects, with no regard for the truth. If this is the case, then we should not believe that the

includes their non-accidental truth), I will assess these explanations as putative explanations of the stronger explanandum.

²⁶ Enoch, “The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope With it,” p. 19.

²⁷ Ibid.

²⁸ Russ Shafer-Landau develops this criticism of third-factor explanations (“Evolutionary Debunking, Moral Realism, and Moral Knowledge,” pp. 32-33.)

proposition that survival is good corresponds with the normative facts (i.e., is true).

Thus, we are not entitled to appeal to this proposition in a vindicating explanation of our moral faculties.

A second reason to think that third-factor explanations do not provide an adequate explanation of the non-accidentality of our moral beliefs is a dialectical consideration that is apparent only in the context of discussing the Causal Objection. The problem is that third-factor explanations do not ground a response to the Causal Objection that generalizes over every substantive *a priori* domain. In other words, they do not provide an explanation that would refute B2 of GCO (which, recall, claims that there is no adequate explanation of the striking fact that many of our *a priori* beliefs are non-accidentally true). When we find the same philosophical puzzle in a range of closely related cognitive domains, it counts in favor of a putative solution if it solves the puzzle in each of them. But third-factor explanations certainly do not; even if they ground an adequate response to the causal objection to (ethical) non-naturalism, they leave us where we began with respect to every other application of the causal objection.

Third-factor explanations cannot be generalized because they depend on idiosyncrasies of the normative domain. In Enoch's account, for example, Darwinian forces are alleged to push us towards the moral facts because one of the aims of evolution—namely, survival—happens to correspond with an important normative fact. But there is no similar correspondence between the aims of evolution and the contents of most of our other *a priori* beliefs, which do not concern survival at all. For example, there is no reason to think that Darwinian forces aim at the fact that nothing can be both red and green all over in the same way that they aim at survival.

Since third-factor explanations do not extend to a priori knowledge in general, they cannot ground a reply to GCO. Thus, even if we were to claim that this kind of evolutionary explanation grounds an adequate reply to the Causal Objection to (ethical) non-naturalism, we would still need to find another explanation to refute the more general causal objection. It seems more likely that this general explanation—whatever it is—also explains the parallel striking fact about our moral beliefs than that an exceptional third factor explains only our moral knowledge.

The second vindicating evolutionary explanation that has been proposed by moral realists would, if successful, vindicate all substantive a priori knowledge. In other words, it would refute A2 of MCO as well as B2 of GCO. According to a Wide Evolutionary Explanation (as I will refer to the second vindicatory attempt), our moral faculties evolved as part of a broader cognitive faculty whose scope encompasses much more than just moral belief. Michael Huemer offers a helpful analogy:

...[W]hy do we have the ability to see stars? After all, our evolutionary ancestors presumably would have done just as well if they only saw things on Earth. Of course, this is a silly question. We can see the stars because we have vision, which is useful for seeing things on Earth, and once you have vision, you wind up seeing whatever is there sending light in your direction, whether it is a useful thing to see or not. Likewise, once you have intelligence, you wind up apprehending the sorts of things that can be known by reason, whether they are useful to know or not. Thus, humans are capable of learning to play chess at an incredibly sophisticated level, despite that the environment of Australopithecus contained no chess boards. If some evaluative truths can be known through reason, we would likely know them whether they were useful to know or not.”²⁹

According to the evolutionary explanation suggested here, evolution selected for some general reasoning ability—presumably because being able to reason was adaptive—and this general ability somehow implied the ability to reliably form beliefs about morality. Thus, the evolutionary history of our moral beliefs just is the evolutionary history of this

²⁹ Michael Huemer, *Ethical Intuitionism* (New York: Palgrave Macmillan, 2005), p. 216.

broader cognitive faculty—just as the evolutionary history of our ability to see stars just is the evolution of the eye. There is no autonomous evolution of our moral faculties.

A Wide Evolutionary Explanation is inconsistent with the empirical premise of the Evolutionary Objection because, in that evolutionary story, there is a privileged segment of evolutionary history in which the contents of our moral judgments were shaped independently of the design of a more general cognitive faculty. We can adapt Huemer's analogy to make the inconsistency more apparent. Suppose that evolution, even before it designed the eye, selected for dispositions which guaranteed that we would form particular beliefs about stars; and, now, even as a general visual faculty controls many of our beliefs about faraway objects, this deviant faculty takes over to control all of our star beliefs. In Street's evolutionary hypothesis, the evolution of our moral faculties is analogous to the autonomous evolution of our star faculties in this imagined evolutionary history. The hypothesis is inconsistent with Huemer's suggestion that the evolution of our moral faculties is implied by the evolution of the broader faculty.

To see why the Wide Evolutionary Objection does not succeed as a response to the Causal Objection, it will be worthwhile to identify one crucial limitation of Huemer's analogy—i.e., one respect in which the evolutionary explanation of our sense of sight (or any perceptual faculty) must be importantly different from the evolutionary explanation of the more general reasoning ability that comprises our moral faculties. The difference is this: the evolutionary explanation of perception incorporates a causal explanation, but the evolutionary explanation of our moral knowledge cannot. Evolution explains perception via the design of a perceptual system that converts the effects of physical objects into mental representations. However, for reasons that should be clear by now,

the evolutionary explanation for moral knowledge cannot incorporate a causal explanation in a similar way. A Wide Evolutionary Explanation of our moral faculties must relate the knower to the facts without the aid of this causal mechanism.

The causal mechanism is important to Huemer's perceptual analogy because it explains why a diverse range of beliefs about the physical world all count as outputs of the same faculty. A general sense of sight includes my ability to see nearby objects such as my feet, and faraway objects such as stars, because the same perceptual apparatus is receptive to the effects of each. If our diverse beliefs about physical objects were not bound together by such a causal mechanism—or some comparable one—then the fact that one belief about a physical object is true would provide no reason to think that some other such belief is true. Our beliefs about stars inherit the credibility of a more general sense of sight because they fall within the scope of the presiding causal mechanism.

In their current state of development, Wide Evolutionary Explanations resemble the evolutionary explanation of perception *minus* the incorporated causal explanation (or any comparable explanatory relation). In particular, Wide Evolutionary Explanations lack an account of a general mechanism that explains why some subsets of our reasoning abilities (such as our moral faculties) count as part of a broader faculty (Huemer's "reason" or Parfit's "rationality"³⁰). In this omission, the Wide Evolutionary Explanations proposed by Huemer and Parfit are representative.

Let us look more closely at Parfit's Wide Evolutionary Explanation. According to Parfit:

³⁰ Huemer, *Ethical Intuitionism*, p. 216; Derek Parfit, *On What Matters*, Vol. 2 (Oxford: Oxford University Press, 2011), p. 494.

“It may be true that, just as cheetahs were selected for their speed, and giraffes were selected for their long necks, human beings were selected for their rationality. That may be how we became able to reason validly, and respond to reasons.”³¹

In his skeletal Wide Evolutionary Explanation, Parfit lines up various accomplishments of reason on a spectrum from primitive to sophisticated. He notes how dogs rely, in the hunt, on a kind of process of elimination; and that early humans must have relied, as potential prey, on simple arithmetic: “Since three lions went into my cave, but only two lions have come out, one lion is still in the cave.”³² Citing these simplest cases of reasoning, Parfit suggests that evolution gradually refined our reasoning abilities to produce the likes of Euclid, Newton, and Godel—as well as, presumably, every average Joe’s capacity for substantive a priori (including moral) knowledge.³³

While Parfit traces an interesting trajectory, he does not propose anything to play the same role in this evolutionary account that the causal relation plays in the evolutionary account of perception—what we might call reason’s “mechanism.” Such a mechanism would explain *how* evolution designed our brains such that we are able to form non-accidentally true beliefs about a diverse range of necessary truths. It would also explain, most importantly, why the evolutionary vindication of a very general ability implies the vindication of the select cognitive faculties that are targeted by causal objections—in particular, our moral faculties.

By omitting from the evolutionary explanation of our reasoning abilities any account of the mechanism of reason, Wide Evolutionary Explanations are sketchy at the

³¹ Parfit, *On What Matters*, Vol. 2, p. 494.

³² Parfit, *On What Matters*, Vol. 2, p. 496. Cp., Richard Joyce, *The Evolution of Morality* (Cambridge: MIT Press, 2007), pp. 182-183.

³³ *Ibid.*

very point that is most important to meeting the explanatory demand raised by causal objections. When searching for a non-causal explanation for substantive a priori knowledge, it is not illuminating to point out that evolution must have designed some general cognitive system that makes us capable of that kind of knowledge. For what causal objections (such as GCO and MCO) question is whether there is any cognitive system capable of this striking epistemic accomplishment.

