

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Why Answer the Epistemic Challenge?

Permalink

<https://escholarship.org/uc/item/7384w2p9>

Author

Pickering, Kirsten Marie

Publication Date

2020

Peer reviewed|Thesis/dissertation

Why Answer the Epistemic Challenge?

By

Kirsten Marie Pickering

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Philosophy

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jay Wallace, co-chair

Professor Niko Kolodny, co-chair

Professor Kathryn Abrams

Fall 2020

© Copyright by Kirsten Marie Pickering 2020

All Rights Reserved

Abstract

Why Answer the Epistemic Challenge?

by

Kirsten Marie Pickering

Doctor of Philosophy in Philosophy

University of California, Berkeley

Professors Jay Wallace and Niko Kolodny, co-chairs.

I describe a debate among metaethicists about how to explain our reliability about value, showing that it is no mere coincidence. The “epistemic challenge” is to give such an explanation.

It can appear that some metaethical views—that is, views about what values and norms are, most basically, like—have an easier time than others in answering the epistemic challenge. This makes the challenge relevant to evaluating such views, since metaethicists agree that successfully giving this explanation is one demand on a satisfying metaethical theory. This consensus is premature, however. Metaethicists have not adequately explained why we should want an explanation of our evaluative reliability at all. Until we have clarified whether and why an explanation of our evaluative reliability is worth having, we should not make such an explanation a goal of a metaethical theory.

Metaethicists sometimes suggest that, unless there is such an explanation, knowledge and epistemically justified beliefs about values and norms would be impossible. However, they are not explicit about why achieving knowledge and epistemic justification in evaluative thought matters. To make this explicit, we should focus on valuable features associated with knowledge and epistemic justification. Our question is whether our evaluative beliefs would lose these features if we lacked an explanation of our evaluative reliability. I consider three such features: the admirability of our evaluative beliefs, their stability, and their reasonableness. In each case,

I argue that the relevant feature is not threatened by a failure to answer the epistemic challenge.

Even if a failure to answer the challenge does not threaten the valuable features of our evaluative beliefs, we might still have a reason to give the sort of explanation desired by the challenger—an explanation that makes this reliability no coincidence. I consider two arguments on behalf of the epistemic challenger for the claim that it is good to have such an explanation. On the one hand, commitment to falsehoods is a bad thing, and we consider some things highly likely to have an explanation. If our evaluative reliability is such a thing, then metaethicists have a reason to answer the epistemic challenge, because otherwise they would be committed to a likely falsehood. I argue, however, that no one is yet in a position to say that our evaluative reliability is highly likely to have an explanation. So we cannot yet claim that this is a reason to explain our evaluative reliability. On the other hand, we often take explanations to have great intellectual value. It seems straightforward to conclude that if we cannot answer the epistemic challenge, we miss out on an explanation, and so on something of intellectual value. I argue that this is not as obvious as it seems, by showing that on one prominent view of explanatory value, we cannot draw this conclusion. If the value of explanation lies in unification, as this view suggests, it is not obvious that explaining our evaluative reliability in a way that makes this “no coincidence” would yield greater value than an alternative explanation. To appeal to explanatory value to motivate answering the challenge, we cannot simply invoke that value, but must show how the relevant explanation better promotes that value than its competitors.

One result of this project is that, for all we have seen, a key question that has interested metaethicists may have no importance. We may have no reason to engage with the epistemic challenge. This should inspire further thought about the proper tasks of metaethics, and about what we should want in a metaethical theory.

For Jim and Yuan, and in memory of Jeffrie Murphy

Contents

Acknowledgements	iii
Preface	v
Evaluative beliefs	v
The epistemic challenge	vii
Appeal to the value of evaluative belief	viii
Appeal to the value of explanation	ix
Some conclusions	xi
1 The Epistemic Challenge Meets the Value Problem	1
1.1 Realism and the epistemic challenge	1
1.2 What is at stake in the epistemic challenge?	6
1.3 The epistemic challenge meets the value problem	18
1.4 Conclusion	25
2 The Epistemic Challenge and Avoiding Accidents	27
2.1 The value of avoiding accidentally true belief	27
2.2 Less worthy of admiration	29
2.3 More vulnerable to loss	37
2.4 More unreasonable	42
2.5 Conclusion	51
3 The Epistemic Challenge and the Value of Explanation	52
3.1 The demand for explanation	52
3.2 Explaining correlations	55
3.3 There's surely some explanation!	57
3.4 The value of explanation	66
3.5 Conclusion	82
Bibliography	83

Acknowledgements

Thanks to my father, Paul Pickering, for valuing my opinions since I was old enough to express them, even when we disagreed. Without our conversations about political philosophy, religion, and how to resolve disagreements about values, I would not have studied philosophy.

Thanks to my mother, Lisa Pickering-Scott, for anything creative I have done in philosophy. Doing creative work requires bravery, and she showed me how to be brave.

For their love and support, thanks to Catherine Pickering, Harry Scott, Jaclyn Vogel, Kelly “Jordan” Scott, Marilyn Scott, Sarah (Pickering) Woodland, Stuart Pickering, and Susie Hutchinson.

I learned a lot from my teachers at Mountain View High School. Thanks to my history teacher, Rob Helsel, for his incredible teaching and his personal encouragement. For her good advice and perceptive criticism, thanks to my French teacher, Amy Hull. Thanks to Ryan Davis, my debate coach and friend, for introducing me to the work of Christine Korsgaard and to the field of philosophy.

As an undergraduate at Arizona State University, I was incredibly lucky to be the student of Jeffrie Murphy, Michael J. White, and Cheshire Calhoun. In their very different ways, they each showed me how to be a good philosopher and how to take pleasure in philosophy. (This dissertation is dedicated to the memory of Professor Murphy, who died this past September. His wit in conversation and relish for literature and philosophy were a delight, as was his low tolerance for bullshit. Our lunches

and meetings were highlights of my intellectual education.) In college, I was also lucky to talk about philosophy with Svetlana Turetskaya, Urusa Alaan, and Bob Zeller and his family. My thanks to all of them.

As a graduate student, it has been a privilege to belong to the philosophical community at Berkeley. For their conversation and example, thanks to Alex Kerr, Barry Stroud, Caitlin Dolan, Clara Lingle, Emily Perry, Janum Sethi, John MacFarlane, Joseph Bjelde, Katie Mantoan, Lisa Clark, Pia Schneider, Ravit Dotan, Sophia Dandele, Tyler Haddow, Umrao Sethi, and Véronique Munoz-Dardé. Thanks especially to John Campbell, for his help and his infectious enthusiasm for philosophy, and to Hannah Ginsborg and Daniel Warren, for accosting me in the Louvre to give me good advice when I still wasn't sure I would finish my dissertation. And thanks most of all to Katharina Kaiser, for her mentorship and encouragement. Her good advice and our discussions of philosophy have shaken me out of my funk many times.

Thanks to my advisors Jay Wallace and Niko Kolodny, for their insights and encouragement over my years in the program. Without their interest, patience, and good humor, I would not have begun, and I certainly would not have finished, this dissertation.

I have also been privileged to belong to another college, serving the students of San Quentin State Prison. Thank you to the staff, students, and Board of Mount Tamalpais College, especially Connie Krosney, Gary Malachi Scott, James King, Jody Lewen, Kathy Richards, Tim Thompson, and Theresa Roeder. You showed me how education can be liberating.

In my meandering career at Berkeley, I often despaired of finishing my dissertation. This was due in part to writing anxiety, and in part to the depression that comes from being stuck. To everyone who helped me in struggling against them, including many mentioned here, thank you.

Finally and most of all, thank you to Jim Hutchinson and Yuan Wu, who listened with interest to every new idea, read every word (usually more than once), helped me again and again to write, and inspired me in countless ways to do better. I cannot describe all that I owe them. But I dedicate this dissertation to them.

Preface

Evaluative beliefs

In ordinary life, questions about value come up constantly. Is *It's A Wonderful Life* a great movie, or not? Would it be wrong to lie to this friend about my plans, in order to protect my privacy? Is it more important to show up to the union meeting tonight, or to stay in and work on this philosophy paper?

All of these are questions about value. They are questions about the evaluative features—*aesthetic, moral, political, intellectual*—of particular objects, actions, and relations in the world. Since similar concrete questions come up again and again, one can get a head start in addressing them by thinking about more abstract evaluative questions. What features make for a great film? Why is it *ever* important to tell the truth? What is valuable about involvement in politics, and how should one weigh this value against a value like understanding? When we make up our minds about either concrete or abstract questions of value, I call this forming an “evaluative belief.”

When faced with evaluative questions of either a concrete or an abstract kind, you might or might not make up your mind and form an evaluative belief about that question. The question might be too difficult. You might instead make some progress by settling a related question—for instance, whether the conflict between public responsibility and personal freedom faced by Jimmy Stewart's character is *interesting*. Or you might lose the moment for making up your mind. You are distracted, and the

time for setting out for the union meeting comes and goes. So you opt for philosophy, but not because you are convinced of its greater importance.

Failing to make up your mind about an evaluative question you face is not always a bad thing. However, it *can* matter quite a lot to have settled evaluative beliefs. Our convictions about value can help us to resist ordinary and extraordinary temptations. Such resistance is not just a matter of willpower, but of having a clear view of the values at stake in what we do. It matters that we do the right thing, and that we achieve something of value in life. In some cases, the only way to do the right thing, or to achieve something of value, is by having a true evaluative belief.

When we set out to form a belief about some evaluative question, we need at least some trust in our own abilities to arrive at the truth. We need to set aside, at least for a time, the fear that our current evaluative outlook is too deeply mistaken for any amount of reflection to put us on the right track. Even when we do not set out to answer such a question explicitly, but simply form evaluative beliefs as needed in daily life, we implicitly trust these abilities. We take ourselves to be, in some sense, reliable in matters of value.

In many contexts where it matters, each of us does question our evaluative reliability. We may come to doubt our impartiality on questions of justice, when these have consequences for our best friends. We may realize that we don't really *get* why we ought to be fair-minded in debates with alienating political opponents, or what is so bad about pursuing a romance with someone under our supervision. In cases like these, conversations with others that we trust—who themselves seem reliable, over time, about these sorts of evaluative questions—can help us to arrive at true evaluative beliefs. (Hard and fast rules, when framed well, can also help.)

But what if there is no real help to be had? What if we are all too far off track, from the beginning, to have a hope? There are many routes to a wholesale worry about our evaluative reliability. This dissertation focuses on one route taken in recent philosophical discussions, the “epistemic challenge.”

The epistemic challenge

In chapter one, I describe a debate in metaethics about our evaluative reliability. By “evaluative reliability,” I will mean our having evaluative beliefs that are largely true. All sides in this debate aim to vindicate the assumption that we are, in fact, evaluatively reliable. Their dispute is over the question: how could such reliability be explained?

When it comes to our beliefs about our physical surroundings, we can explain relatively easily why those beliefs are reliable: we form them by relying on *perception*, and perception puts us in causal contact with our physical surroundings. My belief that there is a table in front of me was formed based on seeing that table, and I see that table through causal contact with the table in front of me: I believe this because it is so. Such a story, however, does not seem available for evaluative beliefs. So, participants in this debate ask, how could they be reliable?

A *complete* explanation of our reliability about our physical surroundings would also need to explain why we have the reliable perceptual faculties that we do. The most promising explanation is clearly an evolutionary one. Here too, it seems that we will have a satisfying explanation. Roughly, we developed reliable perceptual capacities because reliability about the physical facts was *adaptive*: that is, conducive to survival and reproduction. Once again, it seems that our having reliable capacities to get onto the evaluative facts—however exactly these capacities work—will not admit of a parallel evolutionary explanation. While these capacities make a difference to our survival and reproduction, it is hard to see how evaluative reliability would be adaptive, as compared with whatever evaluative psychology leads to taking actions that maximize one’s chances of surviving and reproducing.

The so-called “epistemic challenge” is to give an explanation of our evaluative reliability.¹ Metaethicists think that our prospects for doing so depend crucially on which metaphysical account of the evaluative facts we

¹The challenge is called “epistemic,” both because it concerns the *truth* of our beliefs, and because it is thought to pose a threat to their status as *knowledge* or as *epistemically justified*.

accept: what we think these facts are, most basically, like. For example, if we hold that evaluative facts cannot *cause* anything, unlike facts about tables, then we should not give an account of our evaluative reliability on which we perceive those facts. Metaethicists agree that the demand to explain our evaluative reliability has implications for which view of evaluative facts we should hold: views that can satisfy this demand are, at least in this respect, better.

In chapter one, I argue that this consensus is premature. Metaethicists have not adequately explained why we should *want* an explanation of our evaluative reliability *at all*. The conclusion of chapter one is that until we have clarified whether and why an explanation of our evaluative reliability is worth having, we should not make such an explanation a goal of a metaethical theory.

Appeal to the value of evaluative belief

In chapter two, I start considering why we might want an answer to the epistemic challenge. I begin from the observation that many of those who seek to address the challenge—the “epistemic challengers”—relate the desire for this explanation to the epistemic features of our evaluative beliefs themselves. That is: they suggest that unless there is such an explanation, we will lack evaluative knowledge or justified evaluative belief. Of course, this will only matter if these things are *good*. Our question is whether we lose anything good if we do not answer the challenge: for instance, do we lose a valuable sort of epistemic status?

My task in chapter two is to consider the features that are thought to make knowledge and justification good, and to see whether our beliefs would lose these features if we lacked an explanation of our evaluative reliability: if our evaluative reliability was, in this sense, an accident. (Since philosophers have traditionally associated knowledge and epistemic justification with believing the truth non-accidentally, it seems a real possibility that our beliefs fail of knowledge, or of justification, if our reliability is an accident.)

The first feature I consider is the *admirability* of our evaluative beliefs. When beliefs count as knowledge, or as justified, they are typically thought to be more admirable than beliefs that are true by accident. I argue that given the particular *kind* of accident at issue in the epistemic challenge, even if our beliefs were only accidentally true in this sense, it would not threaten their admirability.

The second feature I consider is the *stability* of our true evaluative beliefs: our tendency not to give them up. When one learns that one's belief, if true, is only true by accident, one often *gives up* that belief. By the same token, if one learns that someone else's belief, if true, could only be true by accident, one will tend not to rely on that person. One might lose out, in that way, on gaining a true belief. I argue, however, that it is not psychologically plausible that coming to appreciate the unanswerability of the epistemic challenge should undermine the stability of our beliefs in either of these ways. Given the centrality of evaluative belief in human life, there is no danger that we will cease to hold such beliefs through recognizing that we are unable to explain our evaluative reliability.

Finally, I turn to the third and most basic valuable feature that beliefs can have: truth itself. Some challengers hold that if there is no explanation of our evaluative reliability, then our evaluative beliefs are highly unlikely to be true at all: there *is* no evaluative reliability to be explained. Once we realize this, we will see that we have no reason to regard our ordinary evaluative beliefs as true—even if, as a psychological matter, we will continue to hold them. In response, I adapt a position take by the philosopher G.E. Moore, in reply to doubts about our knowledge of physical reality. However unlikely it might be that our beliefs are true if our reliability has no explanation, we have to weigh the reasons for doubt that this gives us against the reasons we have to believe the evaluative claims themselves. I claim that it will always be more reasonable to continue to regard our ordinary evaluative beliefs as true than to give them up for the sorts of reasons connected with the epistemic challenge.

Appeal to the value of explanation

In chapter two, I conclude that the epistemic value of our ordinary evaluative beliefs faces no threat from a failure to answer the epistemic challenge. But this does not show that nothing of value is at stake in answering the challenge. After all, *explanations* are thought to have value in their own right, and answering the epistemic challenge just is giving an explanation of our evaluative reliability. In chapter three, I do not argue that this value *isn't* lost if we cannot answer the epistemic challenge, but rather that epistemic challengers have not done enough to show that this value is at stake. I argue that they have more work to do in order to motivate answering the challenge in terms of the demand for an explanation, and I sketch the work remaining to be done.

I consider on behalf of the epistemic challenger, two arguments for the claim that it is good to have an explanation of our evaluative reliability. I raise serious problems for each.

On the one hand, we consider some things highly likely to have an explanation. If our evaluative reliability is such a thing, then any meta-ethical view that rules out there being such an explanation implies a likely falsehood. In answering the epistemic challenge, then, metaethicists would avoid commitment to a likely falsehood: surely a good thing. I argue that no one is yet in a position to say that our evaluative reliability is highly likely to have an explanation. To be in this position, we would need background knowledge that I argue we do not have. So we cannot yet claim that this is a reason to explain our evaluative reliability.

On the other hand, having an explanation of our evaluative reliability might furnish an intellectual good, in the form of the value of explanation. We aspire to explain things, and we are disappointed if no explanation is available. It seems straightforward to conclude that if we cannot answer the epistemic challenge, this is bad for precisely this reason. I argue that this is not as obvious as it seems, by showing that on one prominent account of the value of explanation, the straightforward conclusion cannot be drawn. I describe the “unificationist” view defended by Philip Kitcher, on which the value of explanation is found in the way explanations unify

our picture of the world. Explanations of particular facts only have value as part of a set of explanations with certain common features. Given this more specific conception of explanatory value, there is no guarantee that an explanation of our evaluative reliability will belong to such a set. To motivate answering the epistemic challenge by appeal to the value of explanation, the challenger needs to make determinate claims about how explaining our evaluative reliability would realize that value.

Some conclusions

One result of this project, which emerges from chapters two and three, is that a key question that has interested metaethicists might have no importance. The most salient epistemic consequences thought to follow from a failure to answer the epistemic challenge do not seem to follow. A more speculative attempt to motivate the epistemic challenge by appeal to the value of explanation remains inconclusive. It seems, then, that metaethicists should not treat answering the epistemic challenge as one of the demands on a satisfying metaethical theory. More broadly, they should decline to engage with the epistemic challenge at all, until a reason for doing so has been found.

A related point and a broader moral emerges when we consider chapter one. We noted there that, in their rush to either answer the challenge or show it unanswerable, metaethicists have rarely paused to make clear why it is important for a theory to be able to give an answer. It may be that a great deal of effort has been wasted as a result. This case can serve as a reminder for metaethicists, and philosophers more broadly, to give more careful thought to why it is important to take up a challenge. Barry Stroud has called on philosophers to be more self-reflective about their apparent tasks:

What philosophers need now is a diagnosis or uncovering of what they regard as their problems or their questions, and some understanding of the nature and sources of the kinds

of things they think philosophy should account for.²

Following Wittgenstein, Stroud suggests that philosophers “treat” a question as doctors treat an illness: by starting with a diagnosis. A philosopher’s diagnosis should “identify the assumptions, the demands, the preconceptions, and the aspirations that lead to a question’s having the particular significance it now has for us.” Like a medical treatment based on a hasty diagnosis, the rush to answer a question before appreciating its significance “can make things worse.”³

In my engagement with the “epistemic challenge,” I have stressed a complementary aspect of such a diagnosis: the evaluation of our philosophical demands and aspirations. Are they good ones for us to have, and why? The importance of this question can be made vivid by extending the metaphor of treatment. A key step in a doctor’s diagnosis ought to be the determination that a condition is in fact an *illness*, and that through treatment, we have some hope of achieving a good. A doctor should not subject her patient to treatment without a clear good in view.

²Stroud (2001) pp. 41-42

³Ibid. p. 42.

Chapter 1 The Epistemic Challenge Meets the Value Problem

1.1 Realism and the epistemic challenge

Philosophers have recently turned to the epistemology of value to help settle disputes about the metaphysics of value. A generation ago, philosophers debated whether metaethical theories are defective if they imply that we cannot *empirically test* our views about value.¹ More recently, philosophers have asked whether metaethical theories are defective if they imply that we cannot *explain our reliability* about value. This question is at the heart of the “epistemic challenge for realism,” and is the focus of this dissertation. Before turning to the implications of the challenge for such theories, I will characterize the variety of realism that is its most common target and lay out the shape of the challenge itself.

Metaethicists who embrace *non-naturalist realism* seem to face particular difficulties explaining our evaluative reliability. Non-naturalist realism consists of three commitments. First, values and reasons² exist, in at least the sense that some evaluative claims—claims that attribute evaluative and normative features to things—are true. This presupposes that evalu-

¹See Harman (1977, 1986); Sturgeon (1985); Harman & Thomson (1996); Sayre-McCord (1988)

²I will use the terms “values” and “reasons,” as well as “evaluative” and “normative,” interchangeably. Some hold that one or the other of types of feature is basic—I will remain neutral on this dispute. Nothing I say here should turn on the distinction between these notions, or on whether one is ultimately reducible to the other.

ative claims are *truth apt*, unlike commands or exclamations.³ Second, values and reasons are real, in the sense that whether things have evaluative and normative features is relevantly mind- and language-independent. In particular, whether something has evaluative features does not depend on whether we *believe* or *experience* or *accept* that it has these features, nor on whether we *would* do so under other conditions.⁴ Third, values and reasons are not reducible to “natural” features, where these are identified by their role in good scientific explanations. The irreducibility of values and reasons to natural features is taken to imply both that they can have no role in scientific explanations and that they are causally impotent.⁵ In what follows, by “realism,” I will mean non-naturalist realism.

The epistemic challenge can be posed against this sort of realist in the following way. To begin, the challenger suggests that all of us, realists included, tend to assume that we are not entirely hopeless at discovering the evaluative truth. On the contrary, we assume that we tend to believe the truth about value: in other words, that we are evaluatively reliable, on the whole. We can think of our reliability as a correlation between our evaluative psychology and the evaluative truths. David Enoch describes the correlation this way:

Very often, when we accept a normative judgment *j*, it is indeed true that *j*; and very often when we do not accept a normative judgment *j* (or at least when we reject it), it is indeed false that *j*. So there is a correlation between (what the realist takes to be) normative truths and our normative judgments.⁶

The challenger goes on to say that this correlation cries out for an explanation. Why does what we *believe* about value tend to align with what is *true* about value? This cannot, she suggests, be a mere coincidence.

³The view that evaluative claims are truth apt is “cognitivism.” Oddie (2005) p. 4; Kolodny (2006); Shafer-Landau (2003) p. 17.

⁴Russ Shafer-Landau and Sharon Street refer to this sort of reality as “stance-independence.” See (2003) pp. 13-18; Street (2006) p. 111.

⁵See, e.g. Enoch (2011) pp. 4, 103; Shafer-Landau (2003), p. 55.

⁶Enoch (2009) p. 421.

Yet, the challenger argues, the realist has no choice but to regard our reliability in this way. Conceiving of values and reasons as she does blocks all promising avenues for such an explanation. According to the realist, values and reasons cannot causally impact anything, and so, in particular, cannot have a causal impact on our evaluative psychology. Neither does our evaluative psychology explain the evaluative truths, either through causal impact or through a constitutive relation to the truths. On the contrary, the realist insists that facts about value are both causally and constitutively independent of what we accept or believe, and even of what we *would* accept or believe under other circumstances.

The realist can try to make answering the challenge more tractable by clarifying the explanatory relations among the evaluative facts and those among the relevant psychological facts. For example, if the evaluative facts have an axiomatic structure, a few fundamental evaluative facts will explain many others. In this case, it will be possible to make many true evaluative judgments by deriving them from others. If we *in fact* make our evaluative judgments in roughly this way—by deriving them from a few fundamental evaluative judgements—explaining the reliability of our fundamental evaluative judgments would suffice to answer the challenge. The realist can focus on the narrower question: why do *these* fundamental beliefs about value tend to align with what is true about value?⁷ As we will see, even if our judgments and the evaluative truths have this structure, the realist is thought to face the same obstacles in answering this narrower question.

This is usually brought out through a comparison with the prospects for giving such an explanation in other domains. Consider the parallel ques-

⁷This point is equivalent to one made by Field in posing an epistemic challenge to the mathematical platonist: “as mathematics has become more and more deductively systematized, the truth of mathematics has become reduced to the truth of a smaller and smaller set of basic axioms; so we could explain the fact that the mathematicians’ beliefs tend to be true by the fact that they have been logically deduced from axioms, if we could just explain the fact that what mathematicians take as axioms tend to be true.” Unlike the mathematical judgments of mathematicians, however, our evaluative judgments do not, in general, result from internalizing a systematic science. We would need evidence of another kind—perhaps the evidence of anthropological investigation—to establish that some of our evaluative beliefs derive from more basic ones in a similar way. See Field (1989), p. 231.

tion: why does what we *believe* about our physical surroundings tend to align with what is *true* about those surroundings? Realists about our physical surroundings can respond by appealing to their causal impact on us through perception. The surfaces of physical objects reflect waves of light, and these impact our eyes; such objects give off sound waves, and these impact our ears; sensors in these organs send signals to specialized parts of the human brain designed to interpret them; and so on. While this is merely a sketch, accounts of perception along these lines help us see why forming beliefs about our physical surroundings using perception would reliably yield true beliefs about them.

Even in outline, one might worry that the explanation is incomplete, since it takes for granted our possession of accurate perceptual organs. A complete explanation of why our beliefs tend to align with the truth should also address our *coming to have* such a reliable way of forming beliefs, in the first place. Yet here too, realists about our physical surroundings have an account available. They may appeal to our evolutionary history, specifically the survival value for our ancestors of reliability about their physical surroundings. As Sharon Street observes, when it comes to “truths about a creature’s manifest surroundings—for example, that there is a fire raging in front of it, or a predator rushing toward it”—it is “clear why it tends to promote reproductive success for a creature to grasp such truths: the fire might burn it to a crisp; the predator might eat it up.”⁸ Selective pressures would tend to favor creatures that accurately perceived such truths: more of those creatures would survive and reproduce. To the extent that their reliable perceptual abilities were heritable, we would likely have inherited them.⁹

An explanation along these lines has the potential to render our reliability about our physical surroundings no coincidence. According to these explanations, we believe what we do about them at least in part *because* of how things are with them. In forming beliefs on the basis of what we perceive, we form beliefs about how things are with our physical surroundings through causal contact with those surroundings. Furthermore, we *have* reliable methods of making judgments about the physical world *because* (in some sense) the methods are reliable. To put the point crudely,

⁸Street (2006), p. 130. See also Schafer (2010), p. 472.

⁹See also Setiya (2012) pp. 104-110.

we possess these methods because they were passed on to us, and they were passed on to us because they were reliable.¹⁰

No such explanation appears to be available to the non-naturalist metaethical realist. Such a realist denies that values and reasons can impact us causally, as our physical surroundings do in perception. While some have claimed that recognizing value might have been adaptive for our ancestors, it is hard to make this argument—Street has argued that, by scientific standards, such an explanation is inferior to one that makes no reference to evaluative truth.¹¹ Perhaps the realist could appeal to a common cause—a “third factor”—that explains *both* the evaluative truths *and* our evaluative psychology.¹² Yet given realist assumptions, it is doubtful that a genuinely common cause—a cause that accounts for both of these factors, making their correlation no coincidence—is available. After all, what sort of cause could this be? According to the realist, evaluative features are not the sort of feature that can play a role in explanations in the natural sciences. These presumably include the explanations of interest in the disciplines of psychological and evolutionary biology. Yet when it comes to the explanation of the development and workings of our psychology, these disciplines seem to offer our best hope.¹³ Defenders of realism admit the limits of third factor replies to the challenge. Karl Schafer, for instance, concedes that in attempting to explain our evaluative reliability, sooner or later the realist will be left with two sets of claims:

the non-normative ones that form the foundation of his non-normative theory, and the foundational normative claims about how the normative supervenes on the non-normative...At this stage, the realist will have to concede that there is no explanatory relationship between these two sets.¹⁴

The prospects for a genuine third factor explanation therefore appear doubt-

¹⁰See Setiya (2012) pp. 106-108, for some helpful discussion of how to think about how we come to have adaptive traits.

¹¹Ibid. (2006)

¹²Indeed, many realists who address the challenge adopt this strategy. See, e.g. Enoch (2009); Berker (2014); Wielenberg (2010); Copp (2008).

¹³See, e.g. Tomasello (2015), Cosmides et al. (2019).

¹⁴(2010) p. 483

ful.

In short, explaining why our beliefs tend to align with the independent truths, the realist about value seems to lack explanatory resources had by the realist about mid-sized objects.¹⁵

1.2 What is at stake in the epistemic challenge?

What follows if the metaethical realist cannot explain our reliability about value, so as to show that it is no coincidence? Some urge that, since the coincidence is between our psychology and the truth, our evaluative beliefs would be at best accidentally true, and so *knowledge* of value would be impossible on a realist view.¹⁶ Others worry that our being evaluatively reliable by coincidence would be incredible, too much to believe.¹⁷ If the realist must give up belief in our evaluative reliability, it seems she loses *epistemic justification* for beliefs about value.¹⁸

Importantly, the challenge does not call into question the possibility of evaluative truth, realistically construed. Were the challenger to directly

¹⁵Many versions of the epistemic challenge, and many responses to it, take as background that evolutionary forces have fundamentally shaped our evaluative dispositions, and specifically ask whether there is a satisfying explanation of why such forces would tend to align with the realist's independent evaluative facts. See, e.g. Street (2006), Copp (2008), Enoch (2009), Wielenberg (2010), Schafer (2010), and Shafer-Landau (2012), among others. I agree with Enoch and Street that the role of evolution in the challenge is strictly dispensable. (See Street (2006), p. 155; Enoch (2009), p. 426.) The problem is not specifically that our dispositions were influenced by *evolution*—rather, the problem is that the “best causal accounts of our evaluative judgments, whether Darwinian or otherwise, make no reference to the realist's independent evaluative truths.” Street (2006), p. 147.

¹⁶See Setiya (2012) ch. 3, Wielenberg (2010), Shafer-Landau (2012) p. 4.

¹⁷As Enoch and Setiya note, when construed in this way, the epistemic challenge to metaethical realism is very close to the challenge posed to the mathematical Platonist by Hartry Field. Enoch (2009), p. 421. Setiya (2012), pp. 66-76.

¹⁸See Enoch (2009) pp. 423-424; Shafer-Landau (2012) p. 4; Bedke (2014) p. 103; Locke (2014), Enoch and McPherson (2017), and Clarke-Doane (2017b). In addition to these epistemic worries, such inexplicability might reveal a *theoretical* failure—perhaps a metaethical theory that cannot explain how we became reliable about reasons and values loses “plausibility points.” See Enoch (2009), Berker (2014), Shafer-Landau (2012).

argue against the existence of such truths, she would offer a metaphysical rather than an epistemic challenge. It therefore begs no questions, in this context, to take for granted that there are independent evaluative truths. Given this, we may assume that forming true beliefs about such values and reasons is *possible*.¹⁹ To paraphrase David Enoch, wherever there is truth, there is the possibility of true belief.²⁰

Depending on the formulation, then, the epistemic challenge may threaten evaluative knowledge, epistemically justified evaluative belief, or some related epistemic status in evaluative thought. Many contributors to this literature begin by framing the threat in terms of one or another of these epistemic statuses, before either seeking to meet the challenge or arguing that it cannot be met.²¹

My purpose in this chapter is not to follow these contributors. Instead, I want to understand why the realist ought to show that her metaethical theory is compatible with achieving any given epistemic status in evaluative thought. Addressing this question will reveal a promising way to approach formulations of the challenge, one that focusses squarely on what is at stake in the search for an answer.

¹⁹It *would* beg the question against an important version of the challenge to assume that our evaluative beliefs (understood realistically) are actually on the whole *true*. As will become clear in chapter two, however, I believe that begging this question is precisely what the realist ought to do. The genuine epistemic challenge is *not* to the truth (by and large) of our evaluative beliefs. I will not come to defending this position until chapter two.

²⁰As Enoch puts this: "A class of beliefs is reliable, I take it, if and only if a sufficiently large portion of it is true. If so, wherever there is truth, there is the possibility of reliability. Now, it is not obvious, of course, that there *is* truth in ethics and in normativity, and some metaethicists think no such truth is to be had. But such a claim can certainly not be used as a premise in what is supposed to be an independent [epistemic] challenge to realism." (2009) p. 418.

²¹See, for instance, Shafer-Landau (2003) Ch. 10, and (2012); Street (2006) and (2009); Copp (2008); Enoch (2009); Schafer (2010); Wielenberg (2010); Parfit (2011) 32.114; Setiya (2012); Clarke-Doane (2012); Behrens (2013); Berker (2014); Bedke (2014); Locke (2014); Morton (2016); Huemer (2016); Hayward (2018); Baras (2018).

A shared preference for compatibility

Among metaethicists, there is a clear preference for metaethical theories that are compatible with achieving traditional epistemic statuses, such as knowledge and epistemic justification, in evaluative thought. Let's call this preference for compatibility between one's metaethical theory and traditional epistemic statuses a preference for *epistemic compatibility*.

Realists often express a preference for epistemic compatibility. Some worry that, if realism were to imply that knowledge or epistemic justification is impossible in evaluative thought, that would give us a reason to reject realism. In his answer to the epistemic challenge, Russ Shafer-Landau asserts that, if realism implies that we cannot have moral knowledge, this "leaves realists in a deeply unappealing position."²² David Enoch, in his discussion of the challenge, characterizes accepting the view that "no normative belief is ever justified" as "biting the bullet," and seems to take for granted that realists ought to aim for "a theory of justification and a metanormative theory" that "cohere nicely with each other in a non-skeptical way."²³ In a book-length defense of his own variety of non-naturalist realism, Enoch further suggests that "if the apparatus needed for a rejection of normative skepticism is unavailable to the robust realist...this may count as a reason to reject Robust Realism after all."²⁴

While realists prefer epistemic compatibility, so, apparently, do many contemporary anti-realists. Sharon Street and Kieran Setiya follow up their arguments that realism is epistemically incompatible by arguing that their favored anti-realist metaethical theory *is* epistemically compatible. This, they each suggest, gives us reason to prefer anti-realism over realism. On a realist view, Street urges, our evaluative reliability would be so unlikely that we should think that we probably aren't reliable, a "far-fetched skeptical result."²⁵ Anti-realism avoids such skepticism, she suggests, assuring our reliability through a constitutive connection between the evaluative truth and our evaluative psychology. Setiya, on the other hand, suggests that we must accept a series of theses about the epistemology

²²(2012) p. 1

²³(2009) pp. 417-418

²⁴(2011) p. 5

²⁵(2006) p. 122

and metaphysics of ethics, including his anti-realist metaethical theory, if we are to avoid the conclusion that we are unable to attain ethical knowledge.²⁶

Epistemic compatibility, then, has come to seem important not only to realists, but also to anti-realists. In the literature on the epistemic challenge, epistemic compatibility is treated as a more or less fixed point, an aim for any metaethical theorist.²⁷ This suggests that there is a reason for preferring epistemic compatibility that realists and anti-realists share.²⁸

Why, then, prefer epistemic compatibility? Suppose that a given metaethical theory implied that we can *believe* that some course of action is best, or that some moral principle is correct, or that some piece of art is sublime, but that these beliefs cannot have the traditional sorts of positive epistemic status. Suppose that, according to that theory, we can never *know* the answers to evaluative questions one way or another, nor have *epistemic* justification, such as conclusive evidence of the likely truth of our evaluative beliefs. Why exactly would that count against such a theory?

We should distinguish this question from two similar questions that appear in this literature. First, the question is not about whether epistemic compatibility is *indispensible*. While realists admit some pressure to preserve epistemic compatibility, they often urge that the inability to do so would not settle the question of whether to accept realism. Rather, whether we ought to accept realism depends on its overall plausibility, when compared with other available metaethical positions. Shafer-Landau, among others, makes this point:

If realism makes good sense of our moral practices, yields a plausible metaphysics, has compelling views about moral psychology, and can develop a suitable account of morality's normativity, then it might be able to survive (worst-case scenario)

²⁶(2012) p. 5

²⁷There are exceptions, such as the error theory—a metaethical theory that denies the existence of evaluative features entirely.

²⁸We might contrast prevailing assumptions in this literature with those in play in earlier epistemic challenges, in which the burden of accounting for evaluative knowledge was taken to belong to realism exclusively. See, e.g. J.L. Mackie (1977).

the embarrassment of lacking a plausible moral epistemology. That depends on the theoretical virtues and vices possessed by its competitors.²⁹

Realism's epistemic incompatibility may be an "embarrassment," without providing a decisive reason to reject realism. That this would be an embarrassment at all, however, presupposes that there is *some* reason to prefer a metaethical theory that is epistemically compatible. We are still left with the question: why should we prefer compatibility over incompatibility?

Second, the question is not why the challenge makes trouble for *realism*, rather than for the conception of the epistemic status in play in a given formulation of the challenge. In arguing that realism is incompatible with knowledge or with justified belief about value, challengers typically rely on assumptions about the relevant epistemic status—for instance, on a necessary condition on knowledge that beliefs about non-natural, independent values could not satisfy.³⁰ Realists may therefore dispute the epistemic assumptions in play, leaving open that their theory preserves epistemic compatibility when these statuses are properly understood.³¹ Proceeding in this way, however, still suggests that we have *some* reason to prefer a view of the nature of value that is compatible with knowledge or epistemic justification. But what justifies this preference?

²⁹(2003) p. 234. Here is another statement in this vein, from David Enoch: "We are to opt for the theory that best explains whatever needs explaining, or perhaps—if we can restrict ourselves now to just choice among competing metanormative theories—whatever needs explaining in the vicinity of normative discourse and practice." (2009) pp. 427, 413-414

³⁰The relevant conception need not take the form of such a condition, or of a full or partial analysis of the epistemic status. David Enoch, for instance, relies on the claim that our epistemic justification is defeated without some explanation of our reliability. While this may not be a claim about the nature of epistemic justification, it is a claim about epistemic justification. *Ibid.* p. 424.

³¹See, e.g. Schafer (2010), Copp (2008). Enoch makes a similar point: "if a theory of epistemic justification rules out all normative beliefs as epistemically unjustified, this counts at least as heavily against that theory of justification as it does against the justificatory status of normative beliefs." *Ibid.* p. 417.

Obviously true?

Perhaps we should prefer epistemic compatibility out of concern for the truth. Perhaps we have strong reasons to think it *true* that we can achieve knowledge and epistemically justified belief in evaluative thought. Most straightforwardly, perhaps we have strong reasons to think that we have *in fact* achieved these statuses in some cases. If we had such reasons, that would support a preference for epistemic compatibility, since we want our metaethical theory to be compatible with other claims we take to be true.

A number of commentators do suggest that it is *plausible*, or that we should have *confidence* that some of our evaluative beliefs have the relevant epistemic status. In his formulation of the epistemic challenge, Setiya begins by expressing confidence that we have achieved not only epistemically justified beliefs about ethics, but *knowledge* of ethics:³²

Much of what I hold in ethics I take to be not merely true, and justified, but known. I do not merely think that slavery is wrong, I know it is. I know that women are not the property of men, that there is reason to care about people other than myself, that one should respect the rights of the innocent even at some cost to the greater good. (Ibid. p. 2)

Along similar lines, David Enoch suggests that one ought to have at least as much confidence that some of one's normative beliefs are *epistemically justified* as one does in any general theory of epistemic justification:

...if a theory of epistemic justification rules out all normative beliefs as epistemically unjustified, this counts at least as heavily against that theory of justification as it does against the justificatory status of normative beliefs. For is there really a the-

³²Setiya is interested in the epistemic and metaphysical character of the "ethical," but construes it broadly, as including not only "claims of right and wrong, and of what there is reason to do," but also "reasons for wanting, respecting, or admiring things" and "claims of virtue and vice." "[If] in doubt," he suggests "it's ethical." (2012) p. 3.

ory of epistemic justification in the truth of which you are (and should be) more confident than in the truth of such sentences as “I justifiably believe that torture is prima facie wrong”?³³

Enoch makes a similar assertion about our knowledge of normative claims.³⁴ Once again, if our beliefs in the wrongness of torture and of slavery are clear instances of *knowledge*, or of *epistemically* justified belief, then realism does face a challenge. All else equal, our metaethical theory ought to be compatible with other claims we take to be true.

But is the truth of these claims so clear that realists and anti-realists ought to avoid contradicting them, even absent an explicit argument for them? Should we assume that we have *knowledge* of value, for instance? While philosophers increasingly take the possibility of knowledge for granted in some domains, its possibility in others remains controversial. Claims to mathematical knowledge, knowledge of empirical generalizations, and knowledge of the future have all faced serious philosophical challenges.³⁵ Claims to knowledge of *value* are at least equally controversial. This is in part because, in domains whose claims are subject to persistent disagreement, philosophers have associated the possibility of knowledge with the possibility of expertise. They have also associated the possibility of knowledge with the acceptability of making up one’s mind based on testimony. Both have struck philosophers as objectionable in the domain of value.³⁶

Among non-philosophers, it is certainly not a matter of common sense that anyone has knowledge of value. Unlike, for instance, the claim that we have knowledge of our immediate physical surroundings, it is not true that, before encountering philosophical skepticism, we take for granted that we know what is valuable or what we have reason to do. (On the contrary, the sense of being at sea with respect to questions of value is what drives many students to philosophy.) Many reasonable people would

³³(2010) p. 417

³⁴If “an account of knowledge entails that we can never know that racist discrimination is unjust (and the like), this should be taken as a strong reason to reject that account of knowledge, rather than to reject normative knowledge.” Ibid. p. 419.

³⁵See, e.g., Field (1989), Morrison (2011), Fisher (1975).

³⁶See, e.g. Williams (1995) pp. 203-212; Hills (2010) pp. 169-187.

deny that they personally have knowledge of value, and would express doubt that anyone does. This is reflected in the common sentiment that questions of value are “matters of opinion,” and cannot really be “proved.”³⁷

Perhaps we can make a more convincing case that some of our evaluative beliefs qualify for a more basic cognitive status, such as *epistemic justification*. After all, it *does* seem a matter of common sense that some of us have very good, even decisive, reasons to believe some evaluative and normative claims. Like Setiya, I believe there are decisive reasons to believe that slavery is wrong and that women are not the property of men. More broadly, I am inclined to agree with Ronald Dworkin that we “have considerable evidence...of a capacity to make moral judgments that bring conviction, that are mainly durable, that agree with the judgments of a great many others, and that are amenable to the normal logical combinations and operations.”³⁸ It is an element of the durability of many of our moral and other evaluative judgments that we feel ourselves under no obligation to give them up; on the contrary, we take ourselves to have a strong obligation to maintain them.

Assuming that we ought, on the whole, to believe certain evaluative claims, this does not yet show that our decisive reasons to believe are *epistemic*. A critical role might be played by *other* sorts of reasons.³⁹ For example, it is possible to understand our fundamental convictions about value as articles of faith, or as conditions of personal integrity, or as necessary hopes, and so to see them as supported on non-epistemic grounds, rather than as instances of epistemically justified belief.⁴⁰ We may share Dworkin’s confidence in our capacity to make durable evaluative judgments, and regard ourselves as obliged to maintain these judgments, without having

³⁷In *The Philosophers’ Magazine*, John Corvino notes and deplores the frequency of statements like these in debates on controversial topics. He takes such statements to depend on a suspect distinction between fact and opinion. I would be more inclined to interpret such statements in terms of the distinction between *knowledge* and opinion.

³⁸(1996) p. 118

³⁹These would be reasons for thinking these convictions *true*, just not *epistemic* reasons for doing so.

⁴⁰My point here is that our obligation to believe could be vindicated, even if we do not have decisive epistemic reasons to believe. Whether this would make our evaluative beliefs worse, or leave us anything else to regret, is another question, which I defer to the next section.

taken a position on their epistemic credentials.

If we can raise reasonable doubts about whether we have *actually* achieved the relevant epistemic statuses in evaluative thought, it is hard to see why we should insist that it must be *possible* for us to do so. At the very least, we would need argument to this effect, and we do not typically get such an argument in statements of the epistemic challenge.

These reflections will not convince everyone: some will insist that evaluative thought offers paradigm instances of knowledge and epistemic justification. Given that we *seem* knowledgeable and epistemically justified in evaluative thought, they might urge, the realist is under some pressure to resolve the conflict between this appearance and the challenger's argument that we cannot achieve such statuses. However, I suggest that a significant role is played by something else. That is an assumption about the *value* of such epistemic statuses.

Obviously valuable?

I suggest that metaethicists in this debate feel pressure to answer the epistemic challenge because they believe that it is *very important* to achieve knowledge and epistemic justification in evaluative thought. They may assume that such an achievement would have great *positive* value, as the realization of a hope. Or they may assume that the impossibility of such an achievement would be of great *negative* value, as a kind of disaster. In either case, the impossibility of achieving knowledge and epistemic justification in evaluative thought would be massively disappointing. At bottom, I think this assumption motivates metaethicists' preference for epistemically compatible theories, and so their desire to answer the epistemic challenge.