In each of the next two sections, we will consider two explanatory relations that, like these evolutionary explanations, satisfy the Knower-Fact Constraint. Unlike third-factor evolutionary explanations, however, these putative explanations purport to ground a response to GCO as well as MCO. Moreover, they feature explanatory relations (constitution and a novel kind of causation, respectively) that, if incorporated by some Wide Evolutionary Explanation, might be able to play an explanatory role that is analogous to the role of causation in the evolutionary explanation of perception.

iii. Constitutive Explanations

John Bengson proposes a constitutive explanation to vindicate our intuitions in response to a similar causal objection. The target of this broad causal objection is “Realism,” a “broadly Platonistic’ metaphysics of necessary truths, according to which “what are known are facts about mind-independent abstract entities.”³⁴ To explain how our intuitions about abstracta could be non-accidentally true, Bengson proposes that such entities can partly constitute our intuitions—just as, according to Naïve Realism about Perception, physical objects partly constitute our perceptions. Since this constitutive

³⁴ Bengson, “Grasping the Third Realm,” p. 1.

explanation relates intuitions—understood as psychological phenomena—to abstract entities—facts in the relevant sense—it satisfies the Knower-Fact Constraint.³⁵

In his account, Bengson uses the terms “accidentality” and “non-accidentality” in a slightly different sense than we are using it here. He distinguishes the “doxastic accidentality” that is a property of a belief formed in a fake barn scenario from the “source accidentality” that is a property of one’s sensory experience in a veridical hallucination.³⁶ According to Bengson, the deeper challenge is to explain how our intuitions about intuitions are not source-accidental. For with beliefs that exhibit “doxastic accidentality,” “there is a sense in which one could come to know on the basis of the source state, even if one does not in fact do so on a particular occasion.”³⁷ If it were not for the presence of fake barns, then your perceptual experience of the true barn would be a source of knowledge. In contrast, source accidentality “strikes at the root..., for it renders the source state unable to supply knowledge.”³⁸ Bengson purports to address the deeper explanatory challenge: namely, to show that many of our intuitions about abstracta are not source-accidental.

The sense of non-accidentality that has been incorporated into our two main arguments (MCO and GCO) plausibly implies both of kinds of non-accidentality that Bengson distinguishes. Plausibly, a justified true belief that constitutes knowledge (and

³⁵ In this section and the next, I will use the term “relevant abstract entity” or “relevant abstracta” to stand for the abstract entity that stands in the proposed explanatory relations. I do not take a stand on what kind of abstract entity is best suited to play this explanatory role. Presumably it is a necessarily true proposition or the “abstract fact” that such a propositions represent (I borrow this term from Bengson, “Grasping the Third Realm,” p. 3).

³⁶ Bengson, “Grasping the Third Realm,” p. 6-7, fn 13.

³⁷ Bengson, “Grasping the Third Realm,” p. 6, fn 13.

³⁸ Bengson, “Grasping the Third Realm,” p. 7, fn 13 (continued from p. 6).

thus is “non-accidental” in our sense) does not exhibit doxastic accidentality and is based on a mental state that is not source accidental. I argue here that even if Bengon’s constitutive relation can explain how our moral intuitions could be not source accidental, he does not adequately explain the striking fact that the beliefs based on those intuitions are “non-accidental” in the sense that qualifies them as knowledge.

The easiest way to see this is by focusing on a case in which a belief that clearly does not qualify as knowledge is based on an intuition that is not source-accidental. Let us, then, revisit the case (adapted from Bengon’s essay) of the person with a brain lesion—call him “Leo”—that randomly causes him to have intuitions, some of which are necessarily true. Obviously, the belief based on this intuition is merely accidentally true. It does not constitute knowledge.

Now consider the following question: Is it possible that Leo’s intuition could be constituted by some relevant abstract entity? If so, then that intuition would not be source accidental. Thus, we would have a clear case of a belief that does not constitute knowledge based on an intuition that is not source accidental. This case would show that being based on an intuition that is not source accidental is not sufficient to make a belief non-accidentally true. And this would support my central claim about Bengon’s account that, even if it adequately explains how our intuitions might not be source-accidental, it does not adequately explain the striking fact that many of the beliefs based on those intuitions are not merely accidentally true.

One might object to my central claim on the grounds that it is impossible that Leo’s intuition could be constituted by relevant abstract entities. If it is impossible that intuitions caused randomly by brain lesions could be constituted by relevant abstracta,

then such an intuition must be source accidental. If Leo's intuition is necessarily source accidental, then the belief based on this intuition is not an example of a merely accidentally true belief based on an intuition that is not source accidental.

The impossibility claim, however, is *ad hoc*. No feature of Bengson's account suggests that an intuition caused by such a brain lesion could not be constituted by the relevant abstracta. If any feature of Bengson's account could motivate the claim that this is impossible, it would seem to be what Bengson refers to as the "mechanism underwriting" his constitutive explanation: his account of "property ensurance."³⁹ But property ensurance does not, I will argue, provide any such motivation.

While the details of Bengson's account of property ensurance are complicated, the main idea is easy to grasp. Property ensurance has to do with the way that the constituting material passes on many of its properties to whatever it constitutes. For example, the fragility of the glass "ensures" that the constituted vase will also be fragile. In general: when X constitutes Y, the properties of X often explain the properties of Y. In Bengson's account, property ensurance is important because it explains why true intuitions that are partly constituted by relevant abstracta are non-accidentally true. The relevant abstracta have the property of "being the way the world is."⁴⁰ Property ensurance guarantees that the intuition that is partly constituted by a relevant abstract entity is "not merely accidentally correct."⁴¹

³⁹ Bengson, "Grasping the Third Realm," p. 20.

⁴⁰ *Ibid.*, p. 29.

⁴¹ *Ibid.* I have left out all detail that is not essential to my main point. In Bengson's more thorough account, property ensurance is unpacked in modal terms: "a's having F ensures b's having G iff: necessarily, if a is F and a constitutes b, then b is G" (Bengson, "Grasping the Third Realm," p. 18). Bengson also distinguishes property ensurance from property inheritance (see pp. 18-19). In his schema, property inheritance is a species of property ensurance in which the ensured property is identical to the

While such ensurance explains why an intuition that is constituted by a relevant abstract entity is guaranteed to be true, it does not seem to imply any constraints on true intuitions that can be constituted by these abstract entities. But such a constraint is what would be needed to explain why Leo's true intuition could not be constituted by some relevant abstract entity. Put in a slightly different way: nothing in the nature of the ensurance relation prevents some relevant abstract entity from ensuring that one of Leo's true intuitions is non-accidentally true, too. And such ensurance would "underwrite" the constitutive relation that, according to this objection, is supposed to be impossible.

Since there is nothing in Bengson's account that explains why Leo's intuition cannot be constituted by relevant abstracta, we should assume that such a constitution relation is possible—and thus that the instantiation of Bengson's constitutive relation between an intuition and some relevant abstract entity is not enough to make the belief based on that intuition non-accidentally true in the sense we are using the term in GCO and MCO. While Bengson's constitutive explanation relates the knower's intuition to relevant facts, it does not do so in a way that adequately explains why any of the beliefs based on those intuitions are non-accidentally true.

V. Intellectual Perception

A. Intellectual Perception Introduced

Having seen many putative explanations fail, we might be tempted to concede to defenders of the Causal Objection that there is no non-causal explanation of our

property doing the ensuring. Thus, the fragile glass/vase case is an example of property inheritance, since the property of fragility is the same in both objects. In Bengson's constitutive explanation, the relevant property of the intuition stands in the ensurance (but not inheritance) relation to the relevant property of the abstracta, since "being not merely accidental" (the ensured property of the intuition) is not identical to the property of "being the way the world is" (the ensuring property of the fact) (see p. 31).

substantive a priori knowledge. Such surrender would seem to lead to a kind of radical skepticism. But it could lead in a quite different radical direction as well: to the postulation of causal connections in new and unexpected places, even to an account of substantive a priori knowledge that neatly parallels perceptual knowledge. In this section, I sketch an account of Intellectual Perception that fits this profile.

According to this account, intellectual perception has the same relata as Bengson's constitutive explanation; it relates our intuitions to relevant abstract entities. The crucial difference here is that the constitutive relation is replaced with a causal one. In instances of intellectual perception, relevant abstracta *cause* our intuitions.

Thus, intellectual perception has the same structure as physical perception. The shared structure includes: 1) a representational mental state, 2) the perceived fact, and 3) an appropriate causal relation between the mental state and the fact. For an example of physical perception: when you grasp a tomato, your feeling of the roundness of the tomato is a representational experience, the tomato is the fact that is grasped, and the tomato appropriately causes your experience via the nerves of your hand. Similarly, when you grasp a necessary truth such as that nothing can be both red and green all over, your intuition that the proposition is true is the representational experience, the relevant abstracta is the intuited fact, which appropriately causes your intuition.

To further elucidate my account of Intellectual Perception, I will push the analogy with perception two steps further. First of all, my account obviously implies something analogous to our sense organs—some kind of “intellectual sense organ” that converts the effects of abstracta into intuitions that represent the abstracta. Secondly, there is an illuminating analogy to be drawn between the physical movements that facilitate physical

perception and the mental movements that facilitate intellectual perception. In order to perceive a physical object, you often must strain to bring yourself into appropriate causal contact with it. To see a bird you might have to crane your neck. To grasp a ball you might have to stretch your arm. In my account of Intellectual Perception, entertaining a necessarily true proposition is an analogous form of exertion that, like these physical movements, is aimed at establishing an appropriate causal connection. For example, when you entertain DeMorgan's Law for the first time, that peculiar kind of intellectual exertion is, according to my account, an attempt to establish appropriate causal connection with the relevant abstract entity.