This might seem to accuse metaethicists of a kind of wishful thinking—of choosing their metaethical theory based merely on what they want to be so, rather than on sound theoretical reasons. Such a charge would be too hasty. I suggest that metaethicists' assumptions about the importance of these epistemic statuses lead them to *prefer* epistemic compatibility, and so motivate them to *answer* the challenge. This can be so, without their

preferences influencing their beliefs by, for instance, deceiving them into thinking that they have successfully responded to the challenge. Furthermore, their initial preference for epistemic compatibility need not lead them to overlook reasons to doubt epistemically compatible theories. In general, there is nothing illegitimate about preferring a theory because it is compatible with achieving some important cognitive good. It is legitimate for physicists to prefer a unified account of physical phenomena (“a theory of everything”), for the sake of the valuable sort of understanding that such a theory would make possible, provided that they remain clear-eyed about whether a unified account is viable. Preferring a metaethical theory that leaves room for the possibility of an important form of cognitive success, I suggest, is similarly permissible.⁴¹

Setting this concern aside, do philosophers engage with the epistemic challenge because they take for granted the value of such epistemic statuses? In the literature on the epistemic challenge for realism, the value of knowledge or epistemic justification is occasionally declared, and in a way that seems to motivate the challenge. As we have already seen, Shafer-Landau holds that denying the possibility of moral knowledge “leaves realists in a deeply unappealing position.”⁴² Setiya maintains that conceding our inability to attain ethical knowledge would be “shocking,” and “too great a concession.”⁴³ Erik Wielenberg, a non-naturalist realist, goes a bit further, stating that “if there is no moral knowledge, then, as Jerry Fodor remarked in a different context, ‘practically everything I believe about anything is false and it’s the end of the world.’”⁴⁴ Each of these observations is offered early in posing the epistemic challenge.

More often, however, the commitment to the value of the relevant epis-

⁴¹Being motivated in this way is an expression of optimism that some good is achievable. This need not preclude a concern with whether it is in fact achievable. Provided that one can suspend optimism when needed, to submit it to critical reflection, one need not do so at every moment. One may have pragmatic, as well as theoretical, reasons to take optimistic possibilities more seriously in inquiry.

⁴²(2012) p. 1

⁴³Setiya urges that evaluative knowledge specifically is of great value, so that we ought to lament its loss even if we retained justification: “it is bad enough to admit that we do not know the injustice of torture or slavery, that we have mere justification for these claims.” (2012) pp. 138-142.

⁴⁴(2010) p. 442

temic statuses is expressed indirectly. Take, for example, ubiquitous references to the risk of “skepticism” in this literature. If realism is incompatible with epistemic justification or knowledge, we are repeatedly reminded, realists fall prey to “thoroughgoing skepticism.” If we cannot have “moral knowledge” on realist assumptions, Shafer-Landau states, this “does not refute moral realism, but it leaves realists in the deeply unappealing position of being saddled with a thoroughgoing moral skepticism—a logically coherent position that contains about zero appeal.”⁴⁵ Enoch warns that unless realists can accommodate epistemic justification for our evaluative beliefs, they will pay the “price of a rather thoroughgoing skepticism.”⁴⁶ The case for Setiya’s anti-realist conclusions about the nature of value is simply that “we must accept them if we are to avoid the prospect of ethical scepticism.”⁴⁷

After characterizing the realist’s predicament in this way, these authors move on, without elaborating on the nature of skepticism. By “skepticism,” then, they seem to mean only the impossibility of achieving the relevant epistemic status. The view that knowledge or epistemic justification is impossible in a broad domain of inquiry is indeed a variety of skepticism.⁴⁸ But if this is all that “skepticism” means, these authors repeat themselves. They emphasize that the impossibility of achieving some epistemic status in evaluative thought would saddle realists with skepticism...in the form of the impossibility of achieving that epistemic status in evaluative thought.⁴⁹

Why do they repeat themselves? Engaging in repetition may serve a rhetorical purpose. By invoking “skepticism,” these authors emphasize that the

⁴⁵ (2012) p. 1

⁴⁶ (2009) p. 424

⁴⁷ (2012) p. 5

⁴⁸ As Barry Stroud expresses this: “...the problem is to show how we can have any knowledge of the world at all. The conclusion that we cannot, that no one knows anything about the world around us, is what I call ‘scepticism about the external world’” (1984) p. 1.

⁴⁹ Sharon Street is arguably an exception, in that she fills out the threat in terms of *likely truth*—she suggests that, in failing to answer the challenge, the realist must admit “that many or most of our evaluative judgments are off track” (2006), p. 122. Critics have questioned whether the challenge could lead us to regard our evaluative beliefs as (all things considered) unlikely to be true. See, e.g. Schafer (2010), Berker (2014). We will take up this question in the next chapter.

inability to achieve epistemic justification or knowledge in our evaluative thought would be a *very bad thing*. Similarly, one might greet someone's plan to steal something with the observation, "But that would be *stealing!*" Such a rhetorical use of "skepticism" would not be surprising: unlike the view that knowledge is possible or actual in every domain of inquiry, it is common ground among contemporary philosophers that skepticism is a very bad thing. For one thing, doubting that one can know or justify one's belief about some particular question often leads one to *give up* that belief. In light of this, becoming skeptical about an entire domain makes salient the possibility of giving up an entire class of one's beliefs, a potentially devastating result. Even if we can dismiss the possibility of giving up one's beliefs, accepting skeptical arguments is thought to involve a state of permanent dissatisfaction, or a denial of the needs of reason. Invoking skepticism to highlight the importance of knowledge or epistemically justified belief about value is therefore rhetorically effective. This is so, even though there is some danger that it is *merely* rhetorically effective. If we were to become *agnostic* about whether anything has value, evaluative thought and practice would certainly collapse;⁵⁰ however, we do not yet have an argument that giving up on knowledge or justified belief about value would or could lead to such agnosticism. Similarly, while a state of permanent intellectual dissatisfaction is surely a bad thing, we do not yet have an argument that what reason needs in its dealings with value is knowledge or epistemic justification.

Of course, it might turn out that this *is* what reason needs, or at least what we ought to want. If so, those who mount the epistemic challenge charge realists with putting something important out of reach. Realists seek to answer the challenge in part out of optimism: they believe that realism is true, while also hoping that this does not commit them to giving up any important cognitive achievement in evaluative thought.

⁵⁰See, e.g. Kahane (2017).

1.3 The epistemic challenge meets the value problem

Suppose that metaethicists prefer epistemic compatibility because they assume that achieving traditional epistemic statuses in evaluative thought is of great importance. In this case, the legitimacy of any particular version of the epistemic challenge will depend on showing that the relevant status has this importance in evaluative thought. In fact, the realist's interest in answering the challenge ought to be proportional to this importance. If achieving, say, knowledge in evaluative thought would be of *great* and *distinctive* value—one of the most important goods in life, that no other good could replace—the failure of realism to allow for that status would be devastating. This might indeed seem the “end of the world” from a realist's perspective. On the other hand, if achieving that epistemic status would be of *no* value in evaluative thought, or if it were replaceable by some equally valuable status left untouched by the challenge, the realist has no reason to prefer epistemic compatibility. Not only is it *not* the end of the world, it is not the end of *anything* we have reason to care about.

One might think that the challenger can maintain the great value of traditional epistemic statuses without additional argument. Historically, philosophers have taken knowledge and epistemic justification to have great cognitive value. The distinctive value of such statuses might even provide a starting point for inquiry into the nature of knowledge. As Linda Zagzebski puts the point, “knowledge has to be defined as something we value. We are not interested in a phenomenon with little or no value.”⁵¹ Similarly, Jonathan Kvanvig takes there to be a “presumption in favor of the value of knowledge,” one that is “strong enough that it gives reason to abandon even a counterexample-free account of the nature of knowledge if that account leaves no way open for defending the value of knowledge.”⁵² In the context of the epistemic challenge, why not take for granted that knowledge and epistemic justification have significant value?

Yet as the work of Kvanvig, Zagzebski, and the broader literature on the “value problem” in epistemology suggest, widely accepted philosophical accounts of these epistemic statuses can fail to vindicate the distinctive

⁵¹(2003) p. 26

⁵²(2003) p. 5

value of those statuses. Take, for example, reliabilist accounts of epistemic justification, according to which one's belief is epistemically justified if it was formed using a reliable method. On the most straightforward reliabilist account, being epistemically justified has *instrumental* value: it increases one's likelihood of getting the truth. Like the value of a reliable espresso maker, the value of justification derives wholly from the value of its results.⁵³ It follows, however, that having an epistemically justified true belief is no better than having a true belief.⁵⁴ The value of truth *swamps* the value of epistemic justification, in that once you have a true belief, there is no reason to want that belief to be reliably formed.⁵⁵ Any threat to epistemic justification, construed this way, could only motivate us to answer the challenge if it were *also* to threaten true belief. Our method of forming the belief could be unreliable, while making no difference at all to the value of our belief: one sometimes finds that one got lucky, and on this account of justification, that is all one could want.

Similarly, if we assume with most contemporary epistemologists that the correct account of knowledge must resist Gettier-style counterexamples, it is doubtful that knowledge itself has any distinctive value.⁵⁶ That is, it is unclear that knowledge has value greater than that of other cognitive states in the vicinity of knowledge—for example, than true belief that expresses intellectual virtue, a state possible for a victim in a Gettier case.⁵⁷ Kvanvig argues that we have strong inductive evidence that knowledge does not have distinctive value.⁵⁸ This is because we repeatedly observe

⁵³Zagzebski (2003) p. 13.

⁵⁴Reliabilist accounts of epistemic justification are the most prominent example. However, any view of the value of epistemic justification that makes it merely instrumental to truth is also vulnerable to this objection. Such views are common among non-reliabilists as well. See, for example, Bonjour: "The basic role of justification is that of a *means* to truth, a more directly attainable mediating link between our subjective starting point and our objective goal." Quoted in David (2001), p. 152. See David, pp. 151-152, for a sample of prominent epistemologists who take the goal of truth to explain, in some way, the value of justification.

⁵⁵Kvanvig (2003) ch. 3; Zagzebski (2003)

⁵⁶See Kvanvig (2003) ch. 5; Kaplan (1985); Williamson (2002).

⁵⁷Kvanvig (2010), ch. 5

⁵⁸For this reason, Kvanvig suggests that epistemologists should focus less on knowledge than on *understanding* as the central cognitive value. Ibid. pp. 185-203. Allison Hills has argued that, in the context of ethics in particular, we ought to aim for understanding, rather than for knowledge. (2010) ch. 8-10.

that the most promising accounts of the nature of knowledge—the ones with the most potential to resist Gettier-style counterexamples—are “so gerrymandered and ad hoc” that they leave mysterious why we should care about achieving knowledge.⁵⁹ If this argument is successful, and if we maintain our assumption that knowledge is Gettier-proof, then a formulation of the challenge that threatens knowledge has no force unless it simultaneously threatens something else of value.

Given these controversies, those posing the epistemic challenge for realism ought to address whether and how the threatened epistemic status, as they conceive of it, has value. Only explicit attention to these questions will assure us that the challenge constitutes a genuine threat to evaluative thought, realistically conceived. In fact, the most complete assurance will depend on addressing whether the status has value in *evaluative* thought. After all, if the relevant epistemic status has no value in *evaluative* thought—because, for instance, the characteristic aims of evaluative thought differ in important ways from those of scientific thought—the challenge will fail just as completely as if the epistemic status had no value in any thought.⁶⁰

A case study: Kieran Setiya on the epistemic challenge

Unfortunately, particular formulations of the epistemic challenge often leave these issues untouched—they rarely address the value of achieving the epistemic status under threat. The importance of the relevant status, whether in general or in evaluative thought, remains implicit. In some cases, the epistemic notions in play might reasonably lead one to doubt that they identify anything of value at all.

Take, for example, the epistemic challenge as it appears in Kieran Setiya’s *Knowing Right From Wrong*. Setiya argues that reliability about the realist’s independent values could only be accidental, and so that it is impossible for us to *know* anything about such values: “inexplicably true belief” is

⁵⁹Kvanvig (2003) p. 115. See also Williamson (2002) pp. 30-31.

⁶⁰We should be aware that evaluative thought is distinctive, and that its purposes may be served better by achieving some cognitive status different from knowledge or epistemic justification. See, e.g. Hills (2010) pp. 169-187.

a “threat to ethical knowledge.” He suggests that this matters because of the *value* of evaluative knowledge. To concede that such knowledge is impossible for us would be “shocking,” and “too great a concession.”⁶¹ However, Setiya’s partial account of knowledge leaves room for doubts about its value, both in general and in evaluative thought specifically.

In declaring accidental reliability incompatible with knowledge, Setiya defends a particular version of the common claim that knowledge requires non-accidentally true belief. Such a requirement is meant to avoid Gettier-style counterexamples to an account of knowledge. Such cases feature someone who believes a proposition (call it *p*) and has what we would call very good reasons⁶² for believing that *p*; furthermore, *p* is true. However, they seem not to have knowledge, because they arrived at their true beliefs by accident, relative to the truth.⁶³ Their beliefs are “accidentally true,” and so they do not know that *p*. To count as knowledge, our beliefs must be *non-accidentally true*.

In mounting his version of the epistemic challenge for realism, Setiya argues that knowledge requires that we are not *reliable* by accident. He expresses this as a necessary condition on knowledge:

⁶¹(2012) pp. 88, 138

⁶²They have “justification.” As Gettier notes, having justification for a belief is typically assumed to be compatible with the belief being false. This feature leaves room for the kind of slippage involved in some Gettier cases. (1963) p. 121.

⁶³For example, unknown to them, the reasons that they believe *p* are unrelated to what makes *p* true. A case of this kind is the following—Because Sarah told him that she recently purchased a Ford, Luis justifiably believes that Sarah owns a Ford. On this basis, he infers that either Sarah owns a Ford or Obama is in Boston. As it turns out, Sarah does not own a Ford; she has transferred ownership to her aunt. But his second belief is still true, because Obama just happens to be in Boston. Since the reasons he believes it are unrelated to its truth, Luis doesn’t know that either Sarah owns a Ford or Obama is in Boston. (A variation on Gettier (1963) pp. 122-123.) Or, for example, certain features of their context make alternatives to *p* relevant that their evidence does not rule out. A case of this kind is the following—Yuan is driving down Mulholland Drive, and he happens to see a dog sitting on the side of the road. On this basis, he believes that there’s a dog sitting there. He’s right; there is a dog sitting there. But unknown to him, this stretch of Mulholland Drive is being used to film a live action version of *One Hundred and One Dalmatians*, and the route is still teeming with realistic dog-replicas. Since Yuan’s evidence doesn’t rule out the relevant possibility that what he saw was a dog-replica, he doesn’t know that there’s a dog on Mulholland Drive. (A variation on Goldman (1976) pp. 772-773.)

Condition K: When S knows that p, she knows it by a reliable method, and it is no accident that her method is reliable.⁶⁴

By an “accident,” Setiya means an *explanatory accident*. In this sense of accident, one fact is accidental with respect to another when there is no explanation of either fact in terms of the other. The fact that I use the reliable method of forming beliefs based on perception is no accident, with respect to the fact that perception is reliable, if there is an explanation of the former fact in terms of the latter fact, or vice versa.⁶⁵

Importantly, it is the subject’s *use* of a reliable method that must be no accident. To make this clear: suppose that Marie knows the approximate age of a woolly mammoth fossil through her use of the method of carbon dating. According to condition K, it follows that carbon dating is a reliable method *and* that it’s no accident that *Marie uses* a reliable method. Condition K does not (merely) say that, for Marie to have knowledge, the reliability of *the method* she uses—that is, the reliability of *carbon dating*—must be no accident with respect to the truth. Rather, it says that for Marie to have knowledge, the fact that *Marie uses* a reliable method must be no accident. One’s use of a reliable method is no accident with respect to its being reliable when *either* one’s use of the method explains its reliability, *or* the method’s reliability explains one’s use. As Setiya puts the point: “What is required is an explanatory connection between the use of [the method] and its being reliable, where the explanation may be formal, final, or efficient” (Ibid. 99).

We can now state the epistemic challenge for realism in Setiya’s terms. He argues that, given the special features of the realist’s independent non-natural evaluative facts, such facts can neither explain nor be explained by our use of certain methods of making up our minds about value. On

⁶⁴By a “reliable method,” Setiya means (at least) a method of forming beliefs that results in a true belief in a high enough proportion of cases. He interprets this notion further, in order to avoid worries about individuating methods. Since what follows will not depend on these details, I will not rehearse them here. Ibid. p. 96.

⁶⁵Setiya suggests that quite different types of explanation would satisfy condition K: any of three of Aristotle’s “four causes,” running in either direction, would do. Without further explanation, Setiya states that the fourth, material causation, “seems irrelevant to K.” Ibid. p. 98.

the one hand, the realist holds that the independent evaluative facts are “irreducible and causally inert,” and so cannot enter into causal or functional explanations of why we use our particular methods. On the other hand, “since the facts are constitutively independent of our beliefs, our reliability cannot be constitutively explained.” The realist’s assumptions, then, “leave no room for the satisfaction of K by efficient, final, or formal cause.”⁶⁶ Once again, this matters because condition K is a necessary condition on knowledge. The failure to satisfy K makes “ethical knowledge impossible” on a realist view. And we ought to hope for such knowledge: to deny our capacity for it would be “shocking.”⁶⁷

Yet it is unclear why we should find the impossibility of knowledge in *this* sense shocking. The value of knowledge plays no role in the definition of condition K, nor does it play any role in Setiya’s arguments that condition K is necessary for knowledge. These arguments consist primarily of thought experiments meant to elicit the intuition that, sometimes, when a belief falls short of knowledge, it does so through violating condition K. In one series of cases, Setiya suggests that a person’s reliable disposition to form true beliefs cannot yield knowledge if that disposition is implanted in her without regard to its reliability:

Suppose you aim to implant in me, mechanically—not through persuasion or evidence—the disposition to believe *p* as soon as I am able to comprehend it. You have no idea whether this disposition is reliable; you simply want me to have it. If you are sufficiently powerful, you may succeed...[Even if *p* is a necessary truth,] it is an accident that my belief is true.⁶⁸

In further iterations of the case, Setiya makes clear that, however sophisticated one’s belief-forming disposition, its truth-indifferent source prevents the resulting beliefs from counting as knowledge. Such a belief does not count as knowledge, even if, by chance, the powerful programmer

⁶⁶Ibid. pp. 111-112

⁶⁷Ibid. pp. 112, 138

⁶⁸Ibid. p. 92

gives one the disposition to believe p on the basis of evidence that entitles one to believe p .⁶⁹

It would take us off the track of our inquiry to quarrel with Setiya's intuitions about what counts as knowledge. We can set these intuitions aside in making this point: proceeding in this way is unsatisfying—it does not carry conviction. We are meant to find the loss of evaluative knowledge, formulated in terms of condition K, “shocking.” But this reaction makes sense only if we have some reason to think that, whatever highly valuable epistemic state knowledge turns out to be, satisfying K is necessary to achieve it. Our intuitions about the extension of knowledge do not, by themselves, give us such a reason. As has been remarked in other contexts, such intuitions, far from clarifying the value of knowledge, tend to lead us toward gerrymandered, ad hoc accounts that leave its value inscrutable.⁷⁰ We have, so far, no reason to care about knowledge formulated in terms of condition K, and, in particular, no reason to care about *evaluative* knowledge formulated in terms of condition K.

Indeed, by considering the sorts of cases Setiya offers in support of his necessary condition on knowledge, one could reasonably doubt that a belief that satisfies K has greater value than a belief that does not satisfy K, but is otherwise the same. Consider a highly talented neuroscientist, who programs a subject S to believe p when presented with compelling evidence that p . He does so not because he wants S to believe the truth, but rather on a whim. Presumably Setiya would draw the conclusion that S cannot know that p , because there is no explanation connecting S's use of a reliable method to its reliability. But now suppose that, later in life, the neuroscientist comes to care about the truth. He then programs another subject S* to use an (arguably) identical method: that of believing p when presented with compelling evidence that p . The only difference is that this time, the neuroscientist does it *because* he wants S* to believe the truth.⁷¹ S*'s method seems to satisfy condition K: she uses her method because the neuroscientist wanted her to believe the truth. But it is hard to see

⁶⁹Ibid. pp. 94-95

⁷⁰Agan, see Kaplan (1985); Williamson (2002) pp. 31-32; Kvanvig (2003) ch. 5.

⁷¹Suppose also that p is a conceptual truth, and that the neuroscientist believes p in the usual way one comes to believe conceptual truths, and so by a method that itself satisfies K. (Setiya offers an account of such a method. (2012) pp. 104-108.)

why, as a result merely of a change in the motives of this third person, we should think S*'s belief that p is of any greater value than S's belief that p . Such a change need not result in any difference in the truth-regarding motives of S and S*, nor need it imply any difference in their methods, beyond this historical difference in their causal origins. Why then think that the neuroscientist's change of heart, on its own, makes S*'s belief more valuable?

One might wish to compare the case of the neuroscientist and his subjects to cases of beliefs gained through testimony. In the latter cases, the value of a belief might well depend on the motives of its source, even if the belief the inquirer gains appears otherwise the same. For example, if my informant intends to deceive me, this may devalue even a true belief that I form on the basis of her testimony. But, first, this is a hunch about testimony: in order to use this analogy, we need this hunch to be articulated and substantiated. Second, and more importantly, the analogy is weak: the process of mechanical transmission of some method of belief formation has too little in common with transmission of knowledge by testimony. It's far from clear that the analogy could be sustained once we understand why the value of belief gained through testimony depends on the motives of one's informant. Setiya does not give us the resources to address these issues.

Absent more information about the relationship between satisfying K and achieving a valuable epistemic status, our desire for the latter does not translate to a desire for the former. Motivating acceptance of condition K by appeal to our intuitions about the nature of knowledge, rather than in terms of the value of evaluative knowledge, does not engage with this concern. If we want to carry out the epistemic challenge successfully, we must rely more directly on a characterization of the value of the epistemic status in question.