We should not let the extended analogy with perception induce a misleading mental picture. The relevant abstracta do not occupy a location in physical space from which they continuously radiate abstract effects, like an eternal cell phone tower. And the intellectual exertion that facilitates causal connection with them need not involve physically moving your head. Of course, abstracta have no spatial location, and thus do not send physical signals across a shared three-dimensional space. The kind of causal connection that I am proposing is one in which a kind of immaterial entity—namely, the relevant abstracta—acts on our minds. I claim that a certain focused way of thinking is what enables the abstract entity to have this effect.

B. Intellectual Perception Defended

By now, the general picture of this account of Intellectual Perception should be coming into focus. What is yet unclear is whether we have any reason to accept it. And what may not be entirely clear is how my account of Intellectual Perception grounds an adequate response to GCO or MCO. To show that this account of Intellectual Perception

grounds an adequate explanation of the striking fact that our a priori beliefs are true—and thus refutes GCO—I will need to defend both premises of the following argument:

C1. If relevant abstracta cause our intuitions, the a priori beliefs based on those intuitions are non-accidentally true.

C2. Relevant abstracta cause many of our intuitions.

C3. Therefore, the a priori beliefs based on those intuitions are non-accidentally true.

If I can defend C1 and C2, I will take myself to have provided an adequate explanation of the striking fact that many of our a priori beliefs are non-accidentally true, which would thus refute the more general causal objection. I offer, in what remains, what I consider to be an adequate defense of C1 and a partial defense of C2.

To defend C1, I offer an Argument from Analogy with perceptual knowledge. The terms of the analogy should be fairly obvious by now. Because of the appropriate causal connection between perceivers and physical objects, we are able to adequately explain why our perceptual beliefs are non-accidentally true. My account of intellectual perception posits a parallel causal connection between intuiters and relevant abstracta. If the explanatory demand is satisfied in the case of (physical) perceptual knowledge—and it is—then, given my account of intellectual perception, it is also satisfied in the case of substantive a priori knowledge.

The critic of MCO might claim that there is a special problem extending this account of intellectual perception to solve the causal objection to ethical non-naturalism. The objection claims, in other words, that even if this explanation refutes GCO, it does not refute MCO. Presumably the problem with extending the account is not the scarcity

of relevant moral abstracta. The moral domain includes an abundance of these: moral principles, moral properties, and moral propositions are all arguably abstract. The more pressing worry is that according to the details of my account, if we intellectually perceived moral facts, then the resulting moral metaphysics would cease to qualify as a version of non-naturalism, since it would imply that moral facts are causally active.

An analogy reveals why this objection misses the mark. Suppose that the Ontological Argument is sound and thus that the proposition that God exists is necessarily true. Suppose that, after extensive reflection, you finally grasp this necessary truth. On my account of Intellectual Perception, this implies that some relevant abstract entity (perhaps the necessarily true proposition that God exists) has caused your intuition. However, the truth of my account would not seem to imply that God is “natural” in any interesting sense of the term. If the grasping of theological truths would not serve to naturalize God, then grasping moral truths would not make moral facts natural.

The moral of this theological analogy may be stated as follows. If the downside of positing weird causal relata is that you incur a strong burden of proof, the upside, at least for a defender of non-naturalism, is that it makes the scope of causation so wide that participation in it no longer disqualifies an entity from its non-natural status. In other words: if abstract objects are causally active, then there is non-natural causation.

But the strong burden of proof on my account remains. In the rest of this section, I will discharge some of this explanatory burden by offering a partial defense of C2, the claim that relevant abstracta cause our intuitions. To begin with, I will defend the proposition against the charge that it is incoherent. According to this charge, the claim that abstract entities are causally impotent is an analytic truth. This charge might appeal

to a common characterization of an abstract object as a kind of non-spatial and non-causal entity.⁴²

However, this charge of incoherence is unsuccessful for the same reason that any attempt to draw a substantive metaphysical conclusion from a tautology fails. While it is true that abstract entities are often claimed to be causally impotent, their personalities are not exhausted by this negative quality. (If they were, abstract objects would be philosophically boring postulations.) On the contrary, abstract entities are alleged to do such positive things as *bear truth and falsity* (if propositions are abstract entities), *instantiate* into concrete particulars (if properties are abstract entities), and *constitute* higher-level abstract objects (if concepts constitute propositions), among other things. If it turns out that the entity that plays these roles is causally active, it still makes sense to call it abstract.⁴³ The proposition that is relevant to my account (viz., that the entity that plays these positive roles is causally potent) is not contradictory.

Defending the coherence of C2, however, is a long way from defending its plausibility. To be sure, a thorough assessment of the plausibility of this claim is well beyond the scope of this essay. Such an assessment would have to clarify and address larger metaphysical issues, such as the thesis of the causal closure of the physical. In lieu of a complete assessment of C2, I will perform a more modest supportive task of drawing attention to an unnoticed explanatory gap that my account of intellectual perception is able to fill. This surprising explanatory power is a theoretical virtue of my account.

⁴² See Lewis, "On the Plurality of Worlds," p. 83. See also section 3.2 ("The Causal Inefficacy Criterion") of the entry on "Abstract Objects" in the *Stanford Encyclopedia of Philosophy* (Rosen, Gideon, "Abstract Objects", *The Stanford Encyclopedia of Philosophy* (Spring 2012 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2012/entries/abstract-objects/>).

⁴³ And if it did not, we could systematically replace talk of abstract entities with talk of "schmabstract entities," stipulatively defined to exclude the commitment to causal impotence.

What is my account able to explain? The new explanandum is similar to the one we clarified near the beginning of the paper. There I drew attention to the striking fact that most of our moral beliefs are true. But there is a striking fact that is prior even to this one, a fact so familiar that its striking character often escapes notice. Simply stated, it is that the entertaining of necessary truths is regularly followed by intuitions and beliefs at all. Prior to the epistemic fact that many of these intuitions and beliefs are true (or non-accidentally true) is the psychological fact that intuitions regularly follow entertainings.

One might suspect that there is a straightforward psychological explanation for this familiar psychological fact. However, there is something about this explanandum that resists an easy explanation. In order to understand this, it will be helpful to step back and examine a familiar three-link chain of psychological events. When I entertain a necessarily true proposition, such as that nothing is both red and green all over (the first link), the entertaining is often followed by an intuition (the second), which is then followed by a belief that the proposition is true (the third link). In general, when thinking about necessary truths, the familiar flow of thought is from entertaining to intuition to belief.⁴⁴ In this sequence of psychological events, I draw your attention to an explanatory gap in the transition from the entertaining to the intuition that is not paralleled in the transition from intuition to belief.

We can highlight the relevant explanatory gap by reflecting on a distinction among these three attitudes: between two that are *truth-committed* and one that is *truth-neutral*. Intuitions and beliefs are truth-committed attitudes because they essentially

⁴⁴ Obviously entertaining is not always followed by intuition. For example, one's entertaining of a complex mathematical proposition might not be followed by any intuition at all. To develop my account, I would need to explain why certain complex propositions are "out of reach" to us in the sense that even our most strenuous intellectual exertions fail to establish an appropriate causal connection with the relevant abstracta (the effect of which would be an intuition).

involve a commitment to the truth-value (or lack thereof) of their propositional objects. You have an intuition *that* P is true (or false). And you have a belief *that* P is true (or false). With respect to a few exceptional propositions—such as those that involve referential failure—one might have the intuition that the proposition is neither true nor false. But it is incoherent to claim that you have an intuition or belief regarding P that does not involve any truth-commitment whatsoever. If I were to say that “I have an intuition about P,” without specifying whether the intuition is that P is true or false (or neither true nor false), then I would have left out an essential detail. The attitude of entertaining, however, is different. It does not imply any commitment towards the truth-value of the entertained proposition. *Qua* entertaining the attitude implies the absence of any truth-commitment. If I say that “I am entertaining P,” I have not omitted any detail. In contrast with truth-committed propositional attitudes, entertaining is truth-neutral.

While intuitions and beliefs are both truth-committed, they are, of course, not the same kind of attitude. For our purposes, it will not be necessary to devote much conceptual labor to demarcating these two kinds of truth-oriented propositional attitudes. But so that we can quickly call to mind the difference between intuitions and beliefs, I will say that intuitions are *presentational* while beliefs are *representational*; intuitions present propositions as being true, while beliefs represent them as true.

Because of this crucial difference between the two attitudes—intuitions present a truth-orientation while beliefs represent a truth-orientation—there is a plausible psychological explanation of why beliefs regularly follow intuitions. To explain why a belief that P regularly follows the intuition that P, we can cite the complementary attitudes of these two mental states along with a general disposition to attempt to form

true beliefs. Since most human beings aim (consciously or unconsciously) to form true beliefs, they will tend to represent as true, in their beliefs, what their intuitions present as true. Note that this psychological explanation does not imply that any of our intuitions or beliefs are accurate; even if they are systematically false, the explanation succeeds.

While there seems to be a straightforward psychological explanation of why the belief that P follows the intuition that P, there cannot be an analogous psychological explanation of why entertainings are regularly followed by intuitions or beliefs. Since entertaining is a truth-neutral mental state, there is nothing in the nature of the attitude we can appeal to that, in combination with a general tendency to attempt to form true beliefs, explains why the neutral entertaining P should be followed by a truth-committed intuition that P (or not P). An explanatory gap is implied by the shift from truth-neutral to truth-committed. To fill this gap, we will need to look outside of the psychological properties of the flow of thoughts whose striking familiar order we are trying to explain.

My account of Intellectual Perception fills this explanatory gap with something factive. The truth-neutral propositional attitude (i.e., the entertaining) is regularly followed by truth-committed propositional attitudes (intuitions and then beliefs) because, when we entertain necessary truths, we often grasp the truth. When this happens, according to my account, the relevant abstract entity causes our intuition. The factive abstract entity explains the transition from neutral entertaining to committed intuition.

This explanatory picture will come into clearer focus if we extend the analogy to perception one last time. When you look around, your mind is filled with sensory representations of whatever happens to surround you. These sensory images are definite and have obvious epistemic import. You seem to see a house and not a buffalo. I seem

to see my laptop and not Niagara Falls. That the physical movement of looking around is regularly followed by sensory images with definite contents is describable without mentioning the obvious explanation of why this is so (roughly, that looking around brings us into causal contact with physical objects). Now imagine how curious it would be if this same sequence held (i.e., looking around was followed by definite sensory images) even though you believed that you were entirely causally isolated from the physical world. We would want some explanation of why the neutral act of looking around was so regularly followed by these sensory images that presented the world as being definitely one way rather than another. This is the explanatory situation as regards the regular procession from entertaining to intuition. And just as a physical fact normally fills the gap between looking around and sensory experience, so, according to my account of Intellectual Perception, does a factive abstract entity plug the analogous gap.

In addition to grounding an adequate response to GCO and MCO, my account of Intellectual Perception has the power to explain this often ignored gap in the psychological explanation of a very familiar procession of thought. Each of these accomplishments adds to the overall explanatory power of the account.

V. Conclusion

It is time to review the dialectic we have drawn. According to ethical non-naturalism, moral facts are causally impotent (or at least, we can add now, they do not participate in any *natural* causal relations). According to the Causal Objection, this commitment of non-naturalism implies that there is no explanation for the striking fact that many of our moral beliefs are non-accidentally true.

This challenge to non-natural moral knowledge generalizes, I argued, to all substantive *a priori* knowledge. Because there is no causal explanation for our substantive *a priori* knowledge, it is on an explanatory par with moral knowledge.

After briefly arguing that there indeed must be an explanation for these striking facts about our *a priori* beliefs in general and our moral beliefs in particular, I considered a number of candidate explanations for these striking facts. The problem with explanations that merely relate different properties of the subject to each other is that they fail to relate any of these properties to the relevant moral facts. And the problem with explanations that merely relate different kinds of facts is that they leave out the subject.

Learning from these failures, I turned to consider putative explanations that satisfy the Knower/Fact Constraint. Evolutionary explanations satisfy the constraint but are unsuccessful. The failure of third-factor explanations is overdetermined: they make an illicit appeal to defeated first-order moral beliefs, and, even if they did not, they are unpromising because they do not generalize outside of the moral domain. By contrast, Wide Evolutionary Explanations generalize as far as the implied mechanism can reach. Since they do not specify a mechanism, however, they are basically vacuous if treated as responses to the Causal Objection.

Bengson's constitutive explanation satisfies the Knower/Fact Constraint. Moreover, the explanation is non-vacuous and generalizable across other *a priori* domains. Unfortunately, even if we admit that intuitions and abstracta can be constitutively related in the manner that Bengson suggests, there is no reason to think that the beliefs based on these intuitions are (in our sense) non-accidentally true. By simply replacing Bengson's constitutive relation with a causal one, however, we can explain not

only the striking epistemic fact that many of our a priori (and moral) beliefs are non-accidentally true but also the prior psychological fact that intuitions regularly follow entertainings at all. My account of Intellectual Perception does all of this.

Chapter 3: A Plantingian Pickle for a Darwinian Dilemma: Evolutionary Arguments Against Atheism and Normative Realism

In the last two decades, philosophers have developed evolutionary debunking arguments to target a range of cognitive domains. Surprisingly, two of the most widely discussed arguments have never been discussed at length together. I am referring to Alvin Plantinga's Evolutionary Argument Against Atheism¹ and Sharon Street's Darwinian Dilemma for Normative Realism.² While Plantinga and Street appeal to similar evolutionary considerations in their arguments, they do so in service of drastically different philosophical agendas; Plantinga concludes that it is irrational to believe atheism—a point which obviously favors theism—while the upshot according to Street is that we should embrace normative anti-realism.

That these two arguments have not been discussed together is not only surprising, it is also unfortunate, for, as I will argue, a proper comparison has a huge philosophical payoff: the demise of the Darwinian Dilemma. In this essay, I draw on such a

¹ For various refinements of the argument, see Alvin Plantinga, *Warrant and Proper Function* (Oxford: Oxford University Press, 1993), Ch. 12; Plantinga, *Warranted Christian Belief* (Oxford: Oxford University Press, 2000), pp. 227-240; Plantinga, "Introduction," and "Reply to Beilby's Cohorts," in J. Beilby (ed.), *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism* (New York: Cornell University Press, 2002), pp. 1-12, 204-275; Plantinga, *Where the Conflict Really Lies: Science, Religion, and Naturalism* (Oxford: Oxford University Press, 2011), Ch. 10; Plantinga, "Content and Natural Selection," *Philosophy and Phenomenological Research* 83 (2011), pp. 435-458.

Plantinga refers to his argument as the "Evolutionary Argument Against Naturalism." "Naturalism" is a thesis that is slightly broader than atheism; it is the claim that "there aren't any supernatural beings—no such person as God, for example, but also no other supernatural entities, and nothing at all like God" ("Introduction," p. 1). Since "naturalism" denotes a species of realism in metaethics (see Street, "The Darwinian Dilemma for Realist Theories of Value," p. 27), we will avoid confusion, with a very minor sacrifice of accuracy, by substituting the term "atheism" for "naturalism."

² Sharon Street, "A Darwinian Dilemma for Realist Theories of Value," *Philosophical Studies* 127 (2006), pp. 109-166. For further defense and refinement, see Sharon Street, "Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About," *Philosophical Issues* 18 (2008), pp. 207-228.

comparison to set up a dilemma for the Darwinian Dilemma for Normative Realism. (I will refer to the second-order dilemma as a “pickle” to distinguish it from the first-order one.)

To set up the pickle, suppose that Street’s evolutionary argument entailed Plantinga’s. That would be bad news for Street. For it would imply that her argument for normative anti-realism ultimately supports theism. But atheism is a suppressed premise of the Darwinian Dilemma as well as a commitment of almost all normative anti-realists. If Street’s argument entails theism, then the Darwinian Dilemma is internally incoherent and should be abandoned by almost everyone. Street’s argument succeeds, then, only if there is some good reason to reject Plantinga’s argument that does not count against Street’s own.

I argue that there is not. In rejecting Plantinga’s argument, Street has to make a claim that is either self-defeating—in the sense that it obliges Street to reject a parallel claim of her own argument—or elitist—in that it implies a state of diminished rationality for reflective novices. This is the pickle, I argue, for the Darwinian Dilemma.

I. Two Evolutionary Arguments

Before developing the pickle at greater length, I should introduce the two evolutionary arguments in a manner that emphasizes their underlying similarities. First, I introduce Plantinga’s Evolutionary Argument Against Atheism.

a. The Evolutionary Argument Against Atheism

In this argument, Plantinga begins by assuming atheism (henceforth, A) and a standard Darwinian account of human origins (henceforth, E, for evolution). He then

tries to show how these two propositions (A&E) generate an undefeated defeater for every one of the atheist's beliefs—including, crucially, the belief in A&E.

A core premise of Plantinga's argument is a conditional probability thesis. According to the Probability Thesis, the probability that our cognitive faculties are reliable is low, given A&E. With "R" standing for the proposition that our cognitive faculties are reliable, the Probability Thesis can be abbreviated as follows: Pr. (R/A&E) is low.

To defend the Probability Thesis, Plantinga begins by distinguishing between a belief's neurological properties and its content.³ Adopting a physicalist view of the mind, Plantinga assumes that neurological properties ground contents according to metaphysical or psychophysical laws.⁴ But—and this point is the crux of the argument—neurological properties ground contents *only at some high level of complexity and coordination*. Well before the neurological configurations would have reached this level (given what is probable on A&E), however, evolution would have been selecting for the neurological arrangements that, for entirely physical reasons, cause adaptive responses to ambient stimuli.⁵ And there is no good reason to think that the final modifications, upon

³ Here I summarize Plantinga's most recent defenses of the Probability Thesis (see *Where the Conflict Really Lies*, pp. 318-339, and "Content and Natural Selection"). But Plantinga's earlier defenses of the premise were different. In these, the low probability of R on A&E was partly explained by the possibility of characters such as "Paul the Hominid"—whose weird combinations of systematically false beliefs and strange desires motivated reproductively advantageous behavior. In Plantinga's story, Paul engages in the reproductively advantageous behavior of running away from a tiger because he believes this to be the best strategy for petting it (see *Warrant and Proper Function*, pp. 225-226; and "Introduction," in *Reply to Beilby's Cohorts*, p. 8). In his two most recent defenses of the Probability Thesis, however, Plantinga drops Paul the Hominid and other similar characters from the discussion and foregrounds the deeper worry that belief content probably would not affect behavior if A&E were true.