1.4 Conclusion

Setiya's version of the epistemic challenge illustrates how the prima facie value of an epistemic status is sometimes invoked to motivate our concern

with the challenge, while being overlooked in its detailed formulation.

At the very least, for any epistemic status purportedly under threat in a framing of the challenge, we should ask whether, on the challenger's account of it, it would have value in evaluative thought. If we conclude that it would not have such value, we should reject that formulation of the challenge as without force. This extends to purportedly necessary conditions on epistemic justification or knowledge. When presented with such a condition, we should ask whether it is plausibly necessary to achieve a valuable epistemic status. If it is not, we should once again reject the formulation as without force.

In declining to address formulations of the challenge that employ an apparently valueless epistemic status, we need not assume that the status has been misunderstood. If one makes this assumption for some epistemic status—for instance, if one agrees with Zagzebski that “knowledge has to be defined as something we value”⁷²—one will reject such formulations because they rely on the wrong account of that status. If one does not make this assumption—for instance, if one is open to agreeing with Kvanvig and Mark Kaplan⁷³ in thinking that knowledge has no value beyond that of true, justified belief—one can reject the formulation because nothing hangs on answering it.⁷⁴

⁷²Zagzebski (2003) p. 26

⁷³See Kvanvig (2003), ch. 8; Kaplan (1985).

⁷⁴I take for granted here what I argue in section 2.1 above, namely that our evidence that we have evaluative knowledge, or epistemically justified evaluative belief is ambiguous—there is reason to doubt that we have such knowledge, and the decisiveness of our commitment to many of our evaluative beliefs may well have non-epistemic sources. If this was denied, however, then one could urge the need to answer the challenge out of a concern for the truth.

Chapter 2 The Epistemic Challenge and Avoiding Accidents

2.1 The value of avoiding accidentally true belief

Suppose that the non-naturalist realist cannot answer the “epistemic challenge for realism.” Suppose she must admit that, if we understand evaluative truth as she does, our evaluative reliability would be a coincidence. What follows? Some suggest that this would put knowledge, epistemically justified belief, or a related cognitive status out of reach in evaluative thought. In the last chapter, I argued that this does not yet provide us with an adequate motivation to answer the challenge. We need, in addition, a reason to regard achieving such statuses in evaluative thought as distinctively valuable. As a case in point, I argued that Setiya’s careful formulation of the epistemic challenge as a threat to evaluative knowledge carries no weight, because he does nothing to establish the value of achieving knowledge, in the peculiar sense he presupposes. To succeed in motivating the epistemic challenge through appeal to these statuses, we need to think *both* that (a) achieving the relevant epistemic status in evaluative thought requires answering the challenge, *and* that (b) doing so has value in evaluative thought.

For the purpose of motivating the epistemic challenge, we need not be interested in whether the challenge threatens these traditional epistemic statuses, in particular. Our question is rather whether it threatens *something* of great cognitive value in evaluative thought. I will therefore frame what follows in terms of cognitive statuses whose relation to the epis-

temic challenge is more transparent. Since an answer to (a) above is fairly straightforward for such statuses, we can then focus more squarely on (b): to what extent does falling short of this status threaten anything of value in evaluative thought?

Take, for example, accidentally true belief. For some sense of “accident,” the inability to answer the challenge straightforwardly implies that our evaluative beliefs are accidentally true. Absent some explanation linking the evaluative truth and our evaluative psychology, the alignment between the evaluative truth and our evaluative psychology is an accident or, equivalently, a coincidence.

A number of contributors to this debate suggest that evaluative beliefs are made worse by being *true by accident*. As we have seen, this is precisely the condition on *knowledge* that Setiya takes to demand some answer to the challenge. Others also warn of the danger of accidentally true belief, but frame the threat in terms of beliefs that could easily have been false, or that are unlikely to be true. Justin Clarke-Doane poses the challenge in the former terms; if our belief could easily have been false that presents an “undermining defeater,” in that it “gives us a reason to give up [that] belief” without giving us evidence of its falsity.¹ Formulations of the challenge that emphasize the realist’s commitment to a highly unlikely coincidence of fact and belief, such as those offered by Street and Enoch, take the latter tack.²

Our question is then: is achieving non-accidentally true evaluative belief, in the sense at issue in the challenge, of value? Just for the purposes of addressing this question, let us take for granted that, all else equal, having a true belief about value is valuable, either in its own right or as a means to something else. All else equal, a true belief about directions will get you to your destination when a false belief wouldn’t. Something similar seems true of a true belief *about value*. All else equal, a true belief about value will enable you to act in valuable ways by, for example, leading you to respect, appreciate, or pursue valuable objects when a false belief wouldn’t.

¹Clarke-Doane poses the challenge specifically for T.M. Scanlon’s variety of non-naturalist realism. (2017), p. 849.

²Enoch (2009) p. 418, fn 12.

Conceding this, then, we should note that many valuable things are made no worse by being brought about by accident. Meetings between soulmates are valuable. Accidental meetings between soulmates are not worse than meetings arranged deliberately (and, indeed, they may be better). Poetic genius is valuable. Accidental poetic genius is not worse than poetic genius deliberately fostered (and, indeed, it may be better).³ So, given that true evaluative beliefs are themselves valuable, we can reframe our question in these terms: does true evaluative belief have *less* value when arrived at by accident?

There are at least three promising avenues for arguing that, in general, a belief that is accidentally true has less value than one that is non-accidentally true: by lessening its admirability, by increasing its vulnerability to loss, and by rendering it unreasonable or irrational as a position on the evaluative truth. I will address each in turn, in order to show that none would result from a failure to answer the epistemic challenge. At least as far as these valuable features of evaluative belief are concerned, the threat posed by the epistemic challenge is illusory.

2.2 Less worthy of admiration

We frequently treat skills and skillful actions as worthy of admiration. We admire the skill of the ping pong player who can control the ball no matter its speed or angle of approach, and we admire the particular point that results from that skill. We admire the artistic vision and interpersonal finesse of the director who can produce a great film no matter the quality of her script or cast, and we admire the particular aesthetic choice that results from that skill.⁴ These do not appear to be isolated cases; skills and skillful actions generally inspire admiration.

³I make a parallel point here to one made by Zagzebski, about the difference between goods that are achievements and those that are not: "Some goods are just as good if we do not have to work for them—for example, good health and a safe environment—and some may even be better if we do not have to work for them—for example, love and friendship." (2003), p. 23.

⁴I have in mind the directors such as Douglas Sirk, whose great films were produced under the Hollywood studio system, with scripts and actors not always of their choice.

Virtue epistemologists have suggested that we think of the virtues (and intellectual virtue in particular) as skills whose skillful manifestations are fitting actions and attitudes. As is generally the case for skills and skillful actions, we admire the virtues and the fitting actions and attitudes that arise from them. We admire kind acts that arise from a person's correct sense that, given the circumstances, kindness is called for, and we admire true judgments that arise out of someone's correct sense that, given the evidence, this judgement is called for. According to many virtue epistemologists, knowledge is analagous to the skillful point in ping pong, in that it is a particular manifestation of a skill or excellence. In both cases, what we admire is in part *success*, and in part the *exercise of skill*. A skillful shot is more admirable if it scores the point, and the point is more admirable if it is scored skillfully. Similarly, skillful weighing of evidence is more admirable if it leads to belief in the truth, and vice versa.⁵

This offers us a way of understanding why a true belief has less value when arrived at accidentally: only non-accidentally true beliefs are instances of success *through* skill.⁶ In a case of non-accidentally true belief, the believer exercises her intellectual virtue—by, for instance, correctly weighing her evidence, or reasoning well—and *through* this exercise of skill arrives at the truth. In coming to her true belief non-accidentally, she is like a ping pong player who scores the point through skill. In coming to believe the truth by accident, by contrast, she is like the player who scores the point through luck. She has arrived at a good result, but not as a result of her skill. This may happen even if she acts skillfully. For instance, suppose that, a split second after the ping pong player hits the ball in her signature skillful way,⁷ a bird flies by and nicks the speeding ball with its wing, knocking it off course. Nonetheless, the ball lands on

⁵This is a simplified take on this literature, which is often framed in terms of the value of getting “credit” for one's belief, as one gets “credit” for an action that is attributable to one. The notion of credit originates in an analogy between virtuous successful action and justified true belief. I here frame this in terms of the value of an action or attitude that results from skill, and is therefore admirable. See Sosa (2007) ch. 9; Greco (2000); Riggs (2002) pp. 103-106; Kvanvig (2003) pp. 81-99.

⁶As John Greco puts this point: “According to the account I have in mind, knowledge is a kind of success through virtue. Put another way, knowledge is a kind of success through virtuous agency.” (2009) p. 318. See also Greco (2007).

⁷Suppose that a keen-eyed and experienced observer of ping pong could recognize this as an excellent stroke, one that arose from skill and was highly likely to be successful.

her opponent's side, and he fails to return it. In this case, the player succeeds, but not through her skill.⁸ This makes the player's performance less admirable than it would have been, had she scored the point through skill.⁹

Here, then, is a way to frame the stakes of the epistemic challenge for realism, by way of the notion of accidental truth. The challenger can suggest that, if we cannot explain our being reliable about non-natural, independent values, then all of our beliefs about such values are accidentally true, and so less admirable, resulting as they do from chance rather than from skill. So, on a realist account, our beliefs about value are uniformly less worthy of admiration.

Does this framing succeed? In exploring objections, I will make two assumptions. First, I assume that we are not seriously mistaken in our admiration of skills and skillful actions, and so that these generally merit admiration. Not that we admire, or ought to admire, all skills and their manifestations: these may have further features, such as being morally abhorrent, or trivial, or gained only through wasting a great deal of time, that make them not admirable on the whole. This is compatible with the claim I will assume: that all else equal, skills and exercises of skill are worthy of admiration. Second, I assume that if something *merits admiration*, it is to that extent *good*, and that if something merits *more* admiration than something else in some respect, it is in that respect *better*. That something is worthy of admiration is commonly treated as a sign of its value: in fact, some philosophers have recently suggested *analyzing* value in terms of meriting positive responses, such as admiration, preservation, promo-

⁸Compare this case of Ernest Sosa's: "Take a shot that in normal conditions would have hit the bull's-eye. The wind may be abnormally strong, and just strong enough to divert the arrow so that, in conditions thereafter normal, it would miss the target altogether. However, shifting winds may next guide it gently to the bull's-eye after all. The shot is then accurate and adroit, but not accurate because adroit (not sufficiently). So it is not apt, and not creditable to the archer." (2007), pp. 22-23.

⁹Here is Greco again on this point: "We credit success through ability more than we credit mere lucky success. But we also value success through ability more than we value mere lucky success. In fact, there is a long tradition on which virtuous success, i.e., success through virtue or excellence, is identified as the highest human good". (2009), p. 58.

tion, and respect.¹⁰ Without making any assumption about the correct analysis of value, I will take for granted the weaker claim that if something merits admiration, to that extent it has value.¹¹

Even if we take these claims for granted, the virtue epistemologist's account of the nature of knowledge has inspired several objections, some of which also pose problems for this way of motivating the challenge. First, one might doubt that non-accidentally true belief is generally a manifestation of skill, and so deny that such belief is generally admirable in the way of success through skill. Jennifer Lackey has raised this objection against virtue epistemologists, urging that skill has a minimal role in acquiring many beliefs that are non-accidentally true. Lackey focuses especially on beliefs formed through reliance on testimony, arguing that what explains our success in such cases is usually not our intellectual skill, but rather that of our informant. Take, for example, the case of Morris, a first time visitor to Chicago who alights from his train and immediately asks a random passer-by for the location of the Sears Tower. This passer-by is a Chicago native and knows the city well, and so correctly informs him of its location. Relying on this testimony, Morris forms a true belief about the location of the Sears tower.¹² Here, Lackey suggests, what primarily explains Morris arriving at the truth is not his skill, but rather the fact that his informant has deep familiarity with the city of Chicago. While Morris gains a non-accidentally true belief, this has little to do with his own skill—his “reliable cognitive faculties” are not “the most salient part of the cause” of his true belief.¹³ Lackey emphasizes that we often form true beliefs in this way; these are paradigm cases of achieving knowledge, and so achieving non-accidentally true belief.

If non-accidentally true beliefs are not generally manifestations of skill, and so not generally admirable for that reason, we do not yet have a motivation for answering the epistemic challenge. Even if we answer the challenge, and so demonstrate that our true evaluative beliefs are not accidentally true (in at least one sense), this does not show that such be-

¹⁰See Scanlon (1998); Rabinowicz and Ronnow-Rasmussen (2003, 2004), quoted in Hutchinson (2019).

¹¹Thanks to Jay Wallace, for helpful suggestions about how to frame this point.

¹²(2007) p. 352

¹³Ibid.

iefs are manifestations of skill. We would need something more than a demonstraton of this to suggest that our evaluative beliefs have such a valuable feature.

John Greco has a response to Lackey's argument that is fairly convincing. He argues, first, that Morris *does* manifest an intellectual skill in forming his true belief: that of reliably receiving testimony.¹⁴ Even the ability to do the right thing with testimony requires skill. We underestimate what is needed, because we're so used to this transaction, but it's actually quite an impressive human ability, to be able to inquire after and rely appropriately on testimony. Second, Greco argues that in estimating the role of Morris's intellectual skill in gaining the true belief, a comparison with the contribution of his informant is not really to the point: their endeavor is cooperative, and both contributions serve their epistemic purpose. Greco compares knowledge gained through testimony to the case of a soccer goal scored cooperatively, through an exemplary pass and a competent nudge through the goal posts. The competent nudge is admirable, because it is appropriately involved in the scoring of the goal: "the purposes of soccer playing are well served by the reliable execution of easy goals."¹⁵ We might expand on Greco's point, by pointing out that there is nothing to prevent us from regarding Morris' knowledge as a joint product. Why not evaluate its admirability while taking into account all the skills that produced it, the way we often evaluate performances that are a team effort? If we regard beliefs formed through testimony in this way, non-accidentally true belief will turn out skillful after all.

Rather than objecting that some cases of *non-accidentally* true belief do *not* count as success through skill, one might object that some cases of *accidentally* true belief *do* count as success through skill. That is, one sometimes arrives at the truth through skill, but nonetheless arrives there accidentally, and so any defect in such performances cannot be explained as a failure to succeed *through* skill. A prominent example is Alvin Goldman's fake barn case. In this case, Henry, a traveler driving through an unfamil-

¹⁴This is a skill Morris would lack if he were inclined to accept directions even from very small children and highly intoxicated people, or if he were highly distrustful of trustworthy informants, so that he tended to believe the opposite of whatever he was told.

¹⁵(2007) pp. 64-65.

iar countryside, looks carefully out his window, sees a barn, and thereby believes truly that it is a barn. Little does he know, this is the only real barn in a hundred mile radius; the rest of the countryside is dotted with deceptive barn facades, rather than actual buildings. If his true belief is less admirable than it might be, the problem is not the belief's relation to Henry's skill as a perceiver. Rather, the problem is Henry's environment—he is unwittingly surrounded by fake barns, with only one real barn in the vicinity.¹⁶ One might argue by analogy that we ought to regard Henry's skillful formation of a true belief with just as much admiration as we would if his circumstances weren't deceptive. We can construct cases of successful athletic performance through skill, in which a deceptive environment does not lessen our admiration for the performance. Take again our ping pong player, and suppose her, unknowingly, in a bizarre situation where her opponent has altered the airflows above the table, such that her shot will go awry unless it happens to pass through the only unaffected square inch. She gets lucky in the angle of her shot, and scores a point flawlessly. Despite her good luck, she presumably scores the point through skill, and her performance is just as admirable as it would have been without such machinations.¹⁷

If this objection is successful, then the inability to answer the challenge does not suggest that our evaluative beliefs are any less admirable. True beliefs like Henry's are, in some clear sense, accidentally true—it is just his good luck that he happened to focus on the one real barn in the facade-filled countryside. And yet they are instances of success *through* skill. Perhaps our evaluative beliefs, despite being accidentally true, also count as instances of success through skill. In that case, they are not less admirable as skillful performances. We once again lack a clear motivation to answer the challenge.

Greco also has a response to this objection that seems worth considering. He suggests that there is a disanalogy between the case of Henry and the case of the successful athlete. In the case of the ping pong player, he would contend, the admirability of her skillful performance does not suffer, because we do not consider it part of her skillfulness that she detect

¹⁶(1976) pp. 772-773

¹⁷Kvanvig (2003) and Whitcomb (2007) p. 104, both make this point against the virtue epistemologist. Whitcomb quoted in Greco (2007), p. 84.

this sort of trickery. It is not a part of the skill of playing ping pong to detect and respond differentially to deliberate interference with airflows. In the case of the general skill of forming beliefs through perception, however, the ability to detect and respond differentially to deceptive circumstances *is* thought to be part of the skill. As Greco puts the point, framed in terms of the sport of archery:

The ability to hit a target, like any ability, is defined relative to conditions that are appropriate for that sort of ability. In particular, we do not require that an archer is reliable (relative to an environment) in conditions involving arrow-weighting hoodlums...The situation is different with intellectual abilities, however. Here it does matter how *S*'s performance would be affected by information-tampering hoodlums in the environment. Given the nature and purpose of our knowledge-related abilities, it is centrally relevant whether *S* can reliably negotiate such aspects of her environment.¹⁸

This presumably does not imply that a perceiver in Henry's deceptive circumstances is *blameworthy*, or even criticizable, for failing to make sure he wasn't in fake barn country. As when the bird's flight interferes with one's shot, one's performance may be less admirable through no fault of one's own. Rather, the suggestion is that it is reasonable to view Henry's performance as a defective expression of his skill, given the nature and purpose of intellectual skills. His true belief, although formed through the exercise of perception, is not formed through the skillful exercise of perception. A full reply along these lines would need to assure us that we have reasons to adopt this account of skill, independently of wanting to defend this way of understanding the value of accidentally true belief. I will not attempt a final evaluation of Greco's reply, however, because I think another feature of the virtue epistemologist's account renders it unhelpful in our context.

This account of why accidentally true belief has less value than non-accidentally true belief, whatever its virtues in a general investigation of this question,

¹⁸(2007) pp. 66-67

cannot help in motivating the epistemic challenge for realism. This is because, according to the epistemic challenge, luck intervenes in an entirely different place.

Return to the ping pong player whose successful shot is due to the lucky chance that the bird's wing sent her ball to her opponent's side, rather than to her own skillful performance. Here luck intercedes between the exercise of skill and success. In the analogous case, luck would intercede between the exercise of our dispositions to form evaluative beliefs and our formation of true evaluative beliefs. But this is not the way that luck intercedes, according to the epistemic challenge. The epistemic challenge concerns *not* the link between our evaluative dispositions and our evaluative beliefs, but *rather* something further back in the chain—the link between our evaluative dispositions and the realist's independent evaluative truths. If our beliefs are accidentally true, this is because it is an accident that we have reliable dispositions at all. The problem for the realist is to explain our reliability in a satisfactory way—if she cannot, it is our reliability that is lucky. Nonetheless, relative to the reliable evaluative dispositions we in fact possess, it is not a matter of luck that our performances are successful.

Could we perhaps adapt our account of the badness of accidental truth to this way of being accidentally true? We might, if successful performances through skills acquired accidentally were less admirable. But here the epistemic challenger faces an obstacle: in fact, skills acquired through luck and expressions of such skills do not generally inspire less admiration than skills acquired in other ways. We do not object to skillful performances by child prodigies, or other individuals who gain skills in lucky or inexplicable ways. Furthermore, it does not seem that we *should* find these less admirable. Skills and their manifestations do not appear to deserve any less admiration when they come about in ways that we cannot explain.

However promising the notion of success through skill might prove for explaining why, in some cases, arriving at the truth by accident has less value, its approach cannot help to motivate the epistemic challenge. This is because the threat of accidental truth at issue in the challenge arises in a different place. According to the challenger, it is our *being skillful* that

would be accidental, on a realist view that does not answer the challenge. But having a skill by accident, in the manner of a prodigy, does not seem less admirable, and neither do the manifestations of such a skill.

Yet given such an accident, can we regard our dispositions to believe as skillful, or reliable? Doesn't such an accident shake our confidence that we reliably believe the truth about the evaluative facts? Perhaps the problem with accidentally true belief is not that it is less admirable, but that it undermines our beliefs about value. We will take up one way to understand this possibility in the next section.

2.3 More vulnerable to loss

When one's belief is accidentally true, something about the way one arrived at it left its truth up to chance. I might make up my mind on some question using methods that are pure guesswork. In such cases, I do not make up my mind in a way that assures the truth of my belief. Rather, I leave the truth of my belief up to chance. To take an example, I might truly believe that my bed is covered with leaves, based on a visual experience that is actually a hallucination. Although I may be unaware of this, I arrive at my belief in a way that does not assure its truth, but leaves its truth up to chance.

What is the problem with leaving the truth of one's beliefs up to chance? One problem is that, in a number of situations, if one discovers this fact about a true belief one holds, one will not think it likely to be true and so give it up.¹⁹ Assuming, as we have, that all else equal it is good to have a true belief, we will then be giving up something good. Take again the case of hallucinating and thereby truly believing that one's bed is covered with leaves. If I later discover the origin of my belief in a hallucination and have no further information, I will probably give up the belief, because I will not think it likely to be true.

A parallel point applies to others who stand to benefit from reliance on my true belief. After all, one person's true belief about some subject has

¹⁹See Williamson (2002) pp. 6-7.

value not only for her, but also for someone else who wants to find out the truth about that subject. Individuals rely heavily on others in finding out about the world.²⁰ Our needs as inquirers might even explain the shape of our notions of knowledge and justification, as William Craig has suggested in *Knowledge and the State of Nature*. Craig offers a state of nature account of the concept of knowledge, arguing that its original function was identifying good informants. In such a state of nature, he suggests, an inquirer might want the truth about some matter—say, whether or not a given plant is safe to eat, or whether or not some friend of his is angry at him. The most effective way for the inquirer to find out the truth (about whether or not *p*) may well be to ask someone else with a true belief (about whether *p*).

The features of our present concept of knowledge bear the marks of this basic situation, Craig suggests. To be of use to such an inquirer, an informant should not only have a true belief²¹ about whether *p*, but also a further feature, discernible to someone who does not yet have a true belief about whether *p*, in virtue of which the informant is *likely to be right* about whether *p*.²² In the cases above, further features could be that the informant is an old sage with years of experience with local plant-life, or a close confederate of one's friend who is skilled at interpreting his moods. That the informant has such a discernable feature, distinct from believing the truth, allows an inquirer ignorant about whether *p* to *rely* on that informant, and so form a true belief about whether *p*.

If the informant arrived at the truth by accident, however, a parallel danger exists to that noted in the first person case. That is, there is some

²⁰W.K. Clifford: "no one man's belief is in any case a private matter which concerns himself alone. Our lives are guided by that general conception of the course of things which has been created by society for social purposes. Our words, our phrases, our forms and processes and modes of thought, are common property, fashioned and perfected from age to age...Into this, for good or ill, is woven every belief of every man who has speech of his fellows." (1879), p. 127.

²¹Her attitude toward *p* should be something like belief: it cannot be so wavering, or so inarticulate, that she would not or could not declare it to an inquirer. (1990) pp. 12-13.

²²Craig suggests that this is the original role of the condition on knowledge known as "warrant"—the additional feature needed in order for a true belief to count as knowledge. (1990) pp. 18-19.

information the inquirer could discover, in the light of which she would not regard the informant as likely to have the truth. An informant may well believe the truth about whether an inquirer's friend is angry with her, and have features that make her likely to be right about this—the informant may be the friend's longtime confederate, and have seen the relevant interaction, and so (given just this) be likely to believe the truth about whether he is angry with the inquirer. But, given that her belief is accidentally true, she will also have features in light of which she is not likely to be right—for instance, she may have come to her true belief by relying on a false belief about the relationship between the inquirer and the friend, and so (given both this, and the features already mentioned) not be likely to believe the truth about whether he was angry. If the inquirer discovers this, she will not rely on the informant, and will miss out on a true belief.²³

The above suggests that accidentally true belief is less valuable than non-accidentally true belief, because the former is *more vulnerable to loss*, in an extended sense.²⁴ This is because the accidental truth of a belief implies the existence of some information in light of which it will appear that the belief is not likely to be true. The discovery of such information threatens the believer with loss of a true belief, and the inquirer with losing out on a true belief.

Here, then, is another way to frame the stakes of the epistemic challenge for realism, using the notion of accidental truth. If we understand values as the non-naturalist realist does, the challenger urges that our reliability about such values could *only* be an accident. Street, for instance, says that our being reliable about such values would be “a happy coincidence between the realist's independent evaluative truths and the evaluative directions” actually taken in the development of human psychology,

²³Craig suggests that the concept of knowledge, as we use it now, is “objectivised,” so that we deny an individual knows unless she has features such that any reasonably well-informed inquirer would rely on her in forming a belief about whether *p*. It is valuable for someone with a true belief about whether *p* to inspire reliance among people with many different sources of information. *Ibid.* pp. 87-88.

²⁴This is related to Socrates' reflections on the the value of true opinion: although “true opinions, as long as they remain, are a fine thing and all they do is good, [] they are not willing to remain long...so that they are not worth much until one ties them down by [giving] an account of the reason why.” Plato, *Meno*, 97e-98a.

as a result of influences such as natural selection. But given this, information exists suggesting that all of our evaluative beliefs are not likely to be true. The inability to answer the challenge gives us this information. As Street puts the point, an inexplicable coincidence between the independent evaluative facts and the course of our actual evaluative development would require “a fluke of luck that’s not only extremely unlikely...but also astoundingly convenient to the realist.”²⁵ Those inclined toward realism might carry out the reasoning involved in posing the challenge, and thereby lose their evaluative beliefs. In accepting a realist view, then, their beliefs about value are vulnerable to loss.

This framing is particularly interesting in our context; the possible loss of beliefs about value seems just out of view in a number of presentations of the epistemic challenge. For example, Enoch suggests that our being accidentally reliable about value—our having beliefs that inexplicably align with the evaluative facts—“would be just too miraculous to believe.”²⁶ Presumably a natural result of realizing that something is “too miraculous to believe” is not believing it. Writers such as Clarke-Doane suggest that the epistemic accident at issue in the challenge presents a defeater for our evaluative beliefs, giving us “a reason to give up our belief.”²⁷ Having a reason to give up one’s belief does not imply that one *will* give it up. But presumably a natural result of realizing that one has an unopposed reason to do something is doing it. Wielenberg suggests that the impossibility of moral knowledge would “have dramatic implications for moral philosophy as well as the moral beliefs and practices of most human beings.”²⁸ While he does not elaborate, it is hard to imagine how, beset with controversy as so many moral practices are already, the impossibility of moral knowledge could have such dramatic implications, unless it were to threaten us with the loss or severe weakening of our moral beliefs.

We should take care to distinguish another possible reading of these quotes. Perhaps the problem these authors have in mind is *not* that, if we understand values as the non-naturalist realist does, our beliefs about them are vulnerable to actual loss, but only that in this case we will have “a reason

²⁵ (2006), p. 122

²⁶ (2009), p. 421

²⁷ (2017), p. 849

²⁸ (2010), p. 442

to give up our belief” in such true claims. On this reading, accidentally true belief is less valuable, because we do not wish to hold a belief while also having a reason to give it up. We will explore this possibility in the next section, and our exploration will benefit from strictly separating the disvalue of *having a reason to give up* a true belief from the disvalue of *actually giving up* that belief. What remains is to deal with the danger of loss or of losing out on true evaluative beliefs, which I believe can be done briefly.

First, I think that we must conclude that the realist is in no danger of losing her evaluative beliefs as a result of failing to answer the challenge. To lose one’s evaluative beliefs in this way would be like losing one’s belief that the sun will rise tomorrow as a result of failing to answer Hume’s doubts about unobserved matters of fact.²⁹ In arguing that we have no good reason for expecting that the sun will rise tomorrow, Hume rightly took himself to be entirely powerless to undermine our belief that it will. I take it to be in the spirit of Hume to hold that, even if one’s philosophical view implies that one’s evaluative beliefs could only be true by accident, this has no tendency to lead to the loss of those beliefs. The natural path for the realist to take, in the face of realizing that her beliefs are unlikely to be true, is to continue to believe and to consider herself extremely lucky. Since that is what she will do, there is no real danger of the loss or weakening of her evaluative beliefs.

One might wonder whether the realist will *lose out* on true evaluative beliefs, if she comes to accept that, from a certain perspective, no one’s beliefs about value are likely to be true. Perhaps the realist will not be inclined to seek out good informants in order to learn more about value. This may sound familiar, because it is in fact how most people today treat questions of value. We do not typically consider evaluative questions to be appropriately settled through finding a reliable informant. There is no danger of losing out—we cannot miss out on value that does not appear to be there. (And, to the extent that we rely on the advice of others, what prevents us from regarding them as luckily reliable as well?)

I discussed two ways of filling out the concern to avoid accidental truth—perhaps accidentally true evaluative beliefs are less admirable, or perhaps

²⁹Cite

they are more vulnerable to loss, than non–accidentally true evaluative beliefs. Neither seems plausible of evaluative beliefs that are accidentally true in the sense at issue in the challenge. Lacking a satisfying explanation of our evaluative reliability does not seem to make our beliefs about value less admirable. We will, from a certain perspective, regard ourselves as prodigies in our thinking about values and reasons, but that seems no demerit. Nor does lacking such an explanation make our beliefs more vulnerable to loss. Although the possibility of loss has haunted the debate, our becoming value agnostics is not a likely result of failing to answer the challenge.

But the realist might still be in a bad situation. For example, perhaps if her reliability does not admit of explanation, she *ought* to give up her evaluative beliefs. Or perhaps lacking such an explanation is implausible or objectionable in itself. This chapter will complete our exploration of the disvalue of accidental truth, by considering its implications for what we ought to believe. We will then transition from considering the implications of the challenge for the value of our evaluative beliefs to considering the value of evaluative explanation itself.

2.4 More unreasonable

We will not abandon our evaluative beliefs, even if we cannot answer the epistemic challenge. This seems like common sense. Yet many contributing to this literature frame what is at stake in terms of the tenability of our evaluative beliefs. David Enoch, for example, suggests that without “some explanation, the correlation” between our beliefs and the truths “may just be too miraculous” or “too implausible” to believe. As Enoch makes clear later in his discussion of the challenge, he has in mind a threat to believing evaluative claims *justifiably* or *reasonably*, and only derivatively to maintaining evaluative beliefs at all.³⁰ He worries that we *ought* to abandon belief in our own reliability, if we cannot explain this correlation. This presents us with a different motivation for answering the challenge. Even if, as a psychological matter, our evaluative beliefs are not vulnera-

³⁰(2009), p. 421

ble to loss, as a normative matter, continuing to believe what one ought not believe is itself a bad thing. Further, doing so when aware that one ought not is irrational.

Why would failing to answer the challenge mean that we ought to give up these beliefs? If we cannot explain our reliability about values, then we could only have gotten onto the truth through a “happy coincidence between the realist’s independent evaluative truths and the evaluative direction” taken by our psychological development.³¹ But one might think that such a coincidence is unlikely, and so that we ought to doubt that it has occurred. If we ought to doubt that such a coincidence has occurred, then we ought to doubt our evaluative reliability. Finally, if we ought to doubt our evaluative reliability, then we ought to abandon our evaluative beliefs.

There are different ways to support this suggestion. Perhaps, as an a priori matter, coincidences are unlikely to occur. Whatever the probability that a given evaluative system is correct, and whatever the probability that we should have developed so as to grasp that evaluative system, the probability of both occurring will be lower if they are unconnected than if they are connected. If our beliefs being true would depend on a coincidence, then the truth of our beliefs is also unlikely.³² Or perhaps the metaethical realist, in particular, is committed to viewing the correct values as a small slice of the space of conceptually possible values, and so must regard our hitting upon the correct values as unlikely.³³ In either case, if our beliefs are not likely to be true, we *ought* to abandon them, quite apart from whether this is psychologically possible.

Since acting against what one ought to do is a bad thing, this would make sense of the desire to answer the challenge. Further, if we *realize* that we ought not hold onto our evaluative beliefs, perhaps by realizing that we cannot answer the epistemic challenge, we not only do a bad thing in continuing to believe, but we behave irrationally.

My response to this line of thought is closely related to a choice I made

³¹Street (2006), 122

³²Setiya (2012), pp. 71-72

³³Street (2016) pp. 316-317

in framing this investigation. It may have struck the reader familiar with this literature that I frame the challenge as a threat to the *epistemic* status of our evaluative beliefs—their status as knowledge, epistemically justified, non-accidentally true, or non-accidentally reliable—without acknowledging any threat to what we might call their merely *alethic* status—their truth or reliability. In particular, readers might have noted that, in my last chapter, I distinguished formulations of the challenge that only threaten our beliefs' status as epistemically justified or as knowledge from formulations that *simultaneously* threaten the truth of our evaluative beliefs, and went on to only address the former, and say nothing about the latter. Similarly, in discussing the disvalue of non-accidentally true belief, I assumed that our evaluative beliefs are true, and that such true belief is of value.³⁴ I have not yet, then, taken seriously that the inability to explain the reliability of evaluative thought might call that very reliability into question. I now want to explain why, ultimately, I do not take this seriously.

Of course all parties to this debate admit that the truth or reliability of our evaluative beliefs would remain *possible*, even if ultimately inexplicable. But that alone would not justify my attitude. After all, many urge that, if we admit such inexplicability, the balance of plausibility is strongly against our being reliable about value. In their writings on the epistemic challenge, Street (an anti-realist) and Enoch (a realist) take it that, if there could be no explanation of a correlation between belief and fact, one should not believe there is any such correlation at all.³⁵ This is not because such a correlation is impossible, but because it's so very unlikely. It is, in Enoch's phrase "too miraculous to believe," or, as Street describes it, "a fluke of luck that is...extremely unlikely, in view of the huge universe of logically possible evaluative judgements and truths."³⁶ According to these authors, inexplicable reliability is possible, but not something we can rea-

³⁴I did this primarily in order to isolate the distinctive badness of *accidental* truth. However, one might object that such an assumption instead leaves us unable to understand its badness; it might take all the bite out of accidental truth. If we can assume that we have the truth, then perhaps we should not care about whether we got it by accident. But we *cannot* assume that we have the truth. Only when we leave open whether our belief is true can we understand what is troubling about accidental truth.

³⁵And, perhaps, should even affirm the opposite, by positively *believing* that there is no such correlation.

³⁶Enoch (2010) p. 421; Street (2006) p. 122

sonably suppose to be actual.³⁷ Given that many in this literature take this position, continuing to regard one's evaluative beliefs as on the whole *true*, without further argument, would beg the question against an important version of the challenge.³⁸

I believe that begging this question is precisely what the realist ought to do. My contention is that even if realism (or any other metaethical account) rules out a satisfying explanation of our reliability, this does not create rational pressure on the realist (or any metaethicist) to stop regarding her evaluative beliefs as true.³⁹ If the epistemic challenge poses any kind of threat to evaluative beliefs understood realistically, it can *only* be to their *epistemic* status—their status as knowledge, as non-accidentally reliable, as epistemically justified,⁴⁰ or as theoretically sound. That is why, in this investigation, I focus on these statuses.

An analogy may help to bring out why I think the challenge is only forceful when understood in this way. In some ways, the epistemic challenger regards the metaethical realist as the proponent of the fine-tuning argument regards the atheist. Both charge their targets with positing a massive, miraculous coincidence. The proponent of fine-tuning charges the atheist with positing a coincidence between the *actual* fundamental physical laws and initial conditions of our universe and the very small subset of *possible* laws and initial conditions that allow for the development of life. That these features of our universe fall into this very small subset is just what it means for the universe to be “fine-tuned.” As Robert White describes this:

³⁷This account of the challenge might make it mysterious why authors call it “epistemic” at all. Why take a detour through such complicated notions as knowledge, epistemic justification, or non-accidental reliability, if one simply wants to say that, given realism, our beliefs about values and reasons are highly unlikely to be true? I think what follows helps to explain this.

³⁸These are *beliefs* of ours, and so regarding them as true is our typical state. Our question is whether it is legitimate to continue to regard them as true, without further defense, after being presented with the challenge.

³⁹Once, again, to believe something is to regard it as true. What I contend is that regarding our beliefs in this way is *legitimate* for the realist in the context of the epistemic challenge.

⁴⁰If epistemic justification for a belief amounts merely to the entitlement to regard one's belief as true, then this status also cannot come into question. I do not assume this here.

The universe is said to be extraordinarily 'fine-tuned' for life. The inhabitability of our universe depends on the precise adjustment of what seem to be arbitrary, contingent features. Had the boundary conditions in the initial seconds of the big bang, and the values of various fundamental constants differed ever so slightly we would not have had anything like a stable universe in which life could evolve. In the space of possible outcomes of the big bang, only the tiniest region consists of universes capable of sustaining life.⁴¹

The proponent of the fine-tuning argument contends that, on the atheist's view, there is no satisfying explanation of why our universe should happen to fall within this tiny set of possible outcomes. Without such an explanation, the universe's amenability to life is a highly improbable coincidence—it's just our good luck that the universe allows for complex life forms, like ourselves.⁴² On the theist's view, by contrast, this is no coincidence. God has *designed* the world in a way that accomodates life. The fine-tuned universe, while still perhaps a miracle, is an *intentional, explicable* miracle.

The epistemic challenge has a similar structure. The challenger charges the realist with positing a coincidence between our *actual* evaluative dispositions and the very small subset of *possible* evaluative dispositions that the realist regards as independently correct. She contends that, on the realist's view, there is no satisfying explanation of how our psychology came to align with this very small subset of possible dispositions. Given "the huge universe of logically possible evaluative judgments and truths,"⁴³ our having evaluative attitudes that are roughly on track by coincidence would be extraordinary. On the anti-realist's view, by contrast, this is no

⁴¹White (2000) p. 260. See also Collins (2008) pp. 84-85.

⁴²Here is Collins' elaboration of this coincidence: "one could think of the initial conditions of the universe and the fundamental parameters of physics as a dart board that fills the whole galaxy, and the conditions necessary for life to exist as a small one-foot wide target: unless the dart hits the target, life would be impossible. The fact that...the dart has hit the target, strongly suggests that someone set the dials or aimed the dart, for it seems enormously improbable that such a coincidence could have happened by chance." Collins (2008) p. 85.

⁴³Street (2006) p. 122

coincidence. Since the evaluative truth is a function of our actual evaluative attitudes, the latter can never go too far astray from the former. This explains our approximate reliability in believing the truth about value.

There are several interesting similarities between the fine-tuning argument and the epistemic challenge for metaethical realism.⁴⁴ However, most relevant to my current point is an interesting difference, which emerges when we compare what is treated as common ground in these literatures. In discussions of fine-tuning, it is *not* on the table that either side might have to deny, or become agnostic about, whether the relevant correlation has occurred. The proponent of fine-tuning does not suggest that the atheist must *either* give up her atheism, *or* become a fine-tuning skeptic. Rather, that our universe is fine-tuned for life is taken as common ground—the only question is *how theoretically or intellectually satisfying*, on the whole, competing accounts of the emergence of such a universe prove to be.

That is not to say that the fact that our universe is fine-tuned is a mere assumption. The theist supports this characterization by appeal to a series of findings in physics, and to basic constraints on the development of biological life, which together suggest that only a very narrow range of possible fundamental physical laws and basic parameters could accommodate life. For example, she appeals to evidence that, if the strong nuclear force, which binds protons and neutrons together, had been just a little bit weaker or stronger, life as we know it would not be possible.⁴⁵ Of course, *given* these results, if the universe *weren't* fine-tuned, the atheist wouldn't exist!⁴⁶ One might think this helps to explain why fine-tuning is

⁴⁴For example, those who defend atheism and realism often interpret their task as one of explaining why the relevant correlation, while still in some sense a coincidence, was *not as unlikely* as one might have thought. In response to fine-tuning, for instance, some defenders of atheism speculate that our universe was one member of an infinite multiverse, which was bound to produce at least one universe fine-tuned for life. In response to the epistemic challenge, defenders of realism have speculated that survival is generally good, so that it is not unlikely that creatures formed by evolutionary pressures (if they develop the capacity to reason at all) track the good, or that the features that make a creature capable of judging itself of moral worth also makes it morally valuable, so that it is not surprising that evolved creatures if they develop the disposition judge themselves of moral worth at all, are reliable in their judgments of moral worth.

⁴⁵Collins (2008) p. 85

⁴⁶Thanks to Jim Hutchinson for raising this point.

not called into question for the atheist in this debate, while the truth of her evaluative beliefs *is* called into question for the metaethical realist. If one becomes agnostic about whether slavery is wrong, or whether women have the same moral worth as men, one does not thereby call one's own existence into question.

But this line of thought does not take the analogy far enough. Agnosticism about whether the universe is fine-tuned *is* intelligible given the existence of the atheist and other complex life forms. For instance, it is possible that our current scientific consensus is deeply mistaken, and that life could exist under a much wider range of fundamental physical laws and basic parameters⁴⁷ than current thinking in physics and biology suggest. And why stop there? Perhaps the atheist is actually a disembodied mind, existing in a realm ungoverned by any physical laws. These are possibilities, although ones that we dismiss in regarding our current scientific methods as reliable. That the universe is fine-tuned remains common ground only if all parties in this debate may assume the reliability of these methods.

In the face of the fine-tuning argument, why make this assumption? Presumably because we think it more reasonable, on the whole, to call into question *either* atheism *or* the abhorrence of positing coincidences, *before* entertaining doubts about the reliability of our current scientific methods. These methods are, it seems, the best ways we have of learning about the physical universe. We do not consider our disinclination to posit coincidences a *better* guide to the physical truth than these methods.

By the same token, in considering whether the truth of the realist's evaluative beliefs can possibly come into question as a result of the challenge, each of us must ask: what do I consider a better guide to the evaluative truth? Do I consider ordinary approaches to settling evaluative questions—molded as these are by personal history, cultural practice, and our evolutionary development—a better guide? Or do I consider the general disinclination to posit coincidences between fact and belief a better guide?

⁴⁷Perhaps actually 95% of possible combinations of physical laws, basic parameters, and starting conditions would accommodate life. The fact that our universe accommodates life might not then be a coincidence that cries out for explanation.

In the spirit of G.E. Moore, we might approach this question by considering the provenance of the evaluative judgments we regard as most certain.⁴⁸ Take, for instance, the *prima facie* badness of pain. As things stand, there is simply *nothing surer* in evaluative life than that pain is *prima facie* bad. This is true, for all we can tell, and if we have to suppose ourselves evaluative prodigies in order to maintain the truth of our opinion on this point, we ought to do it. Our ordinary methods of addressing such questions are what lead many people to judge that pain is bad, that friendship is valuable, and that suffering and deprivation ought to be alleviated. When it comes to making judgments about the evaluative truth, we ought to trust those methods more than the general disinclination to posit coincidences.