⁴ Plantinga allows that content might supervene on *disjunctions* of neurological properties (to accommodate multiple realizability) or *conjunctions* of neurological properties and environmental properties (to allow that meaning is not entirely "inside the head") (*Where the Conflict Really Lies*, p. 323; "Content and Natural Selection," p. 443).

which content finally begins to supervene on these arrangements, would often result in beliefs that are true. Instead:

“This new [content] will be implied with causal or metaphysical necessity by the relevant [complex neurological] property which, we may assume, is adaptive; but that doesn’t give us the ghost of a reason for assuming that the content thus accruing to the structure is *true*. Here natural selection is obliged to take potluck; it selects for adaptive [neurological] properties, but must then accept the content properties, true or false as the case may be, that supervene on them.”⁶

While the Probability Thesis is supported by evolutionary considerations, the other core premise of Plantinga’s argument is an *a priori* epistemological principle. The principle concerns the epistemic consequences of believing the Probability Thesis for the “Reflective Atheist.” By definition, the Reflective Atheist believes A&E and also accepts the Probability Thesis (presumably as a result of being impressed by Plantinga’s defense of it). According to this principle, if the Probability Thesis is true, then the Reflective Atheist has an undefeated defeater for R.

On behalf of this principle, Plantinga offers an argument from a series of analogies.⁷ In one of these, Sam believes that he has taken a pill that causes global cognitive malfunction in almost everyone who ingests it. Sam should believe, then, that the probability is high that it has caused global cognitive malfunction in him.⁸ Intuitively, Sam is no longer entitled to trust his cognitive faculties. And the epistemic position of the Reflective Atheist is relevantly similar to Sam’s: she also believes that her

⁵ Plantinga, *Where the Conflict Really Lies*, p. 324.

⁶ Plantinga, “Content and Natural Selection,” p. 444.

⁷ For this analogy and other similar ones, see Plantinga, “Reply to Beilby’s Cohorts,” pp. 206-208, 224, 240; *Where the Conflict Really Lies*, p. 341. In “Reply to Beilby’s Cohorts,” p. 240, Plantinga is explicit that his support for this premise is analogical: “I would argue that [the Probability Thesis] is a defeater for R by way of analogies.”

⁸ Plantinga, “Reply to Beilby’s Cohorts,” p. 206.

cognitive faculties are under the influence of a force that probably has rendered them unreliable. By analogy, the Reflective Atheist also has a defeater for R.

This thought experiment also helps us to see why the defeater of R is not easily defeated itself. The problem is that the sources of evidence to which we would normally appeal to defeat the defeater are blocked by its pervasive scope.⁹ After ingesting the pill, Sam's cognitive faculties might still seem to Sam to be working. Sam seems to remember a long track record of success with these cognitive faculties; Sam apparently has been forming true beliefs with his cognitive faculties for his entire life. But all of these sources of evidence—his seeming that his cognitive faculties are working, his apparent memories and new sensory experiences—are all products of the suspect cognitive faculties. Once he has reason to think that his cognitive faculties are probably unreliable—which generates the original defeater—he cannot then rely on these cognitive faculties to prove their reliability. The Probability Thesis is plausibly an *undefeated* defeater because it preemptively strikes down all defeater-defeaters in this fashion.

From these two premises, it follows that the Reflective Atheist has an undefeated defeater for R. We can reconstruct the argument we have developed so far:

- A1. The Probability Thesis: Pr. (R/A&E) is low.
- A2. If the Probability Thesis is true, then the Reflective Atheist has an undefeated defeater for R.
- A3. Therefore, the Reflective Atheist has an undefeated defeater for R.

These three propositions constitute the core of the Evolutionary Argument Against Atheism. To avoid the intermediate conclusion A3, which obviously counts against atheism, it is necessary for Street to reject A1 or A2. Since my focus will be on

⁹ Plantinga, "Introduction," pp. 11-12.

these first two premises, I will quickly run through the remainder of Plantinga's argument.

To get from A3 to a claim that fairly obviously implies that reflective belief in atheism is self-defeating, we will need to add:

A4. If the Reflective Atheist has an undefeated defeater for R, then she has an undefeated defeater for all of her beliefs—including her belief in A&E.

A5. Therefore, the Reflective Atheist has an undefeated defeater for all of her beliefs—including her belief in A&E.

And to get from here to the final conclusion that reflective belief in atheism is irrational, we can add:

A6. If the Reflective Atheist has an undefeated defeater for her belief in A&E, then reflective belief in atheism is irrational.

A7. Therefore, reflective belief in atheism is irrational.

b. The Darwinian Dilemma for Normative Realism

Having reconstructed the Evolutionary Argument Against Atheism, I now turn to consider our second evolutionary debunking argument: Sharon Street's Darwinian Dilemma for Normative Realism. Briefly: normative realism is the thesis that normative facts are "stance-independent" in the sense that they are not constituted by any agent's attitudes towards non-normative facts; normativity is "objective" in a robust sense.¹⁰

The dilemma for normative realism arises given the hypothesis that Darwinian forces had a tremendous influence on the contents of our normative judgments.¹¹ Either

¹⁰ Street, "A Darwinian Dilemma for Realist Theories of Value," pp. 2-4.

¹¹ Street qualifies the original hypothesis by claiming that the influence of Darwinian forces on the contents of our normative judgments was "indirect": what were directly selected for were "basic evaluative

these Darwinian forces stand in a truth-tracking relation to the (realistically-construed) normative facts or they do not.

On the first horn of the dilemma, the realist claims that the Darwinian forces do stand in this truth-tracking relation.¹² The problem with this horn, according to Street, is that it is implausible on empirical grounds. As an empirical hypothesis, the truth-tracking hypothesis competes with other empirical hypotheses about the evolution of our normative beliefs. According to the rival “Adaptive Link Account,” our normative judgments were selected only for their motivational effects and not, additionally, because they were true.¹³ Street claims that, because the Adaptive Link Account is simpler, clearer, and more illuminating than the truth-tracking hypothesis, it is superior according to standard criteria for evaluating empirical hypotheses.¹⁴

The other horn of the dilemma is to concede that these Darwinian forces did not track the truth of the normative facts.¹⁵ In that case, it would be a massive coincidence if many of our normative judgments were true. Street offers an analogy:

“On [the non-truth-tracking] view, allowing our evaluative judgments to be shaped by evolutionary influences is analogous to setting out for Bermuda and letting the course of your boat be determined by the wind and tides: just as the push of the wind and tides on your boat has nothing to do with where you want to go, so the historical push of natural selection on the content of our evaluative judgments has nothing to do with evaluative truth. Of course every now and then,

tendencies” in our evolutionary ancestors. Unlike these basic tendencies, “full-blooded evaluative judgments” are linguistically infused and are held reflectively. According to the qualified (indirect) hypothesis, we inherit these tendencies which in turn directly affect the contents of our normative judgments (*ibid.*, pp. 10-13).

¹² For Street’s discussion of this horn of the dilemma, see, *ibid.*, pp. 16-26.

¹³ *Ibid.*, p. 18.

¹⁴ *Ibid.*, p. 21.

¹⁵ For Street’s discussion of this horn of the dilemma, see, *ibid.*, pp. 13-16.

the wind and tides might happen to deposit someone's boat on the shores of Bermuda."¹⁶

While the conclusion of Plantinga's argument is that reflective belief in atheism is irrational, because it is self-refuting, the conclusion of Street's is that normative realism is implausible, because it implies either an unsupported empirical claim (i.e., the truth-tracking hypothesis) or a crippling epistemological problem (i.e., the coincidence).

The similarities between Street's Darwinian Dilemma and Plantinga's Evolutionary Argument Against Atheism are greater than is immediately evident; underneath Plantinga's probabilities and Street's coincidence is a common underlying structure. For this reason we can transpose Plantinga's argument into a parallel Darwinian dilemma for atheism or—as I will do here—reconstruct Street's argument to include a conditional probability thesis and a premise about defeaters.

The similar underlying structure is most apparent when we line up these two arguments on one particular point: Street's empirical case against the truth-tracking account and Plantinga's defense of the Probability Thesis. In this important segment of the arguments, each philosopher appeals to evolutionary considerations to support a similar kind of claim: that the mechanisms that explain what we believe within the target domains were not selected because any of the resulting beliefs are true. For Plantinga, certain neurological arrangements were selected only for their physical effects on our nervous systems—incidentally, some of these arrangements ground belief contents. For Street, certain psychological tendencies were selected only for their motivational effects—and incidentally, these tendencies predispose us to make normative judgments

¹⁶ Ibid., p. 13.

with particular contents. In neither scenario does evolution select for the correspondence of these contents with the relevant facts.

Once we have lined up the arguments on this point, we are in a good position to notice the close relationship between Street's coincidence claim and Plantinga's Probability Thesis. For Plantinga, it is the similar absence of a truth-tracking relationship between (i) the neurological properties that determine our beliefs contents and, (ii) the kinds of facts represented by the contents, that explains why the probability of R is low. And the low probability, in turn, explains why it would be Street's kind of coincidence if many of these beliefs were true. In general, the absence of the truth-tracking relation that Street describes explains the kind of low probability that Plantinga posits, which in turn explains why it would be a coincidence if many of the relevant beliefs were true. Each of these epistemic concepts constitutes a link on the same explanatory chain.