The realist is not in a position to simply dismiss the disinclination to posit coincidences, even on the assumption that her realism precludes explaining our reliability. In that case, she still faces a choice: she could *either* suppose herself an (unlikely) evaluative prodigy, *or* give up realism. If she maintains her realism, she must justify doing so despite the cost (if any) of supposing herself an (unlikely) evaluative prodigy. But she deserves fair play. By putting the truth of the realist's evaluative beliefs on the table, the epistemic challenger raises the stakes in an illegitimate way. When faced with a dilemma between believing oneself inexplicably reliable about value and giving up one's realism, one might not immediately know what to think. How important, really, is it to explain *how* we get the truth about value? Is being unable, in principle, to give this explanation *worse* than denying that values are relevantly independent of us? It is far from clear. On the other hand, if our options are abandoning such realistic interpretations of value, or holding that we ought to remain agnostic about whether pain is bad, we know what to think.⁴⁹

I deny that the realist confronts the second set of options, if she cannot

⁴⁸(1939)

⁴⁹I consider this to be in the spirit of Ronald Dworkin's reply to the epistemologist who suggests that lacking an explanation of the workings of our faculty of moral judgment threatens skepticism. It can appear that Dworkin illegitimately brackets realism, pitting the demand for explanation against the demand for belief in moral claims. But he is right: the confusion is in the framing of the challenge, not in his reply. However, as the next chapter shows, there are other prospects for motivating the epistemic challenge in terms of explanatory value. (1996) pp. 117-118.

answer the epistemic challenge. This is not to deny that the realist ought to take seriously the possibility of radical evaluative error. On the contrary, she must, by her own lights. The realist holds that, fundamentally, evaluative truth is not a function of our evaluative attitudes or language. Thus, unlike some of her opponents, she must admit that, however sure a guide to the evaluative truth our current attitudes or linguistic conventions might appear, they could always be overturned, revealed as illusory. She ought to admit this, *even of our attitudes to the badness of pain*.⁵⁰

However, that does not mean that the realist ought to suspend her evaluative beliefs in the face of just *any* possibility of error. In particular, she need not do so in order to avoid accepting inexplicable reliability. Like the atheist confronted with the fine-tuning argument, when confronted with the epistemic challenge the realist may continue to trust what appear to her the best ways of getting at the truth about this subject. Doing otherwise would be to treat the dispositions that warn us against positing coincidences as surer guides to the evaluative truth than the dispositions that guide us in our judgment of the *prima facie* badness of pain. The realist need not take this position. Indeed, she ought to reject it. There is, as things stand, no surer guide to the evaluative truth than our attitudes about the badness of pain.⁵¹

Even if we cannot answer the challenge, it does not follow that we ought to deny our evaluative reliability. We have views about what is so in evaluative matters: about what has value and what gives us reasons to act. Fragmentary and provisional as these views may be, the disinclination to posit coincidences does not seem to offer a better guide to the evaluative

⁵⁰Some self-described realists may balk at this. I think they should not. The goodness of pain may be unimaginable for contemporary human beings, but for the realist, the limits of our imagination are, like our attitudes and language, not the final word on the matter.

⁵¹Could an inability to answer the epistemic challenge suggest the falsity of a large number of *less sure* evaluative beliefs, even if not the belief that pain is bad? Perhaps not all of the evaluative claims we accept will seem as sure as the badness of pain, when matched against the unacceptability of massive coincidences. This requires more thought, but initially, the prospects are murky. The challenge concerns our evaluative dispositions as a whole. One might think that whatever *general disposition* gets me to the conclusion that pain is bad is going to be more reliable than the one that disinclines me to accept coincidences.

truth. That maintaining these views commits us to a coincidence does not therefore give us a good reason to abandon them. Given this, failing to answer the challenge does not show that we ought to give up our evaluative beliefs.

2.5 Conclusion

We have not seen a reason, associated with the epistemic status of our evaluative beliefs, for preferring to have an answer to the epistemic challenge. The inability to answer the challenge does not seem the sort of thing that would make our evaluative reliability less admirable. Given the psychological importance and independent authority of our evaluative beliefs, the challenge puts no serious psychological or normative pressure on us to give up all of our evaluative beliefs.

However, there may still be a cost to regarding the correlation between our evaluative psychology and the evaluative truth as a coincidence. As theorists of value, we might remain unsatisfied—why shouldn't we demand an explanation, for its own sake? We will take up this possibility in the next chapter.

Chapter 3 The Epistemic Challenge and the Value of Explanation

3.1 The demand for explanation

What else of cognitive value might we miss by failing to answer the challenge? So far, we have considered valuable features of our evaluative beliefs that seemed to depend on their non-accidental truth: their admirability, their security, and their reasonableness as positions on what is true. We did not find cause for concern. It seemed that our evaluative beliefs would retain these valuable features, *even if* our reliability is an accident in the sense at issue in the challenge. We still lack a reason to address the epistemic challenge. But perhaps looking for something *else* that is threatened is a mistake: perhaps the good thing lost is simply *the explanation itself*.

Some who address the epistemic challenge appear to urge the importance of explanation for its own sake. They suggest that the correlation involved in our evaluative reliability “begs for an explanation.”¹ Enoch, for example, describes how “striking” it is for a high proportion of one’s beliefs about normative matters to be true, and for a high proportion of the truths about normative matters to be among one’s beliefs. This “striking correlation calls for explanation.” While “unexplained and even unexplainable correlations are not...impossible,” a metaethical theory that implies such an unexplainable correlation “loses plausibility points, perhaps to

¹Street (2006) p. 125

the point of unacceptability.”²

For many philosophers, this would be enough to motivate answering the challenge: if something calls for an explanation, presumably one ought to provide one. However, Enoch goes on to urge further consequences for our ordinary evaluative beliefs. He claims that “absent some such explanation, the correlation would be just too miraculous to believe.”³ If we cannot believe in our evaluative reliability, this undermines our ordinary normative beliefs: after denying that the correlation holds, one “can no longer hold these normative beliefs justifiably (and so, perhaps nor can [one] hold them as beliefs at all).”⁴ In chapter two, I argued that the failure to explain the challenge would not have these further consequences for our evaluative beliefs. *Even if* our evaluative reliability admits of no explanation, we are not in danger of losing these beliefs, nor of being unreasonable in regarding them as true. By the same token, inexplicability would not give us adequate reason to deny our evaluative reliability. If that is right, Enoch is wrong to claim that inexplicability has these additional consequences. But why can’t the demand to explain a striking correlation stand on its own as a motivation for answering the challenge? Even taking our evaluative reliability for granted, we might still have reason to want an explanation of that reliability.⁵

We might make sense of this demand “standing on its own” in two different ways. The first way begins from the realization that, in some circumstances, it strikes us as overwhelmingly likely what we observe has an explanation. (In such cases, we are tempted to exclaim that there *must* be some explanation.) For example, suppose that I am walking through the woods, and come across an arrangement of rocks in the shape of the sentence, “Hello Kirsten!” I will assume that there is some explanation; I will rightly dismiss the suggestion that the rocks fell into this shape by chance. Likewise, perhaps it is obvious that some *correlations* we observe

²(2009), pp. 421, 427

³Ibid. p. 421

⁴Ibid. p. 427

⁵By a “striking correlation” Enoch might just mean a correlation that, if inexplicable, would be “too miraculous to believe.” In this case, I want to urge that there is more to the demand for explanation than Enoch suggests: we may face pressure to explain a correlation, quite apart from whether it would otherwise be “too miraculous to believe.”

are highly likely to have an explanation, so that any theory that implies that they are mere coincidences—that is, that they are inexplicable—is committed to a likely falsehood. The correlation that constitutes our evaluative reliability might be one of these. If so, and if realism rules out any answer to the epistemic challenge, realism is committed to a likely falsehood.⁶

The second way starts from a readiness to hope for what is good for us, rather than a deference to what appears likely to us. We have a general intellectual interest in explaining the world. If a theory implies that a correlation is inexplicable, doesn't it imply that we cannot satisfy that interest, in at least that case? If so, and if realism rules out any answer to the epistemic challenge, then realism implies that we miss out on something good. This would not show that the realist's theory is *false*. However, it would make sense of why the realist would find the inability to answer the challenge disappointing, and so motivate her to show that realists can answer it.

As I argued in chapter one, metaethicists who address the epistemic challenge seem motivated to do so by the assumption that *something* of great cognitive value is at stake in providing an answer. As we also noted in chapter one, something would certainly be at stake, if the realist's inability to answer the challenge committed her to something false: to be committed to something false is a bad thing, cognitively speaking. Our first way of reading the explanatory demand appeals to our desire to avoid such a bad thing. Our second way appeals to our desire to achieve a good: the value of explanation itself. Once again, since our question is why we ought to care to answer the challenge at all, the focus on cognitive value is appropriate.

⁶I treat the fact that one is evaluatively reliable like the fact that the rocks are arranged in the shape of a sentence: one accepts that this is the case on the basis of one's ordinary ways of confirming facts like this, and asks whether this is likely to be a coincidence. If this is unlikely enough, this motivates rejecting assumptions that rule out an explanation (as metaethical realism might). Thanks to Niko Kolodny for pressing me to make this clear.

3.2 Explaining correlations

Before considering reasons to hope for an explanation of the correlation involved in our evaluative reliability, we need a clearer picture of what explaining a correlation involves. On one way of thinking, the correlation at issue holds because certain evaluative claims have two properties: they are true and they are believed by us. Because of this there is a correlation between the fact that certain evaluative claims are true and the fact that those evaluative claims are believed by us. On the one hand, it is *true* that pain is prima facie bad, that friendship is good, that persons have rights, and so on. On the other hand, we *believe* that pain is prima facie bad, that friendship is good, that persons have rights, and so on. What is it to explain something like this?

In some sense of explanation, there will be an explanation of a correlation whenever there are explanations of its correlata. Suppose the day of my graduation is the only rainy day in May. In that case, one day has two properties: being the day of my graduation and being the only rainy day in May. Because this single day has two properties, there is a correlation between these two facts: that this is the day of my graduation and that this is the only rainy day in May. (I will generally construe correlata as facts. While I sometimes find it convenient to speak of correlations between properties, events, or things, it should be straightforward to translate this into talk of facts.) There may be an explanation of why *that* day was the day of my graduation—in terms of my completion of the coursework in that semester, the thought process of the registrar who set the graduation date, and so on—as well as an explanation of why *that* day was the only rainy day in May—in terms of local weather conditions, rates of precipitation and evaporation in the area, and so on. But the correlata themselves afford us an explanation of the correlation, in a sense, since the correlation is nothing more than the correlata holding. The correlation holds *because* the correlata hold.

This is not what the epistemic challenger has in mind by an explanation of the correlation, however. She would not suggest that there is any difficulty giving *this* sort of explanation. In Street's words:

[The normative realist] may explain each *side* of the coincidence in as much depth as [she] likes—going into wonderful normative depth about why family and friendship are valuable, and wonderful scientific depth about why we were selected to think this. But all this goes nowhere toward explaining the thing that really needs to be explained, namely the *coincidence itself*.⁷

What does Street mean by an explanation of “the coincidence itself”? I suggest that she means an explanation of *the correlation* between the evaluative truths and our evaluative beliefs that *makes* the correlation no coincidence. The apparent difficulty for the realist is that her view only allows for what Elliot Sober has called “a coincidence explanation” of the correlation—that is, an explanation on which the correlata have no explanatory connection to each other: “neither causes the other, and they do not have a common cause.”⁸ When Street talks of explaining the coincidence, and Enoch talks of explaining the correlation, it seems they have in mind an explanation of the correlation that posits an explanatory relationship *between the correlata*.⁹ From this point on, unless I indicate otherwise, this is what I will mean by an “explanation of a correlation.”

Such an explanation might posit a variety of explanatory relations between the correlata: perhaps one explains the other, or perhaps some third fact explains both. For example, the correlation between smoke and fire is explained if the fire explains the smoke. The correlation of my arrival at Moses Hall at 4 pm and your arrival at 4 pm is explained if we both

⁷(2016) p. 322

⁸Sober is talking about the correlation and explanation of *events*, and so appeals to the notion of “cause.” While it can seem inapt to talk of facts causing other facts, we can adapt the notion of a “coincidence explanation” to one fact explaining another fact. (2012) p. 362-363.

⁹Enoch lays out the three directions of explanation here: “On a (robustly) realist view of normativity, it can’t be that our normative judgments are causally or constitutively responsible for the normative truths...[and] normative truths are causally inert, [so] they are not causally responsible for our normative beliefs. Nor does there seem to be some third-factor explanation available to the realist.” (2009) pp. 421-422. This also seems to be what a variety of other epistemic challengers have in mind by explaining the correlation, although for some (such as Setiya) this is only a necessary condition. See Setiya (2012).

read the colloquium announcement, and that explains each of us arriving when we do. However, there being an explanation of each of the correlata does not suffice for an explanation of the correlation, in this sense. For example, we may arrive at Moses Hall at the same time, but one thing explains why I arrive at that time, and something entirely different explains why you arrive at that time. No matter how extensive the explanations of the correlated facts, absent an explanatory connection between them, the *correlation itself* is not explained, in the special sense at issue here.

A correlation is no coincidence—or, equivalently, no accident—if and only if it has an explanation, in this sense. Given this understanding of the explanation of the correlation between our evaluative beliefs and the evaluative truths, our question is: do we have reason to hope that our metaethical view can allow for one?

3.3 There's surely some explanation!

One might think that *some* correlations simply must have an explanation or, more modestly, are very likely to have an explanation. In that case, any theory that implies that they lack an explanation is likely false. Perhaps the correlation between our evaluative psychology and the evaluative truths is one such correlation.

In Hartry Field's statement of a parallel epistemic challenge posed to the realist's counterpart in the philosophy of mathematics, one can detect something like this suggestion. Field doubts that, on a realist understanding of mathematics, there could be an explanation of the correlation between our mathematical judgments and the mathematical facts. Given the realist's view of mathematical entities, it seems "that we cannot explain the mathematicians' beliefs and utterances on the basis of the mathematical facts," nor can we explain the mathematical facts "on the basis of those beliefs and utterances," nor can we identify "some common cause producing both."¹⁰

If mathematical realism does not allow for such an explanation, Field

¹⁰Field (1989) p. 231

thinks, the realist has two options. One is to deny our reliability. He urges that this is unsatisfactory, for reasons we discussed in the prior chapter.¹¹ Since I have argued that we ought to maintain belief in our evaluative reliability in the face of the epistemic challenge, we can set this option aside in our context.

The realist's *other* option is to accept our reliability, but to claim that it is "a brute fact that needs no explanation."¹² But this is "dubious," according to Field:

[There] is nothing wrong with supposing that some facts about mathematical entities are just brute facts, but to accept that facts about the relation between mathematical entities and human beings are brute and inexplicable is another matter entirely.¹³

Explanation is demanded, when it comes to our mathematical reliability, because of the *kind* of fact at issue. When it comes to facts about the relation between mathematical entities and human beings, we should doubt that they are brute.¹⁴

But why so? Perhaps by a "brute" fact, Field has in mind a fact about a fundamental feature of reality: a fact central to the correct account of the world, that explains many other facts about the world without itself admitting of explanation. On such a picture, the correlation that constitutes

¹¹That is, to deny our reliability would take us down a path that threatens the stability or reasonableness of our mathematical beliefs. After all, "to maintain a class of beliefs while holding the meta-belief that most of those beliefs are false seems plainly unsatisfactory." Ibid. p. 232.

¹²Ibid. pp. 231-232

¹³Ibid. p. 232

¹⁴John Burgess and Gideon Rosen have suggested that such a doubt could only threaten our justification for beliefs about mathematical entities: "Though Field maintains that his challenge is 'not to our ability to justify our mathematical beliefs'...the implicit suggestion [...] can hardly be anything but this, that if the reliability thesis cannot be explained, then continued belief in claims about mathematical entities is unjustified." (1997) p. 42. I suggest a different reading: Field is just saying that *inexplicable* mathematical reliability is dubious.

our reliability would be something like a fundamental law of physics. This is indeed a controversial position, as Thomas Nagel observes: “The thought that the relation between mind and the world is something fundamental makes many people in this day and age nervous.”¹⁵ In fact, this is a position that is hard to defend. We view some features as fundamental, because facts about them appear to explain so much else that we observe. This is one mark of fundamentality. When one considers the range of phenomena that such a fundamental relation between mind and world would explain, they appear much more limited than those explained by the fundamental laws of physics. While the laws of physics seem to have an explanatory relation to the features of brute matter, to the behavior of animals, to the behavior of human beings, and even to events in consciousness, our evaluative reliability would presumably explain only the epistemic credentials of our many particular evaluative beliefs. Perhaps Field is thinking that since our evaluative reliability could not be such a fundamental feature, we are forced to conclude that evaluative reliability has an explanation.

Yet taking the correlation involved in our evaluative reliability to be fundamental is not our only option, if we deny that it has an explanation, in the sense at issue. We encounter such unexplained correlations about the most ordinary, nonfundamental matters. As Street observes, there are many correlations that are “mere coincidences—coincidences regarding which no further explanation is required, and indeed regarding which we think it would be positively confused to insist upon one.”¹⁶ One is unlikely to find an explanation of the correlation between being the day of one’s graduation and being the only rainy day in May; between Angela Bassett’s and Madonna’s birthdays; or being an overworked front-line healthcare worker and being the winner of the Surrey lottery. Unexplained correlations like these are perfectly ordinary, and we do not suspect that they involve fundamental features of the world.

Is there some *other* reason to think that the correlation involved in our evaluative reliability *must* have an explanation or, more modestly, is very likely to have an explanation? We sometimes feel entitled to assume that some correlation has an explanation. Suppose that I am reading philoso-

¹⁵(1997) p. 130

¹⁶(2016) p. 311

phy papers by different authors, and I find that their lengthy statements of a certain argument are phrased in precisely the same way. I will assume that there is some explanation for the fact that author A put the argument in *this* way, and so did author B; that is, I will assume that their identical phrasing is no coincidence. Perhaps one borrowed the statement from the other, or there is a third person who is the ultimate source of both statements, or they created the statement together, and so on. In such a case, I would be right to dismiss the suggestion that the authors independently hit on the same way of stating the argument. After all, there is a nearly endless variety of ways that people can put a philosophical argument into words. (To be clear, this is a variety of possible *formulations* of a given argument, not a variety of genuinely different arguments for some conclusion. It is not particularly surprising, given that two people are thinking about a similar question in philosophy, that they make the same point.) Given all the equally eligible ways of putting an argument, it is highly unlikely that two people would put an argument in exactly the same way by coincidence.

Why is dismissing this possibility a reasonable reaction? Only, I suggest, because I have relevant background knowledge of the array of eligible possibilities one faces in framing a long argument, and so of the low probability that such a correlation would come about by chance.¹⁷ One can see the relevance of background knowledge by considering what my reaction ought to be to the suggestion that a similar correlation in a domain unfamiliar to me is a coincidence. For example, I am unfamiliar with the process of stating a novel proof of a theorem in set theory. If two authors produced identical statements of a previously unknown proof, I would not be entitled to reject the hypothesis of mere coincidence. For all I know, there could be just a few eligible ways to state any given proof, given the language and formal constraints of set theory. In that case, given the small number of ways they could have formulated their proofs, that they are the same by coincidence is not especially unlikely. I cannot rule this out, given my limited understanding of what is involved in stating a proof in set theory; it would therefore be unreasonable of me to insist that there *must* be some explanation of the correlation, that this *cannot* be a coincidence.

¹⁷This is a necessary, not a sufficient condition, for being reasonable in dismissing the suggestion. If the only available explanations of the correlation are extremely far-fetched, for example, then I may have to reassess.

For the epistemic challenger to have this sort of justification for dismissing the possibility of a coincidence between our evaluative attitudes and the evaluative truth, she too must have an understanding of the relevant range of possibilities. As the spread of conspiracy theories in the wake of the 2020 U.S. presidential election reminds us, confidence that there *must* be an explanation for some correlation is not a reliable sign that one has the relevant background knowledge. People are often confident that some correlation *must* have an explanation, despite having very little idea of the possibilities in play. For example, suppose that three counties experienced a major error in their vote count in this election, and that all of these counties used Dominion voting machines. It would be very easy to find an amateur detective who, relying only on this information, would insist that the correlation between having major voting errors and using Dominion voting machines *must* have an explanation. Perhaps the machines' common software was hacked to produce the errors, or (a more innocent explanation) perhaps the design of Dominion machines is defective or unintuitive for users. It can even appear naive to take seriously the hypothesis of coincidence. However, further information might reveal that a coincidence is not particularly unlikely. Perhaps half of all counties in the country use Dominion voting machines, for instance.¹⁸ In that case, if using a Dominion voting machine has no explanatory relationship to having a major voting error, we should expect that, when such errors occur, about half the time they will occur in counties using Dominion machines. Given just this information, the probability of all three major errors happening in these counties is $1/2 * 1/2 * 1/2 = 1/8$. This is not particularly unlikely—it is equivalent to a fair coin coming up heads three times in a row.¹⁹ In this case, and in many others, we ought to gather more information before making any judgment about the probabilities.

We are broadly evaluatively reliable; the relevant correlation holds. Does

¹⁸Suppose this is for an innocuous reason, such as that the U.S. market for voting machines only supports a few manufacturers.

¹⁹The broader moral is that *coincidences* are not always unlikely. Note also that there being a small number of possible outcomes does not *explain the correlation*, in our sense. Establishing that a conjunction of facts is not unlikely is not the same as establishing that their conjunction is no coincidence. Take the case of flipping two coins, and both coming up heads. There is no explanatory connection between the fact that one coin came up heads and the fact that the other coin came up heads; the outcomes are independent. Still, both coins will come up heads one out of four times.

our background knowledge suggest that this was unlikely to have come about by coincidence? In particular, does an epistemic challenger have sufficient knowledge of the background possibilities to insist on an explanation? Or is she like our amateur detective, in moving too quickly from correlation to causation (in an extended sense of causation)?

I submit that no one has background knowledge of the relevant possibilities adequate to insist that there must be an explanation here. Remember that the correlation to be explained is between the evaluative claims we believe and the evaluative claims that are true. How likely was this correlation to come about by coincidence? It is hard to know even how to get started in answering this question. However, the literature on the epistemic challenge suggests one starting point. The evaluative claims we now believe are the result, at least in part, of our existing evaluative dispositions. These in turn are the result of various causal forces, including the selective pressures that produced the evaluative psychology we inherited from our ancestors.²⁰ We can think of the evaluative dispositions we inherited as dispositions to make certain sorts of evaluative judgments, rather than others. So we might make progress on our question by asking whether it was highly unlikely that we would have broadly reliable evaluative dispositions, given the possibilities for the evaluative facts and evaluative dispositions in play.

What are the possible evaluative facts? Many regard them as *metaphysically necessary*: if so, there is only one way they could have been.²¹ For now, let us suppose that there is just one possible set of evaluative facts. What, on the other hand, are the possible evaluative dispositions? Perhaps these consist of *evolutionarily* possible dispositions for creatures like us: possible evolutionary histories of such creatures, in which they develop the disposition to make evaluative judgments at all. If these are

²⁰Many prominent epistemic challengers stress the relevance of evolutionary psychology to the explanation of these evaluative dispositions. See especially Street (2006), Enoch (2009), Wielenberg (2010), Joyce (2006), Copp (2008), and Shafer-Landau (2012).

²¹T.M. Scanlon regards “pure normative facts” as necessary: such facts do not vary with non-normative claims nor do they “vary ‘on their own.’” (2014) p. 41. In giving an answer to the epistemic challenge, Enoch depends, at one point, on the claim that “[f]undamental normative truths are presumably necessary in a fairly strong sense.” (2009) p. 433.

the possibilities, then for all we know, this coincidence was not highly unlikely. Perhaps in the vast majority of possible evolutionary histories, creatures like us develop evaluative dispositions that are broadly reliable. Holding fixed that they develop the capacity for evaluative judgment at all, perhaps creatures like us would develop broadly reliable evaluative dispositions 75% of the time, and broadly unreliable evaluative dispositions 25% of the time.