Plantinga follows the explanatory chain one link further—by specifying what happens to the justification of a set of beliefs when you believe that the cognitive faculties that produced them are likely to be unreliable. As the epistemological premise of his argument suggests: you acquire a defeater for your belief that the relevant cognitive faculties are reliable.

A plausible interpretation of Street's argument includes a corresponding claim about defeaters. In the passage that supports this interpretation, Street states:

“Barring such a coincidence, the only conclusion remaining is that many or most of our evaluative judgments are off track. This is the far-fetched skeptical result that awaits any realist who takes the route of claiming that there is no relation between evolutionary influences on our evaluative judgments and independent truths.”¹⁷

¹⁷ Street, “A Darwinian Dilemma for Realist Theories of Value,” p. 14.

In this passage, Street suggests that the normative realist who opts for this horn of the dilemma should believe that most of her normative judgments are “off track.” Surely this realization generates a defeater for the belief that one’s normative faculties are reliable.

One more parallel between the two arguments is implied by Street’s explanation of why the normative realist cannot restore trust in her normative faculties via “rational reflection” or “reflective equilibrium.”¹⁸ Recall that, in Plantinga’s argument, the defeat of R is undefeated because all of the cognitive faculties we would normally rely on to restore trust in our cognitive faculties are in a state of defeat themselves. This is a consequence of the all-encompassing scope of R.

Even though the defeater implied by Street’s argument is local—it leaves all of our non-normative faculties intact—these are worthless for restoring trust in our normative faculties, on the plausible assumption that you cannot derive (non-trivial¹⁹) normative claims from non-normative ones. As Street claims: “what rational reflection on evaluative matters involves, inescapably, is assessing some evaluative judgments in terms of others.”²⁰ Thus, the only kind of reflection that would promise to vindicate our normative faculties—reflection to the conclusion that some normative judgment or other is true—must enlist some other (discredited) normative judgment. Since trust in our normative faculties cannot be restored by rational reflection, the defeater of our

¹⁸ Ibid., p. 16.

¹⁹ A.N. Prior, “The Autonomy of Ethics,” in P. Geach and A. Kenny (eds.), *Papers in Logic and Ethics* (London: Duckworth, 1976), 88-96. Prior draws our attention to trivial counterexamples such as 1) Tea-drinking is common in England. Therefore, tea-drinking is common in England, or all New Zealanders ought to be shot. The conclusion is normative insofar as an “ought” appears in the disjunction. Clearly, however, the person who had reason to distrust her normative faculties would not be able to restore trust in them via this kind of inference from “is” to trivial “ought.”

²⁰ Street, “A Darwinian Dilemma for Realist Theories of Value,” 15.

normative faculties, which is implied by the cosmic coincidence, is plausibly an undefeated defeater.

Having noted these similarities, we are in a position to reconstruct Street's argument in a manner that resembles Plantinga's. What I will call the "Coincidence Thesis"—the parallel of Plantinga's Probability Thesis—claims that the probability is low that our normative faculties are reliable given normative realism and evolution (again, E). With "NR" referring to the claim that our normative faculties are reliable, we can state the Coincidence Thesis as follows:

B1. The Coincidence Thesis: Pr. (NR / normative realism & E) is low.

Now, we can add an epistemological principle that parallels Plantinga's defeater premise. In this premise, the "Reflective Normative Realist" is defined as the person who accepts normative realism and E and believes the Coincidence Thesis:

B2. If the Coincidence Thesis is true, then the Reflective Normative Realist has an undefeated defeater for NR.

B3. Therefore, the Reflective Normative Realist has an undefeated defeater for NR.

In the next section of the essay, I will consider whether Street can reject either core premise of Plantinga's argument without undermining B1 or B2. Before we get there, however, I will complete the reconstruction by showing how to get from B3 to the conclusion that normative realism is implausible:

B4. If the Reflective Normative Realist has an undefeated defeater for NR, then she has an undefeated defeater for all of her normative beliefs.

B5. Therefore, the Reflective Normative Realist has an undefeated defeater for all of her normative beliefs.

B6. If the Reflective Normative Realist has an undefeated defeater for all of her normative beliefs, then normative realism is implausible.

B7. Therefore, normative realism is implausible.

I will refer this argument as the *Reconstructed Darwinian Dilemma*, to distinguish it from Street's original formulation. Having reconstructed Street's original dilemma within the scaffolding of Plantinga's Evolutionary Argument Against Atheism, we are in a prime position to notice when an objection to Plantinga's argument implies an equivalent objection to Street's.

Before considering the relevant objections—which together make the pickle—it will be beneficial to address one point of potential confusion. Since the Reconstructed Darwinian Dilemma and Plantinga's Evolutionary Argument neatly parallel one another in the first five premises, we might wonder why they begin to diverge at the sixth premises—the former veering to the conclusion that normative realism is implausible, and the latter to the conclusion that reflective belief in atheism is irrational. One might have thought that the parallel arguments would either both support a kind of anti-realism or both support theism.

The discrepancy is easily explained when we uncover the suppressed premises of the arguments. In Plantinga's argument, a suppressed premise is the denial of a kind of global anti-realism. If reality is simply the projection of our cognitive faculties, then arguably the probability is high that our cognitive faculties are reliable, no matter how they were formed. Conversely, a suppressed premise of Street's argument is the denial of

theism, or more specifically the claim that god has engineered the cosmic coincidence she describes.

These suppressed premises suggest a broader point about this style of evolutionary argument. In general, there are two ways to save a cognitive domain from radical skepticism in light of the conviction that it would be a cosmic coincidence if our beliefs about that domain tracked relevant stance-independent facts. One way is to embrace anti-realism—by claiming that the facts are (in some sense) projections of our cognitive faculties. The other way is to accept the theological conviction that God has engineered the coincidence between our beliefs and the stance-independent facts. Without flagging the dismissed alternative, Plantinga favors theism over global anti-realism while Street opts for normative anti-realism over the postulation of a divine engineer.

While we could extend the parallel between the Evolutionary Argument Against Atheism and the Reconstructed Darwinian Dilemma by making these suppressed premises explicit, it would not be worth the effort. It will suffice to note that, when I claim that Street's argument supports theism, I am assuming (with Plantinga) that theism is more plausible than global anti-realism.

II. The Pickle

a. The poisonous end of the pickle

The Reconstructed Darwinian Dilemma supports theism unless Street can reject A1 or A2 of Plantinga's argument without undermining B1 or B2. I will consider, in reverse order, the implications of rejecting A1 and A2.

Suppose that Street rejects A2, which is, to recall:

A2: If the Probability Thesis is true, then the Reflective Atheist has an undefeated defeater for R.

The problem with swallowing this end of the pickle is that it forces Street to reject the parallel premise of the Reconstructed Darwinian Dilemma:

B2: If the Coincidence Thesis is true, then the Reflective Normative Realist has an undefeated defeater for NR.

In this section, I defend what I will call the “Poisonous Conditional”: if A2 is false, then B2 is false as well.

The primary reason for accepting the Poisonous Conditional is that A2 and B2 are motivated by parallel considerations. These have already been suggested. As we have already noticed, the antecedents of A2 and B2 (the Probability Thesis and the Coincidence Thesis, respectively) make a similar kind of claim. They both aim to show that Darwinian forces were unlikely to have tracked the truth of certain relevant facts. If the truth of one of these claims would generate a defeater for the relevant cognitive faculties, then the truth of the other one would, too. And again, parallel considerations motivate the claim that the relevant defeaters are undefeated defeaters. The defeater of R is undefeatable because of its global scope and the defeater of our normative beliefs is undefeatable because of the isolation of the normative realm vis-à-vis our non-normative faculties.

Because the two premises are supported by parallel considerations, most objections to A2 will apply, *mutatis mutandi*, to B2. For example, one kind of objection that has been raised against Plantinga’s argument as well as Street’s Darwinian Dilemma

appeals to the impressive amount of warrant we have for the beliefs in the reliability of the relevant cognitive faculties. As applied to Plantinga's argument, this objection claims that the warrant we have for the belief that R is "colossal."²¹ In comparison, our warrant for believing the Probability Thesis is weak. According to this objection, this disparity in warrant prevents the belief in the Probability Thesis from defeating the belief that R.

This objection might make the mistake of treating undercutting defeaters as rebutting defeaters; the amount of warrant for an undercutting defeater does not need to be as great as the warrant for the putatively defeated belief in order to undercut it. My main point here, however, is not to criticize the objection but rather to note its generalizability. If this objection succeeds in refuting A2, then a parallel objection succeeds in refuting B2. For the amount of warrant we have for the belief that our normative faculties are reliable is very impressive, too. Maybe not quite as great as the amount of warrant we have for believing R. But it is unlikely that the difference between "very impressive" and "colossal" will be enough to uphold the disparity between the two arguments that Street needs—unless, of course, there is a greater disparity between the quality of our evidence for the Probability Thesis and the quality of our evidence for the Coincidence Thesis. I delay discussion of these premises, however, until the next section.