This is mere conjecture, of course. The point is that we simply do not know the various ways that evolutionary processes *could* operate on creatures like us under varying circumstances, while also leading us to develop the capacity to make evaluative judgments. One might delineate such possibilities by understanding *one* actual instance of such a process, and understanding this *well enough* to anticipate how interventions on that process might affect the reliability of a creature's evaluative dispositions, without preventing their emergence altogether. Scholars and scientists have offered hypotheses about the emergence of some human capacities for evaluative judgment—capacities for moral judgment, for example.²² Yet controversies continue even in how to give a satisfactory description of what would count as having a moral capacity.²³ Given that even what falls within the relevant range of capacities remains to be determined, we do not understand the evaluative development of creatures like us sufficiently to determine the range of possibilities. And so, if the relevant possibilities are *metaphysical* in the case of the evaluative truths, and *evolutionary* in the case of our evaluative dispositions, we do not yet have background knowledge adequate to support the insistence that our evaluative reliability could have been no coincidence.

Yet are these the relevant ranges of possibilities? Some epistemic challengers have made other suggestions. Street, for one, has treated *logically* or *conceptually* possible evaluative dispositions and evaluative truths as

²²See, for example, De Waal (2006); Korsgaard (2006); Tomasello (2015); Buchanan (2020).

²³On some accounts, such as that offered by Frans de Waal, moral capacities centrally involve the ability and inclination to cooperate with others and respond to their needs. Recent writers have followed Christine Korsgaard, in taking a creature's guidance of itself by norms and rules as essential to moral capacities, with pro-social inclinations playing a peripheral, or co-equal role. See De Waal (2006) pp. 52-55; Korsgaard (2006), pp. 105-112.

relevant. In arguing that our being evaluatively reliable by coincidence is “extremely unlikely,” Street appeals to the “huge universe of logically possible evaluative judgments and truths.”²⁴ She urges that

the universe of logically possible evaluative judgments is huge, and we must think of all the possible evaluative judgments that we *don't* see – from the judgement that infanticide is laudable, to the judgement that plants are more valuable than human beings, to the judgement that the fact that something is purple is a reason to scream at it.²⁵

If these are relevant possibilities, it seems that our evaluative judgments and the evaluative truths *could* have been almost *anything*. In particular, whether the disposition to make a given evaluative judgment was an evolutionary possibility for creatures like us is irrelevant to whether that judgment is a relevant possibility. Given Street’s view of the relevant possibilities, it would indeed be astounding if our *actual* evaluative judgments aligned with the *actual* evaluative truth by coincidence. The unlikelihood of such a coincidence might well support the demand for an explanation.

But why should we treat *these* as the relevant possibilities, as opposed to the metaphysically possible evaluative truths and the evolutionarily possible dispositions to evaluative judgment? While Street does not address this question explicitly, at certain points she suggests that the *realist*, in particular, is committed to the taking seriously all logically possible evaluative judgments in determining the likelihood of evaluative reliability by coincidence.²⁶ She suggests that this follows from the realist’s commitment to the radical independence of the evaluative truths from our evaluative judgments. Because of this commitment, the realist cannot rely on her *actual* judgments, and those of others, to justify dismissing bizarre logical possibilities as she considers the likelihood of becoming reliable by coincidence. The realist can only dismiss these possibilities if she has

²⁴(2006) p. 122

²⁵(2006) p. 133

²⁶See Street (2016) pp. 313, 315.

some “non-trivially-question-begging reason” to think that *her* evaluative judgments, and not these others, are among the true ones.²⁷

I am unsure what to make of this defense of the relevance of logical or conceptual possibilities, but it seems insufficient.²⁸ First, suppose the realist must admit that, in *some* sense of possibility, there are countless possible evaluative judgments and evaluative truths. Even so, this does not imply that these possibilities are relevant to determining the likelihood of being evaluatively reliable by coincidence. Second, Street invokes the realist’s view that the evaluative truth is independent of our actual judgments to say why, for her, the logical possibilities are relevant. But then why not bring the realist’s other views to bear? Why can’t the realist appeal to her metaphysical view that the evaluative truth is necessary, or to her scientific view that, for creatures like us, only evolved evaluative dispositions and their resulting evaluative judgments are genuine possibilities, and so narrow the space of relevant possibilities? It is entirely unclear what one can and cannot appeal to in sketching the space of possibilities. Even at this meta level, I suggest, we lack the background knowledge necessary to support the insistence on an explanation. We have a way of determining the logical possibilities for evaluative judgment and evaluative truth—through the sort of imaginative exercise Street offers. So, if these were the relevant possibilities, that would solve our earlier problem in delineating evolutionary possibilities. However, we would have to know that the logical possibilities *are* the relevant possibilities, and we still lack this knowledge.

In summary: we saw that, in ordinary cases where one is entitled to insist that a correlation is highly likely to have an explanation, one’s background knowledge supports that insistence. However, it seems that no one has equivalent background knowledge when it comes to the correlation constituted by our evaluative reliability.

²⁷Ibid. p. 315

²⁸Throughout her work on the epistemic challenge, Street contends that the realist ought to suspend belief in her own evaluative reliability in answering the challenge. This may explain why she claims that the realist cannot appeal to any of the evaluative claims she believes in determining the relevant possibilities. But we have rebutted the contention that the epistemic challenge demands this suspension of belief.

Suppose that metaphysically possible evaluative truths and evolutionarily possible evaluative dispositions are the relevant possibilities in determining whether the correlation was unlikely to occur by coincidence. On this assumption, I have argued that we need a better grip on the ways these evaluative dispositions might have emerged, before we can conclude that this coincidence would be highly unlikely. (For all we know, once we have this background knowledge, our coincidental reliability will not turn out not to be particularly unlikely.) Epistemic challengers might disagree on two counts. First, they might argue that we have a more secure grasp on how our evaluative capacities evolved than I suggest. In this case, they would need to produce that information, and put it to use in delineating the relevant possibilities. Second, and more fundamentally, they might deny that *these* are the relevant possibilities for determining the likelihood of this coincidence: they might urge, with Street, that it is rather the logical or conceptual possibilities that matter. However, I have urged that none of us know how to make this choice, either. So it seems that a lot of work remains in order to pursue this strategy. Without the relevant background knowledge, challengers cannot insist on an explanation for the sort of reason we can do so in ordinary cases.

I know of no other way to insist that the correlation involved in our evaluative reliability is highly likely to have an explanation. Perhaps the challenger could provide one, however. If so, it is their burden to articulate that entitlement.

I will now move on to an entirely different way of framing the demand for an explanation for its own sake: not in terms of what is needed to avoid commitment to what is likely to be false, but rather, in terms of what is needed if we are to hope for a cognitive good.

3.4 The value of explanation

Since it doesn't seem that we can dismiss the possibility of coincidence outright, on the grounds of its implausibility, should we nonetheless *hope* for an explanation of the correlation between our evaluative psychology and the evaluative facts, because of the positive value of such an expla-

nation? If so, then we once again have a reason to answer the challenge: something of importance hangs on explaining this correlation.

What, in general, is the value of explaining correlations? Suppose that I have observed, over a period of time in the past, that my joints ache on hot days. I might want to discover an explanation of the correlation: that is, an explanatory connection between the fact that on certain days my joints ache and the fact that those days are hot. Why might I want this? On the one hand, the explanation might be a source of power over these events. If I have an explanation of the correlation between hot days and my joint pain, this might give me the ability to *predict*, and even to *control*, the aching of my joints. If hot days cause my joint pain, for instance, this may enable me to predict that the correlation will continue into the future, as well as to intervene on my joint pain (by, for instance, moving to a colder climate). On some accounts the value of explanation consists in such abilities to predict and control.

When it comes to answering the epistemic challenge, however, it is hard to see how explaining the correlation of our evaluative judgments and the evaluative truth could confer such abilities. For one thing, one predicts and seeks to control events that have not yet occurred. Coming to have our current evaluative psychology is not such an event. Regardless of whether our reliability is a coincidence, certain causal forces have *in fact* brought human evaluative psychology to this point. Nothing, it seems, remains to predict or to control.²⁹ In any case, no one writing in this literature mentions either sort of power in addressing the epistemic challenge. Since I doubt that this is what they have in mind, I will not pursue this line of thought further.

Explanations are also thought to have *intellectual* value, however, quite apart from whether they give us any power over the phenomena explained. Even if I have no hope of intervening on my joint pain, explaining *why* my

²⁹Suppose that what is explained is not (merely) our evaluative psychology, but rather the evaluative truth. Perhaps our psychology explains the truth, or perhaps both are explained by a common cause. Would we gain power over evaluative truth by grasping that explanation? As in the explanation of our evaluative psychology discussed in the main text, this presumably depends on the explanans, which might or might not be manipulable by us. It would not be, if the evaluative truths were necessary.

joints flare up on hot days may give me a sort of intellectual satisfaction: a sign of the value of such explanations. Perhaps we miss out on something of intellectual value, if this correlation has no explanation.

This line of thought inspires a very simple argument that we do miss out on something of intellectual value if we cannot answer the epistemic challenge. If there is no answer, then the correlation between our evaluative judgments and the evaluative truth has no explanation. But explanation is of intellectual value. So, if anything—including this correlation—has no explanation, we miss out on something of intellectual value. We can conclude that, if there is no answer to the epistemic challenge, we miss out on something of intellectual value. Given this, the realist, and other metaethicists, have a reason to hope for an answer to the challenge.³⁰ This simple argument appears commonsensical. On its own, it seems to motivate answering the challenge.

However, I will argue that it presupposes a peculiar and disputable way of conceiving of the value of explanation. In fact, this conception is not only disputable: it is *disputed*. As I will show, on at least one prominent theory of explanation, the argument does not go through. This is the *unificationist* theory of explanation, defended by Philip Kitcher.³¹ I will show that, on Kitcher's view of explanation, it is possible to lack an explanation of the correlation involved in our evaluative reliability, while lacking nothing of intellectual value. I will not argue in favor of this theory of explanation and its conception of explanatory value, nor will I reply to prominent objections. My goal is not to establish the unificationist theory of explanation. Rather, my goal is to better define the burden that the epistemic challenger faces, if she wishes to appeal to the intellectual value of explanation in motivating the challenge. I will show that, on one prominent, plausible view of explanation, we might lack an explanation of the relevant correlation while lacking nothing of value. Assuming that I succeed in showing this, the challenger will need to offer a definite view of the value of explanation—one that will make sense of hoping for an

³⁰A number of realists concede that, if their view implied that there is no explanation of the correlation involved in our evaluative reliability, this on its own would lessen the attractiveness of realism. They may have in mind the kind of loss currently at issue. See Enoch (2009) p. 435; Shafer-Landau (2003) p. 234.

³¹Kitcher (1989)

answer to the epistemic challenge—that takes one of two forms. They might offer an account of the intellectual value of explanation that competes with the unificationist view, or they might offer an argument for further claims that, when combined with the unificationist view, make sense of hoping for such an explanation. Until either attempt is made, challengers have not met their burden of identifying something of intellectual importance at stake in their challenge.

Unification and the value of explanation

According to Kitcher, explanations have value because they *unify* our picture of the world, allowing us to see “connections, common patterns, in what initially appeared to be different situations.”³² We unify our picture of the world—our beliefs about what is so—by seeing how *more* of what is so is explained in *fewer*, relatively *stringent* ways. In unifying our beliefs, we see that “there are patterns of derivation that can be applied again and again” to arrive at conclusions about what is so. The value of explanation is realized only by grasping *multiple* explanations and seeing how they fall under such common patterns.³³ Explanations have explanatory value only because they belong to a set of explanations that makes unification possible.³⁴

Unification is a matter of degree. Explanations are *most* valuable, then, when they belong to the set of explanations that *best* unifies our beliefs.³⁵ Since we ought to prefer *more* explanatory value to *less*, we should have no qualms about abandoning an explanation that could only belong to a less unifying set of explanations.

³²Ibid. p. 432

³³Ibid. p. 437.

³⁴Ibid. pp. 430, 437

³⁵Kitcher would not put this point in quite this way. For Kitcher, something is acceptable as an explanation only if it belongs to the explanatory store. Since our focus is on explanatory *value*, and this comes in degrees, it is helpful to use “explanation” more loosely. I will use “an explanation” to refer to what Kitcher would call a “candidate explanation” or “derivation” of that phenomenon, and the “most valuable explanation” to refer to what belongs in the explanatory store. Ibid. p. 431.

To get a clearer idea of when an explanation belongs to a set of explanations, and of when some set *best* unifies our beliefs, we should briefly review some of the details of Kitcher's theory of explanation. Kitcher follows many traditional theories of explanation in taking explanations to have the form of *deductive arguments*.³⁶ An explanation is an "argument whose conclusion describes" the fact to be explained. As in any argument, an explanation consists of steps classified either as premises or subconclusions, and a conclusion that follows from these in a particular way. Take, for example, this argument:

Alpha Centauri 1

1. All far stars are twinklers when viewed from Earth.
 2. Alpha Centauri is a far star.
- So,
3. Alpha Centauri is a twinkler when viewed from Earth. (1, 2)

This is an explanation of *why Alpha Centauri is a twinkler when observed from Earth*. Its conclusion describes what is to be explained. The argument has a determinate structure: steps 1 and 2 are premises that entail the conclusion through universal instantiation and conditional elimination.

These features, discernable in the argument in isolation, do not yet reveal whether the argument has explanatory value.³⁷ That depends on

³⁶C.G. Hempel's well known theory of explanation, for example, construes explanations as deductive arguments. These are the conditions Kitcher places on "ideal explanations." In giving explanations in daily life, Kitcher takes us to roughly approximate this sort of argumentative structure, once we make explicit various shared assumptions. He also allows that we may have to accept some irreducibly probabilistic theories, although he regards these as having less explanatory value. See *Ibid.* pp. 448-452.

³⁷This is not to deny that, whenever we grasp a deductively valid argument whose premises are true, this has *some* kind of intellectual value. Such arguments are a potential means of acquiring a true belief from other true beliefs. However, acquiring a true belief that *p* is not yet to acquire an explanation of *why p*. See, e.g. Aristotle, *Posterior Analytics*, Book I, ch 13, 78a22-78b10.

whether the argument instantiates a *pattern* that is more broadly applicable. As mentioned above, we achieve greater unification by revealing that *more* of what is so is explained in *fewer* ways. Take, for example, the fact that Alpha Centauri is a twinkler. In searching for a valuable explanation of this fact, it is relevant that other stars, planets low in the sky, city lights, steadily burning flames, and sheets of illuminated metal are also sometimes twinklers. Many sorts of objects twinkle in the right circumstances. Each will have its own explanation, in the form of an argument whose conclusion describes that particular phenomenon. However, these explanations might fall under the same *pattern*.

Let me elaborate on Kitcher's notion of pattern, by showing how two particular explanations can share a pattern, in this sense. Consider first another explanation of the twinkling of Alpha Centauri:

Alpha Centauri 2

- A1. If light traveling from a source to a point of observation is highly refracted and the source has an angular diameter of between 0.0005 and 0.05 arcseconds from the point of observation, that source is a twinkler when viewed from that point of observation.
- A2. Light traveling from Alpha Centauri to the Earth travels through Earth's atmosphere.
- A3. Light that travels through Earth's atmosphere is highly refracted.
So,
- A4. Light traveling from Alpha Centauri to Earth is highly refracted.
(A2, A3)
- A5. Alpha Centauri has an angular diameter of between 0.0005 and 0.05 arcseconds from Earth.
So,
- A6. Alpha Centauri is a twinkler when viewed from Earth. (A1, A4, A5)

This argument describes specific objects and properties: Alpha Centauri and its size, the Earth's atmosphere, the Earth. It derives the claim that Alpha Centauri is a twinkler from our position by appeal to that star's apparent size—its "angular diameter"—when viewed from Earth, and to what happens to light traveling from Alpha Centauri through Earth's atmosphere. In all these ways, the argument differs from the following explanation of the twinkling of city lights:

City Lights

- C1. If light traveling from a source to a point of observation is highly refracted and the source has an angular diameter of between 0.0005 and 0.05 arcseconds from the point of observation, that source is a twinkler when viewed from that point of observation.
- C2. Light traveling from a city light in San Francisco to the Oakland hills travels through a great deal of air pollution.
- C3. Light that travels through a great deal of air pollution is highly refracted.

So,

- C4. Light traveling from from a city light in San Francisco to the Oakland hills is highly refracted. (C2, C3)
- C5. A city light in San Francisco has an angular diameter of between 0.0005 and 0.05 arcseconds from the Oakland hills.

So,

- C6. A city light in San Francisco is a twinkler when viewed from the Oakland hills. (C1, C4, C5)

While C1 of this argument is the same as A1, it is a different argument from **Alpha Centauri 2**. It describes an entirely different intervening medium, point of observation, and source of light. Still, the arguments have a lot in common. We can see what they have in common by formulating what Kitcher calls a "schematic argument," as well as instructions

for how to fill out the schematic argument to produce any particular argument.

Schematic argument

S1. If light traveling from a source to a point of observation is highly refracted and the source has an angular diameter of between 0.0005 and 0.05 arcseconds from the point of observation, that source is a twinkler when viewed from that point of observation.

S2. Light traveling from S to O travels through M.

S3. Light that travels through M is highly refracted.

So,

S4. Light traveling from S to O is highly refracted. (S2, S3)

S5. S has an angular diameter of between 0.0005 and 0.05 arcseconds from O.

So,

S6. S is a twinkler when viewed from O. (S1, S4, S5)

Filling instructions

- For S, fill in a source of light.
- For M, fill in a medium.
- For O, fill in a point of observation.

We can now say with greater precision what is meant by a “pattern.” By a “pattern,” Kitcher has in mind a schematic argument, its filling instructions, and a “classification”: a specification of the “inferential characteristics” of the argument. These characteristics include which steps are premises and which are inferred, and what rules of inference are used.³⁸ In this case, the classification designates S1-S3 and S5 as premises; S4 as

³⁸Kitcher (1989) p. 432-433

inferred from S2 and S3; S6 as inferred from S1, S4, and S5; and certain rules of first order logic as the rules of inference.

We can see that both **Alpha Centauri 2** and **City Lights** fall under the argument pattern above. (Let us call this Pattern A.) We can formulate either argument by starting with Pattern A's schematic argument and following its filling instructions. The arguments also have the same classification: the steps in each that correspond to those in the schematic argument have the same status as premises or as inferred, and the rules used to draw the inferences are the same.³⁹ **Alpha Centauri 1** does not instantiate Pattern A—one could not formulate this argument by starting with the schematic argument and following the filling instructions. **Alpha Centauri 1** instantiates other argument patterns, of course, including one (call it Pattern B) obtained by replacing the name "Alpha Centauri" in steps 1 and 2 with a variable fillable with any object.⁴⁰

Pattern A makes many demands on its instantiations; in Kitcher's terms, it is relatively "stringent."⁴¹ To instantiate this pattern, an argument must describe a source of light, the source's size, a medium, and a point of observation; it must take the generalization in S1 as a premise; its steps must satisfy relatively specific inferential relations to each other, and so on. Intuitively, Pattern A makes *more* demands on its instantiations than does, for example, the argument pattern produced by replacing each step of Pattern A's schematic argument with propositional variables, fillable by any proposition.⁴² This different pattern would make fewer demands on its instantiations, and these explanations, as a class, would have less in common with each other. This has implications for the unifying power of that pattern and its set of instantiations: grasping how several explanations instantiate a less stringent pattern is less unifying, all else equal, because such a pattern captures a less substantial similarity.⁴³

³⁹Ibid. pp. 432-433

⁴⁰The schematic argument for Pattern B is: S*1. All far stars are twinklers. S*2. S is a far star. S*3. S is a twinkler. (S*1, S*2).

⁴¹Ibid. p. 433

⁴²E.g. by replacing the current contents of S1 with "If p and q , then r " and doing the same for all the other steps.

⁴³As Kitcher notes, the unifying power of an argument pattern involves balancing the demand for stringency and the demand to explain more types of phenomena. Too much stringency produces a "'pattern' which is its own unique instantiation." The set

How do we compare the value of two explanations of a phenomenon—for instance, that of **Alpha Centauri 1** and **Alpha Centauri 2**? This depends on comparing the argument patterns they instantiate, both in terms of stringency and in terms of the range of phenomena their instantiations explain. Greater explanatory value can be achieved by using fewer patterns that are more stringent, and whose instantiations explain more types of facts.⁴⁴ Using these criteria, it seems that **Alpha Centauri 2** has greater explanatory value than **Alpha Centauri 1**. The former instantiates Pattern A, while the latter instantiates Pattern B.⁴⁵ While these patterns are perhaps equally stringent, the set of explanations that instantiate Pattern A derive more types of facts than those that instantiate Pattern B: not only the twinkling of far stars, but the twinkling of any source of light observed through a medium that refracts light traveling through it. (Of course, it's possible that achieving even greater explanatory value requires abandoning *both* Pattern A and Pattern B in favor of some third pattern applicable to optical phenomena beyond twinkling.) There is no reason, on Kitcher's view, to employ more argument patterns without any corresponding increase in the types of facts explained. **Alpha Centauri 1** and **Alpha Centauri 2** are explanations of the same phenomenon, instantiating different argument patterns—if we were to adopt both patterns, this would not lead to more types of facts being explained. Since this is so, these particular explanations are in competition; we should abandon the one with less explanatory value.

This completes our discussion of explanatory value, as Kitcher conceives it. According to Kitcher, we can understand explanatory value in terms of grasping how more types of facts are explained in fewer, more stringent ways, where this can be understood using his notions of an argument pattern and its instantiations. When it comes to explanatory value, a set of explanations that unifies one's beliefs to a greater degree is better than any set of explanations that unifies them to a lesser degree. Even if the latter set contains an explanation of a particular fact that is *not* explained in the

of explanations that best unifies our beliefs will make the "best tradeoff" between the criteria of unification. Ibid. p. 433.

⁴⁴Ibid. p. 432.