While most objections to A2 apply, *mutatis mutandi*, to B2, there is one objection that applies uniquely to Plantinga's principle. This objection appeals to the global scope

²¹ Alvin Plantinga attributes this objection to William Alston, Timothy O'Connor, William Lane Craig, and "unnamed" others. (See the unpublished essay "Naturalism Defeated?," p. 16, available at http://www.calvin.edu/academic/philosophy/virtual_library/articles/plantinga_alvin/naturalism_defeated.pdf). Derek Parfit, *On What Matters*, Vol. 2 (Oxford: Oxford University Press, 2011), pp. 521-522, makes a similar point in a broad discussion of epistemological objections to normative realism.

of Plantinga's argument to show that it is self-defeating.²² Since this objection purports to undermine Plantinga's argument without posing any parallel threat to Street's argument, it poses some threat to the pickle, and thus we will need to consider it carefully.

According to this objection, Plantinga's argument is self-defeating because, once the atheist obtains a defeater for R, she obtains a defeater for her beliefs in the premises of Plantinga's argument. And once she loses her justification for endorsing Plantinga's argument, her justification for believing R—and, by extension, A&E—is restored. Understood as an objection to A2, this objection claims that, even if the Probability Thesis were true (as the antecedent of A2 claims), the Reflective Atheist would not have an undefeated defeater for R (as the consequent claims), because the defeater of R generates a defeater for itself.

In assessing this objection, the first thing to note is that the dialectical loop to which it draws our attention does not end with the restoration of R. Once the Reflective Atheist's trust in R is restored, then she is again in a position to form a justified belief in the Probability Thesis—and again acquire a defeater for R and, by extension, A&E. And, again, the defeat of R will undermine her justification for endorsing Plantinga's argument. Apparently, the Reflective Atheist is stuck in a kind of dialectical loop in which the defeater of R continues to generate a defeater for itself and then a defeater for that defeater, *ad infinitum*. If the defeater of R is not undefeated, as A2 states, in the sense of being *never* defeated (for it defeats itself), it is undefeated in the sense that it is

²² For the most thorough discussion of this objection, see Plantinga, "Naturalism Defeated?," pp. 55-58.

never *ultimately* defeated—for it always reappears, at every new iteration of the recursive epistemic structure, to defeat its defeater.

If believing the Probability Thesis plunges the atheist into this dreadful dialectical loop, then Reflective Atheism is not a stably rational position. It is not stably rational in the sense that the justification of belief in atheism is alternatively defeated, and restored, defeated, and restored, *ad infinitum*. I do not mean to imply that Reflective Atheists are rationally obliged to constantly change their doxastic attitude towards atheism from belief to disbelief, as though the recursive structure of the defeater loop were extended in time. The synchronic structure of the loop, however, does rationally oblige the atheist to take a skeptical doxastic attitude towards R—either to suspend judgment towards it or, at best, split the difference between suspending judgment and belief.

This last claim draws intuitive support from a slight extension of the analogous story of Sam. Sam is stuck in the same kind of dialectical loop as the Reflective Atheist. He has a vivid memory of having taken a pill that, he believes, probably has disabled his cognitive faculties. But if he believes that the pill has probably disabled his cognitive faculties, then he is not justified in trusting this memory. So then he shouldn't believe that he took the pill. But then he should trust his memory. The circle continues.

It seems fairly obvious that the rational response to this recursive loop is for Sam to suspend judgment about R or, at best, adopt some complicated semi-skeptical doxastic attitude to R. To claim that all of Sam's good reasons for his beliefs about the pill are self-defeating—and thus that Sam is justified in simply trusting his cognitive faculties again—is not skeptical enough.

This objection to A2 attempts to show that Plantinga's argument is self-defeating because the defeater of R defeats the justification for believing Plantinga's argument. While this is technically true, the larger dialectic of which this defeat is a part implies that reflective belief in atheism is less than fully rational. Like Sam in the story, the Reflective Atheist stuck in this loop is not entitled to simply believe R. Thus, this objection does not succeed in vindicating the rationality of reflective belief in atheism.

Because A2 and B2 have similar motivations, and most objections to A2 have a parallel to B2, and the one that does not fails to provide a robust vindication of atheism, the Poisonous Conditional is plausible: if A2 is false, B2 is false as well. To protect B2, then, Street must accept A2.

b. The sour side of the cucumber

If Street accepts Plantinga's epistemological principle (A2), then she will have to attack the Probability Thesis (A1) in order to avoid an intermediate conclusion (A3) that would be devastating to atheism. At first blush, this line of attack seems promising. Certainly, the quasi-empirical argument that Plantinga offers on behalf of the Probability Thesis is different from the one Street offers in her defense of the Coincidence Thesis (B1). In his defense of the Probability Thesis, Plantinga appeals to considerations from the philosophy of mind and evolutionary theory to support the hypothesis that evolution selects for neurological properties with no regard for content. This, he takes it, is the most plausible view of evolution's cognitive effects once we adopt the view that god does not exist.

There are many ways of objecting to Plantinga's support for the Probability Thesis that do not threaten the empirical grounds of the Coincidence Thesis.²³ I thus concede that it is consistent to reject the Probability Thesis while maintaining the Coincidence Thesis.

If Street accepts A2, though, then her denial of the Probability Thesis carries one implausible implication. As I shall argue, it sanctions a kind of "epistemic elitism" according to which rational belief in atheism depends on a combination of scientific and philosophical expertise—of the kind that is needed to assess Plantinga's quasi-empirical argument for the Probability Thesis.

In this context, it is not *ad hominem* to note that the best critics of the Probability Thesis are relative experts in philosophy or biology. Even if these experts are justified in rejecting the Probability Thesis, it does not follow that novices, who by definition are not fluent with the relevant evidence, are also justified in rejecting this controversial premise.

To the contrary, I contend that a large class of novices should suspend judgment about the Probability Thesis if they ever entertain it. The relevant novices are atheists who are not well informed about evolutionary biology or philosophy of mind.²⁴ For my purposes, it is not important to demarcate the line between "novice" and "expert" in any precise way. What is important is that the class of novices is non-negligibly large. And, with any reasonable demarcation, it is. Many people do not have the kind of expertise required to assess the Probability Thesis. Let us refer to members of this class of

²³ For critiques of the Probability Thesis, see Brandon Fitelson and Elliott Sober, "Plantinga's Probability Arguments Against Evolutionary Naturalism," *Pacific Philosophical Quarterly* 79 (1998), pp. 115-129; Evan Fales, "Darwin's Doubt, Calvin's Calgary," in Beilby (ed.), *Naturalism Defeated?*, pp. 43-58.

²⁴ The relevant novices also lack testimonial evidence that the Probability Thesis is false, or possess such evidence but are not justified in regarding it as determinative of what they ought to believe.

uninformed novices who actually entertain the Probability Thesis as “Reflective Novices.”

In defense of the claim that Reflective Novices ought to suspend judgment about the Probability Thesis, I offer an analogy. Consider some proposition about the phylogeny of Walruses: for example, the claim that walruses are descended from seals. If you don't know much about evolutionary biology, then it is clear that you should suspend judgment about this proposition. You have no idea. The Probability Thesis is a relevantly similar proposition. It concerns not the evolutionary history of walruses but rather the much more complicated evolutionary history of our brains and minds. As Plantinga's defense of the Probability Thesis indicates, a proper assessment of the proposition requires expertise in both evolutionary biology and the philosophy of mind. Lacking such expertise, the Reflective Novice ought to suspend judgment on the Thesis.

Suspending judgment, I take it, is a species of not believing. As such, it is different from either believing or disbelieving a proposition. As a species of not believing, it is different from the kind of not believing that is entailed simply by never having considered a proposition. Instead, it is a kind of considered withholding of judgment.

What are the epistemological consequences of suspending judgment about the Probability Thesis? Remember that, when exploring this side of the pickle, it is fair to assume that Plantinga's epistemological premise (A2) is true. This premise implies that *believing* the Probability Thesis generates an undercutting defeater for R (for the proponent of A&E). As a general rule about defeaters, though, S's *believing* some proposition generates an undercutting defeater for S's belief that Q if and only if S's

suspending judgment about that proposition generates something nearly as bad as a defeater for S's belief that Q: what I will call a "diminished defeater." Some examples will be helpful here.

In the first one, I adapt a classic barn façade case.²⁵ In the original version, you are driving down the highway, and seem to see a barn, and so form the belief that there is a barn. But then, on good evidence, you come to believe that the highway is lined with barn facades. This belief defeats your justification for believing that there is a barn.

Modifying the example, suppose that you do not simply believe, but rather suspend judgment about, the proposition that the highway is lined with barn facades. Maybe the evidence that bears on this proposition is ambiguous or (for whatever reason) just too complicated for you to assess. Intuitively, such suspension does damage the epistemic status of your belief that there is a barn.

Consider another modified classic example.²⁶ In the original classic, you go into a factory and seem to see that a bunch of widgets are red. But then the supervisor tells you that they are being irradiated under red light (it helps the company detect hairline fractures) and thus will appear red even if they are not. If you believe the supervisor, you obviously obtain a defeater for your belief that the widgets are red.

Adapting the example, however, suppose that you do not simply believe the supervisor. Maybe there is another supervisor you have equal reason to trust. She assures you that the widgets are not being irradiated today. Since you don't know which supervisor to believe, you ought to suspend judgment about whether the widgets are

²⁵ Alvin Goldman, "Discrimination and Perceptual Knowledge," *The Journal of Philosophy* 73 (1976), pp. 771-791, at p. 772.

²⁶ Plantinga, *Warrant and Proper Function*, p. 230, adapts this example from John Pollock, *Contemporary Theories of Knowledge*, (Totowa: Rowman and Littlefield, 1986), p. 38.

being irradiated. Having to suspend judgment towards this proposition undercuts your justification for believing that the widgets are red. You would be irrational to persist in this belief.

That said, you may be justified in taking a more fine-grained doxastic attitude towards this proposition (e.g., you don't simply believe that the widgets are red, but you "half-way believe" it with a credence level (say) of .25); or you may be justified in some similar belief with a finer-grained content (e.g., you don't believe simply that the widgets are red, but that there is a near 50% probability that they are red). Either way, it is clear that the mandate to suspend judgment leaves you with much less positive epistemic status than you began with. To flag the difference in epistemic impact between belief and suspension of judgment towards the same proposition, I will refer to the defeater that results from the latter as a "diminished defeater."

If the Reflective Novice suspends judgment about the Probability Thesis, then she will acquire one of these diminished defeaters for her belief that R is true. And since all of the Reflective Novice's beliefs result from her cognitive faculties, this diminished defeater will generate a diminished defeater for all of her beliefs—including the belief in A&E. The Reflective Novice's belief in atheism will not be strictly self-defeating but rather "self-diminishing," in the sense that appropriate reflection on the Probability Thesis should generate a diminished defeater for her belief that A&E.

But it is implausible—I call it "epistemic elitism"—to claim that novices occupy a epistemological position so precarious that, just by thinking about Plantinga's argument, they can lose their justification for all of their beliefs. And that is what is mandated by the foregoing considerations, in tandem with A2.

To protect the rationality of novices, Street needs to reject Plantinga's epistemological principle (A2). She needs to claim that the Probability Thesis is not the sort of proposition that generates defeaters (if one believes it) or diminished defeaters (if one suspends judgment on it). But this novice-friendly option for rejecting Plantinga's argument we already excluded in the previous section when we considered "the poisonous end of the pickle."

III. Conclusion

Remember the original worry: Street needs to be careful that her argument does not generalize over Plantinga's argument, and thus supports theism (if not global anti-realism). To avoid such overgeneralization, Street needs to reject either Plantinga's epistemological principle (A2) or his Probability Thesis (A1). She should not reject A2. For, as I have argued, parallel considerations support A2 and B2. If A2 is false, then B2 is false as well. Rejecting A2 is "poisonous" in the sense that it undermines B2 of the Reconstructed Darwinian Dilemma.

So Street should accept A2. But once she has accepted A2, there is something "sour" about the remaining strategy for resisting Plantinga's argument (i.e., by rejecting A1). Even if experts have good reason to reject A1, novices do not. Ignorant of biology and philosophy, Reflective Novices should suspend judgment about A1. But if believing the Probability Thesis generates a defeater for R—as A2 claims—then suspending judgment about it generates a diminished defeater for R. Thus if Street accepts A2, then she makes a commitment which implies that the rationality of novice atheists is extremely fragile: they cannot even entertain Plantinga's argument without having the

rationality of all of their beliefs diminished. Only expert atheists can consider Plantinga's argument and leave with their rationality intact. Such epistemic elitism is surely unacceptable.

Because there is no way to reject the Evolutionary Argument Against Atheism that does not have some implausible implication for the Reconstructed Darwinian Dilemma, we should conclude that the success of the latter argument implies the success of the former. If evolutionary considerations support normative anti-realism, then they also support theism, and if they support theism, then they do not support normative anti-realism. Therefore, evolutionary considerations—of the sort Street identifies at least—do not support normative anti-realism. The philosophical pay-off of comparing Plantinga's and Street's evolutionary arguments is the discovery of a new strategy by which the normative realist can resist Street's Darwinian Dilemma.

Bibliography

- Audi, Robert. *Moral Perception*. Princeton: Princeton University Press, 2013.
- Ayer, A.J. *Language, Truth, and Logic*. New York: Dover Publications, 1952.
- Bedke, Matthew. "Intuitive Non-Naturalism Meets Cosmic Coincidence." *Pacific Philosophical Quarterly* 90 (2009), 188-209.
- Beilby, James, ed. *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. New York: Cornell University Press, 2002.
- Bengson, John. "Grasping the Third Realm." Forthcoming in *Oxford Studies in Epistemology*, Vol. 5, edited by T. Szabo Gendler and J. Hawthorne.
- Bonjour, Laurence. *In Defense of Pure Reason*. Cambridge: Cambridge University Press, 1998.
- Casullo, Albert. *A Priori Justification*. Oxford: Oxford University Press, 2003.
- Clarke-Doane, Justin. "What is the Benacerraf Problem?" Forthcoming in *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, and Infinity*, edited by F. Pataut.
- Cowan, Robert. "Perceptual Intuitionism." *Philosophy and Phenomenological Research* 86 (2013), 1-30.
- Cullison, Andrew. "Moral Perception." *European Journal of Philosophy* 2 (2010), 159-175.
- Cuneo, Terence. "Reidian Moral Perception." *Canadian Journal of Philosophy* 3 (2002), 229-258.
- Dretske, Fred. *Seeing and Knowing*. London: Routledge and Kegan Paul, 1969.
- Enoch, David. "The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope With It." *Philosophical Studies* 148 (2010), 413-438.
- Fales, Evans, "Darwin's Doubt, Calvin's Calgary," In *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. New York: Cornell University Press, 2002.
- Field, Hartry. *Realism, Mathematics, and Modality*. Oxford: Basil Blackwell, 1989.

- Fitelson, Brandon and Elliott Sober, "Plantinga's Probability Arguments Against Evolutionary Naturalism." *Pacific Philosophical Quarterly* 79 (1998), 115-129.
- Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis* 23, 1963.
- Goldman, Alvin. "Discrimination and Perceptual Knowledge," *The Journal of Philosophy* 73 (1976), 771-791.
- Huemer, Michael. *Ethical Intuitionism*. New York: Palgrave Macmillan, 2008.
- Joyce, Richard. *The Evolution of Morality*. Cambridge: MIT Press, 2007.
- Lewis, David. *On the Plurality of Worlds*. Oxford: Blackwell Publishers, 1986.
- Mackie, John. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin (1977), 38-42.
- McBrayer, Justin. "Moral Perception and the Causal Objection." *Ratio* 23 (2010), 159-175.
- "A Limited Defense of Moral Perception." *Philosophical Studies* 149 (2010), 305-320.
- McGrath, Sarah. "Moral Knowledge by Perception." *Philosophical Perspectives* 18 (2004), 209-228.
- Parfit, Derek. *On What Matters*. Oxford: Oxford University Press, 2011.
- Plantinga, Alvin. *Warrant and Proper Function*. Oxford: Oxford University Press, 1993.
- *Warranted Christian Belief*. Oxford: Oxford University Press, 2000.
- "Introduction." In *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. New York: Cornell University Press, 2002.
- "Reply to Beilby's Cohorts." In *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. New York: Cornell University Press, 2002.
- *Where the Conflict Really Lies: Science, Religion, and Naturalism*. Oxford: Oxford University Press, 2011.
- "Content and Natural Selection." *Philosophy and Phenomenological Research* 83 (2011), 435-458.

- Pollock, John. *Contemporary Theories of Knowledge*. Totowa: Rowman and Littlefield, 1986.
- Prior, A.N. "The Autonomy of Ethics." In P. Geach and A. Kenny (eds.), *Papers in Logic and Ethics*, edited by P. Geach and A. Kenny. London: Duckworth, 1976.
- Ross, David. *The Right and the Good*. Oxford: Oxford University Press, 1930.
- Shafer-Landau, Russ. *Moral Realism*. Oxford: Oxford University Press, 2003.
- "Evolutionary Debunking, Moral Realism, and Moral Knowledge." *Journal of Ethics and Social Philosophy* 7 (2012), 1-37.
- Schechter, Joshua, "The Reliability Challenge and the Epistemology of Logic." *Philosophical Perspectives* 24 (2010), 437-464.
- Shaffer, Michael and Michael Veber, *What Place for the A Priori?* Chicago: Open Court, 2011.
- Siegel, Susanna. "Which Properties are Represented in Perception?" In *Perceptual Experience*, edited by T. Szabo Gendler and J. Hawthorne. Oxford: Oxford University Press, 2006.
- Skarsaune, Knut. "Darwin and Moral Realism: Survival of the Fittest." *Philosophical Studies* 148 (2010), 413-438.
- Smith, Michael. *The Moral Problem*. Oxford: Blackwell Publishing, 1994.
- Street, Sharon. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127 (2006), 109-166.
- "Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About." *Philosophical Issues* 18 (2008), 207-228.
- Sturgeon, Nicholas, "Moral Explanations." In David Copp and David Zimmerman (eds.), *Morality, Reason, and Truth*, edited by David Copp and David Zimmerman. Totowa, NJ: Rowman & Allanheld, 1985.
- Väyryen, Pekka. "Some Good News and Some Bad News for Ethical Intuitionism." *Philosophical Quarterly* 58 (2008), pp. 489-511.
- Wielenberg, Erik. "On the Evolutionary Debunking of Morality," *Ethics* 120 (2010), 461-464.