⁴⁵Each also instantiates countless other argument patterns, but their explanatory value will just depend on the pattern they instantiate that does best by the "criteria of unification." Ibid. p. 434.

former, one does not thereby miss out on any explanatory value in opting for the former. The “worth” of explanations “cannot be appreciated by considering them one-by-one but only by seeing how they form part of a systematic picture” of the world.⁴⁶ To have good reason to think that we lose out on something of intellectual value by lacking an explanation of a particular fact, we need reason to think that this explanation belongs to the set that best unifies our beliefs.

Back to the epistemic challenge

Let us return to the epistemic challenge. Our question was whether anything of intellectual value is at stake in giving an answer to the challenge. In particular, is the intellectual value associated with explanation at stake? It seemed so. After all, if there is no answer to the epistemic challenge, we lack an explanation of something, namely, an explanation of the correlation between the evaluative truth and our evaluative psychology. At least with respect to that correlation, we miss out on something of value.

We can now see that this is not so straightforward. As we have just observed, on Kitcher’s view of the value of explanation, lacking some particular explanation need not imply the loss of anything of explanatory value. A particular explanation is only valuable as an occasion for unifying our beliefs. This only happens when we grasp that the explanation is substantially similar to other explanations. That is, a particular explanation becomes an occasion for unifying our beliefs if it instantiates a stringent argument pattern that many other explanations instantiate. If *this* is the way that explanation has intellectual value, we cannot infer that we lose out on this value if a particular fact has no explanation. So on this view, the simple argument fails.

Furthermore, on the unificationist view, explaining a correlation so as to make it no coincidence has no special explanatory value. What matters is not minimizing coincidences, as such, but rather that *more* is explained using *fewer* and *more stringent* explanatory strategies. In this vein, it may help to remember Sober’s notion of a “coincidence explanation” of a cor-

⁴⁶Ibid. p. 430

relation, on which its correlata have no explanatory relationship to each other. A coincidence explanation of the correlation between my joint pain and hot days, for example, might (1) derive that I have joint pain on certain days from claims about my diet on those days, and (2) derive that those days are hot from claims about local weather conditions, and (3) conjoin the conclusions of these arguments to derive the correlation. This is not an explanation of the correlation, in *our* sense—a sense that makes the correlation no coincidence—but it is such an explanation in *Kitcher's* sense.⁴⁷ More importantly, a coincidence explanation may be *best*: it may belong to the set of explanations that *best unifies* our beliefs. So the question left open for the epistemic challenger is: why think that, by drawing an explanatory connection between our evaluative beliefs and the evaluative truths, we will achieve greater unification than we would otherwise?

Perhaps the epistemic challenger can argue that, when it comes to our evaluative reliability, we have good reason to think that this sort of explanation *would* achieve greater unification. After all, the challenger might point out, our evaluative reliability—the rough alignment of our evaluative beliefs with the evaluative facts—implies *multiple* correlations of evaluative belief and evaluative truth. This raises the hope of achieving greater unification, through discovering *many* (non-coincidence) explanations of these correlations, that instantiate one relatively stringent argument pattern.

To illustrate this, let's make use of a candidate answer to the epistemic challenge: Street's explanation of the correlation involved in our evaluative reliability. This explanation relies on her anti-realist account of evaluative truth. Here is a reconstruction of this account:

Anti-realism: For any evaluative claim *p* believed by an agent A, it is true (for A) that *p* if and only if A's belief that *p* would stand up to scrutiny in terms of A's other evaluative beliefs.⁴⁸

⁴⁷In answering the challenge, then, it does not follow that we explain *more* than we would otherwise, from Kitcher's point of view. An explanation acceptable to the challenger and a coincidence explanation of the correlation are equally explanations of the correlation for him.

⁴⁸According to Street's anti-realist position, the only evaluative truths are those that

One can construct an explanation for a particular correlation between an agent's evaluative belief and an evaluative truth (for that agent) using this account.⁴⁹ For example:

Thinking makes it so

- T1. John believes that torturing others for fun is wrong and that belief would stand up to scrutiny in terms of John's other evaluative beliefs.
- T2. For any evaluative claim p believed by an agent A , it is true (for A) that p if and only if A 's belief that p would stand up to scrutiny in terms of A 's other evaluative beliefs.
- T3. It is true (for John) that torturing others for fun is wrong if and only if John's belief that torturing others is wrong would stand up to scrutiny in terms of John's other evaluative beliefs. (T1, T2)
- T4. It is true (for John) that torturing others for fun is wrong. (T1, T3)

This is an explanation, in Kitcher's sense, of *why it is true for John that torturing others for fun is wrong*. It is also an explanation that makes this correlation of an evaluative belief and an evaluative truth no coincidence: it provides an explanatory connection between the fact that John *believes* that torturing others for fun is wrong and the fact that it is *true* (for John) that torturing others for fun is wrong. By giving many similar explanations, one can provide an answer to the epistemic challenge. Assuming that actual human agents are fairly consistent in their evaluative outlook, the majority of their evaluative beliefs will stand up to scrutiny in terms of their other beliefs. Since this is so, positing an explanatory connection between the evaluative truths and an agent's evaluative beliefs is possible in many instances.

are true "for an agent." I will continue to include this relativization, but this feature of Street's view will not play a role in our discussion. (2006) pp. 152-154.

⁴⁹I am indebted to Selim Berker's helpful discussion of Street's anti-realism in articulating this explanation. See Berker (2014) pp. 232-235.

By the same token, **Thinking makes it so** might well contribute to unification. This particular explanation falls under a pattern whose instantiations explain many other evaluative truths. The pattern—call it Pattern T—emerges by replacing “John” with a variable fillable by any agent, and replacing “that torturing others for fun is wrong” with a propositional variable fillable by any evaluative proposition. Using this single pattern, one can derive a variety of claims about evaluative truths. So, answering the epistemic challenge in Street’s way might well furnish us with a set of explanations that fall under a broadly applicable, and so unifying, argument pattern. Other answers to the epistemic challenge, including those that seek to explain our evaluative beliefs in terms of the evaluative facts and those that seek to explain both in terms of some third factor, might well instantiate such an argument pattern.

I believe this is a promising line of thought for the challenger. However, employing Kitcher’s account of explanatory value to motivate the challenge requires something more. Specifically, we require a positive answer to this question: is *more* explanatory value to be expected from the set of explanations that connect our evaluative beliefs to the evaluative facts, than from *competing* sets of explanations? Sticking with our example, Pattern T’s instantiations will explain why various evaluative claims are true for an agent, by appealing to the agent’s belief in those evaluative claims. But there are many apparently *competing explanations* of why evaluative claims are true, which instantiate different argument patterns, and do not connect our evaluative beliefs to the evaluative facts. Some set of these might well do better, as far as unifying our beliefs. As a step toward elaborating on this, consider this competitor of **Thinking makes it so**:

Unnecessary pain

- P1. If an action inflicts pain unnecessarily, then that action is wrong.
- P2. Torturing others for fun inflicts pain unnecessarily.
- P3. Torturing others for fun is wrong. (P1, P2)
- P4. If p, then it is true that p.

P5. It is true that torturing others for fun is wrong. (P3, P4)

This is an explanation of *why it is true that torturing others for fun is wrong*. One stringent argument pattern that **Unnecessary pain** instantiates—call this Pattern U—can be obtained by replacing “Torturing others for fun” with a variable fillable by any action.

Comparing the explanatory value of **Thinking makes it so** and **Unnecessary pain** is not straightforward.⁵⁰ Remember that it is no objection to **Unnecessary pain** that it does not explain the correlation between our evaluative belief and the evaluative truth, in the sense of making this no coincidence. What matters is whether **Unnecessary pain** belongs to the set of explanations that best unifies our beliefs.⁵¹ Whether it does depends on a comparison with its competitors, in terms of the stringency and the applicability of the argument pattern it instantiates.

For Pattern T, Pattern U, and the explanations that fall under them, it is not at all clear which pattern and its instantiations will offer greater explanatory value. On the one hand, Pattern U’s instantiations explain the truth of evaluative claims about the wrongness of actions, while Pattern T’s instantiations explain the truth of evaluative claims more broadly. On the other hand, unlike Pattern U, Pattern T’s instantiations will only explain the truth of evaluative claims that are *believed by some agent*. No instanti-

⁵⁰This comparison might appear especially difficult, since the conclusions of some pairs of explanations that instantiate Pattern T and Pattern U appear to conflict. E.g. for some hypothetical agents who believe that torturing others for fun is not wrong, an explanation instantiating Pattern T will imply that this is *not wrong* (for them), while an explanation instantiating Pattern U will imply that this is *wrong*. But even if *some* of the implications of competing sets of explanations conflict, provided there is still a great deal of overlap, we can compare their explanatory value. (Even if we currently believe one implication and disbelieve the other, we might have reason to compare their explanatory value, since an increase in explanatory value can motivate changing one’s mind. See Kitcher (1989) section 7.4.)

⁵¹**Unnecessary pain** might still play a role in a *coincidence* explanation of the correlation. This would depend on explanations instantiating Pattern U (such as **Unnecessary pain**) as well as explanations of our evaluative beliefs that do not instantiate Pattern U. This approach commits us to at least one further argument pattern, beyond Pattern U. However, since psychological phenomena also stand in need of explanation, this argument pattern may well be broadly applicable. Thanks to Jay Wallace for pressing me to clarify this point.

ation of Pattern T is helpful in deriving the truth of an evaluative claim that no agent currently believes.⁵² In light of this, Pattern U's instantiations may explain more types of facts, after all. They explain the truth of claims about the wrongness of actions, regardless of whether anyone believes them. Such claims may come in a far greater variety than evaluative claims that are believed by some agent.⁵³ Tying our explanation of the evaluative truth to our current psychology, in the way characteristic of many answers to the epistemic challenge, threatens to restrict the range of evaluative truths explained.

This completes our discussion of explanatory value, as a motivation for answering the epistemic challenge. This discussion, like our discussion of the unlikelihood of being reliable by coincidence, was inconclusive. However, it indicates what is needed to take explanatory value as one's reason to answer the challenge. The challenger will need to offer a more definite view of the value of explanation. They might do so by accepting the unificationist picture of explanatory value, and arguing that explaining the correlation is likely to be superior, on this view, to providing no such explanation. Our discussion shows how they might begin to give this argument. Alternatively, the epistemic challenger might offer a competing account of the intellectual value of explanation, on which explaining our evaluative reliability is of unambiguous value. What they cannot do is motivate the challenge through a vague appeal to the value of explanation. Such value is not obviously at stake in seeking an explanation of the correlation between our evaluative judgments and the evaluative truths.

⁵²This might seem like a problem peculiar to our reconstruction of Street's anti-realism. A different reconstruction may do better. Consider **Anti-realism 2**: For any evaluative claim p and agent A , p is true (for A) if and only if A would come to believe p after scrutinizing p in terms of A 's other evaluative beliefs. An argument pattern using **Anti-realism 2** would have broader applicability, but it would raise a different question: would instantiations of such a pattern count as answers to the epistemic challenge? This is an account of what is true in terms of what we would believe, if we were to consider an evaluative claim in the light of our other evaluative beliefs. Can this account help to explain why our actual evaluative beliefs tend to align with the evaluative truth? Once again, work remains to be done.

⁵³Pattern U might also be more *stringent* than Pattern T. Explaining the truth of all evaluative claims using a single pattern seems a bit like explaining the truth of all physical claims using a single pattern: perhaps the instantiations falling under such a pattern will not have enough in common.

3.5 Conclusion

In this chapter, we have considered the demand for explanation, just in itself, as a motivation for answering the epistemic challenge. We found two ways of spelling out this demand. First, a challenger might argue that the correlation that constitutes our evaluative reliability is also overwhelmingly likely to have an explanation. If so, one should answer the epistemic challenge in order to avoid commitment to a likely falsehood. Second, a challenger might appeal to our general intellectual interest in explaining the world. If a theory implies that the correlation involved in our evaluative reliability is inexplicable, perhaps it implies that we cannot satisfy that interest, in at least that case. I have argued that neither case is straightforward for the epistemic challenger. Once again, we have no clear reason to seek an answer to the epistemic challenge.

Bibliography

- [1] Aristotle. *Posterior Analytics*. In *The Basic Works of Aristotle*, pp. 110-186. Ed. McKeon. New York: Random House. 1941.
- [2] Bedke, Matthew. 2014. "No Coincidence?" In *Oxford Studies in Metaethics, Volume 9*, pp. 102-125. Ed. Shafer-Landau. Oxford: Oxford University Press. 2014.
- [3] Behrends, Jeff. 2013. "Meta-normative Realism, Evolution, and Our Reasons to Survive." *Pacific Philosophical Quarterly*. 94: 486-502.
- [4] Berker, Selim. 2014. "Does Evolutionary Psychology Show that Normativity is Mind-Dependent?" In *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*, pp. 215-250. Eds. D'Arms and Jacobson. Oxford: Oxford University Press. 2014.
- [5] Buchanan, Allen. 2020. *Our Moral Fate*. Cambridge: The MIT Press.
- [6] Burgess, John and Rosen, Gideon. 1997. *A Subject with No Object: Strategies for Nominalistic Interpretation of Mathematics*. Oxford: Clarendon Press.
- [7] Clarke-Doane, Justin. 2012. "Morality and Mathematics: The Evolutionary Challenge." *Ethics*. 122: 313-340.
- [8] ——— 2017. "Objectivity and reliability." *Canadian Journal of Philosophy*. 47: 841-855.
- [9] Clifford, W.K. 1879. "The Ethics of Belief." In *Reason and Responsibility: Readings in Some Basic Problems of Philosophy (Fifteenth*

- Edition*), pp. 125-128. Eds. Feinberg and Shafer-Landau. Boston: Wadsworth. 2013.
- [10] Collins, Robin. 2008. "The Fine-Tuning Argument." In *Readings in the Philosophy of Religion*, pp. 84-95. Ed. Kelly James Clark. Peterborough: Broadview Press. 2008.
- [11] Copp, David. 2008. "Darwinian Skepticism about Moral Realism." *Philosophical Issues*. 18: 186-206.
- [12] Corvino, John. 2015. "The Fact/Opinion Distinction." *The Philosophers' Magazine*. Date published: March 4, 2015. Date accessed: December 15, 2020. <<https://www.philosophersmag.com/essays/26-the-fact-opinion-distinction>>.
- [13] Cosmides, Leda; Guzman, Ricardo; and Tooby, John. 2019. "The Evolution of Moral Cognition." In *The Routledge Handbook of Moral Epistemology*, pp. 174-228. Eds. Zimmerman, Jones, and Timmons. New York: Routledge Publishing. 2019.
- [14] Craig, William. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- [15] David, Marian. 2001. "Truth as the Epistemic Goal." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, pp. 151-169. Ed. Matthias Steup. Oxford: Oxford University Press. 2001.
- [16] De Waal, Frans. 2006. *Primates and Philosophers: How Morality Evolved*. Ed. Macedo and Ober. Princeton: Princeton University Press.
- [17] Dworkin, Ronald. 1996. "Objectivity and Truth: You'd Better Believe It." *Philosophy & Public Affairs*. 25: 87-139.
- [18] Enoch, David. 2009. "The epistemological challenge to metanormative realism: how best to understand it, and how to cope with it." *Philosophical Studies*. 148: 413-438.
- [19] ———. 2011. *Taking Morality Seriously*. Oxford: Oxford University Press.

- [20] Enoch, David and McPherson, Tristram. 2017. "What do you mean 'This isn't the question'?" *Canadian Journal of Philosophy*. 47: 820-840.
- [21] Field, Hartry. 1989. *Realism, mathematics and modality*. Oxford: Basil Blackwell Ltd.
- [22] Fisher. 1975. "Fallibility and knowledge of the future." *Philosophy and Phenomenological Research*. 36: 44-58.
- [23] Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis*. 23: 121-123.
- [24] Goldman, Alvin. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy*. 73: 771-791.
- [25] Greco, John. 2007. "The Nature of Ability and the Purpose of Knowledge." *Philosophical Issues*. 17: 57-69.
- [26] ——— 2009. "The Value Problem." In *Epistemic Value*, pp. 313-321. Eds. Haddock, Millar, and Pritchard. Oxford: Oxford University Press.
- [27] ——— 2003. "Knowledge as Credit for True Belief." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 111-134. Eds. DePaul and Zagzebski. Oxford: Oxford University Press. 2003.
- [28] Harman, Gilbert. 1977. *The Nature of Morality*. New York: Oxford University Press.
- [29] Harman, Gilbert and Thomson, Judith Jarvis. 1996. *Moral Relativism and Moral Objectivity*. Cambridge: Blackwell Publishers.
- [30] Hills, Allison. 2010. *The Beloved Self: Morality and the Challenge from Egoism*. Oxford: Oxford University Press.
- [31] Hume, David. 1888. *A Treatise of Human Nature*. Eds. Nidditch and Selby-Bigge. Oxford: Clarendon Press. 1978.
- [32] Hutchinson, Jim. 2019. "What can't what is true be valuable?" *Synthese*. ISSN: 0039-7857.

- [33] Joyce, Richard. 2006. *Evolution of Morality*. Cambridge: The MIT Press.
- [34] Kahane, Guy. 2017. "If Nothing Matters." *Nous*. 51: 327-353.
- [35] Kaplan, Mark. 1985. "It's Not What You Know that Counts." *The Journal of Philosophy*. 82: pp. 350-363.
- [36] Kolodny, Niko. 2006. "Objectivity in Ethics." In *The Encyclopedia of Philosophy, Second Edition*. Ed. Borchert. Detroit: Macmillan Reference USA. 2006.
- [37] Kitcher, Philip. 1989. "Explanatory Unification and the Causal Structure of the World." In *Scientific Explanation*, pp. 410-505. Eds. Kitcher and Salmon. Minneapolis: University of Minnesota Press. 1989.
- [38] Korsgaard, Christine. 2006. "Morality and the Distinctiveness of Human Action." In *Primates and Philosophers: How Morality Evolved*, pp. 98-119. Ed. Macedo and Ober. Princeton: Princeton University Press. 2006
- [39] Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- [40] Lackey, Jennifer. 2007. "Why we don't deserve credit for everything we know." *Synthese*. 158: 345-361.
- [41] Locke, Dustin. 2014. "Darwinian Normative Skepticism." In *Challenges to Moral and Religious Belief: Disagreement and Evolution*, pp. 220-236. Eds. Bergmann and Kain. Oxford: Oxford University Press. 2014.
- [42] Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. London: Penguin Books.
- [43] Moore, G.E. 1939. "Proof of an External World." *Proceedings of the British Academy*. 25: 273-300.
- [44] Morrison. 2011. "Skepticism about Inductive Knowledge." In *The Routledge Companion to Epistemology*, pp. 445-453. Eds. Pritchard and Bernecker. New York: Routledge. 2011.

- [45] Nagel, Thomas. 1997. *The Last Word*. New York: Oxford University Press.
- [46] Oddie, Graham. 2005. *Value, Reality, and Desire*. Oxford: Oxford University Press.
- [47] Parfit, Derek. 2011. *On What Matters, Volume Two*. Oxford: Oxford University Press.
- [48] Sayre-McCord. 1988. "Moral Theory and Explanatory Impotence." In *Essays on Moral Realism*, pp. 256-281. Ithaca: Cornell University Press. 1988.
- [49] Plato. *Meno*. In *Five Dialogues (Second Edition)*, pp. 58-92. Tr. Grube and rev. Cooper. Indianapolis: Hackett Publishing Company, Inc. 2002.
- [50] Rabinowicz, Wlodek and Ronnow-Rasmussen, Toni. 2003. "Tropic of Value." *Philosophy and Phenomenological Research*. 66: 391-423.
- [51] ——— 2004. "The strike of the demon: on fitting pro-attitudes and value." *Ethics*. 114: 391-423.
- [52] Riggs, Wayne. 2002. "Beyond Truth and Falsehood: The Real Value of Knowing That P." *Philosophical Studies*. 107: 87-108.
- [53] Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge: The Belknap Press of Harvard University Press.
- [54] ——— 2014. *Being Realistic About Reasons*. Oxford: Oxford University Press.
- [55] Schafer, Karl. 2010. "Evolution and Normative Scepticism." *Australasian Journal of Philosophy*. 88: 471-488.
- [56] Setiya, Kieran. 2012. *Knowing Right from Wrong*. Oxford: Oxford University Press.
- [57] Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.
- [58] ——— 2012. "Evolutionary Debunking, Moral Realism, and Moral Knowledge." *Journal of Ethics & Social Philosophy*. 7: 1-37.

- [59] Sober, Elliot. 2012. "Coincidences and How to Reason About Them." *European Philosophy of Science Association Proceedings*. 1: 355-374.
- [60] Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford University Press.
- [61] Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies*. 127: 109-166.
- [62] ——— 2016. "Objectivity and Truth: You'd Better Rethink It." In *Oxford Studies in Metaethics, Volume 11*, pp. 293-333. Ed. Shafer-Landau. Oxford: Oxford University Press. 2016.
- [63] Stroud, Barry. 1984. *The Significance of Philosophical Scepticism*. Oxford: Oxford University Press.
- [64] ——— 2001. "What is Philosophy?" In *What is Philosophy?*, pp. 25-46. Eds. Ragland and Heidt. New Haven: Yale University Press. 2001.
- [65] Sturgeon, Nicholas. 1984. "Moral Explanations." In *Ethical Theory: Classical and Contemporary Readings*, pp. 526-537. Ed. Pojman. Belmont: Thomson Wadsworth. 2007.
- [66] Tomasello, Michael. 2015. *A Natural History of Human Morality*. Cambridge: Harvard University Press.
- [67] Whitcomb, Dennis. 2007. "Intellectual goods: An epistemic value theory." Doctoral dissertation. Cambridge: Rutgers University.
- [68] White, Robert. 2000. "Fine-Tuning and Multiple Universes." *Nous*. 34: 260-276.
- [69] Wielenberg, Erik. 2010. "On the Evolutionary Debunking of Morality." *Ethics*. 120: 441-464.
- [70] Williams, Bernard. 1995. "Who Needs Ethical Knowledge?" In *Making sense of humanity*, pp. 203-212. Cambridge: Cambridge University Press. 1995.

- [71] Williamson, Timothy. 2002. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- [72] Zagzebski, Linda. 2003. "The Search for the Source of Epistemic Good." *Metaphilosophy*. 34: 12-28.