

Evaluative Realism and the Argument from
Queerness

Luke Townend

PhD

University of York

Philosophy

September 2020

Abstract

Purely Evaluative Moral Realism is the view that there are moral facts, but they are always and only concerned with goodness and badness, without being concerned with reasons or obligations. In this book, I argue that such a view has the virtue of avoiding the Argument from Queerness for Moral Error Theory. Firstly, I lay out and examine the Argument from Queerness to show how it is a threat to moral realism. Next, I isolate and argue against one key premise of the Argument from Queerness; that all moral facts are specially normative, which is to say concerned with a distinctive type of reason for action. Denying this claim creates conceptual space for Purely Evaluative Moral Realism. I then describe this view and demonstrate how it avoids the Argument from Queerness. In addition to all this, I defend Purely Evaluative Moral Realism from objections to show that it does not collapse into any sort of incoherence and is a legitimate rival to more familiar metaethical views. In doing so, I argue that Purely Evaluative Moral Realism is compatible with a surprisingly wide range of theories within substantive ethics, and that it allows us to account for many of our intuitions about moral language and moral reasons, so long as we accept that moral reasons may be more like our other reasons than is usually thought, and that they lack the features distinct to special normativity.

Contents

Abstract	2
Contents	3
Preface	6
Acknowledgements	11
Author's Declaration	12
1 The Argument from Queerness	13
1.1 Moral Error Theory	13
1.2 The Structure of the Argument from Queerness	18
1.3 Debunking Accounts	22
1.4 Queerness Arguments	33
1.5 The Queerness Argument from Moral Knowledge	47
1.6 The Queerness Argument from Motivation	50
1.7 The Queerness Argument from Supervenience	60
1.8 The Queerness Argument from Special Normativity	72
2 Normativism and Non-Normativism	87
2.1 The Normativity of Moral Facts	88
2.2 The Connection between Norms and Values	97
2.2.1 Moorean Consequentialism as a Conceptual Truth	98
2.2.2 Values to Be Respected	100
2.2.3 From Personal Goodness to Moral Norms	101

2.2.4	The Buck-Passing Account of Value	104
2.3	The Case for and against the Buck-Passing Account	108
2.3.1	The Ontological Parsimony Argument	109
2.3.2	The Heterogeneity Argument	111
2.3.3	The Open-Feel Argument	116
2.3.4	The Demystification Argument	118
2.3.5	The Redundancy Argument	120
2.3.6	The Similarity Argument	122
3	Purely Evaluative Moral Realism	127
3.1	What Purely Evaluative Moral Realism Is and What It Isn't	128
3.1.1	Commitments of Purely Evaluative Moral Realism	128
3.1.2	Contrasting Purely Evaluative Moral Realism with Similar Views	133
3.2	Purely Evaluative Moral Realism and the Argument from Queerness	136
3.2.1	Purely Evaluative Moral Realism and the Argument from the Queerness of Special Normativity	136
3.2.2	Purely Evaluative Moral Realism and the Argument from the Queerness of Supervenience	137
3.3	Objections to Purely Evaluative Moral Realism	140
3.3.1	Restriction on Substantive Theories	141
3.3.2	Queerness Arguments	142
3.3.3	Purely Evaluative Moral Realism and Action	147
4	Moral Facts and the Guiding of Action	155
4.1	Purely Evaluative Moral Realism and Substantive Ethics	156
4.1.1	Scalar Utilitarianism	157
4.1.2	Virtue Ethics	160
4.1.3	Deontology by Way of Honouring Values	165
4.1.4	Divine Command Theory	168
4.2	Realism about Non-Special Moral Norms	172
4.3	Objections to Realism about Non-Special Moral Norms	181

<i>CONTENTS</i>	5
4.3.1 Problems with the Analogy between the Moral and Epistemic Domains	182
4.3.2 Denial of the Self-Consciously Moral Motive	188
Conclusions and Summary	195
A Virtues of Purely Evaluative Moral Realism	197
A.1 Dissolving the Is-Ought Problem	197
A.2 Free Will and Moral Responsibility	199
Bibliography	203

Preface

Could we account for morality by making only evaluative claims—statements about what is good and bad—without any reference to reasons or duties? In what follows, I will defend just such an account, which I call Purely Evaluative Moral Realism. This enquiry is chiefly one into metaethics. First-order or substantive ethics (what we typically refer to simply as ‘ethics’) is concerned with which things in the world fall into moral categories like right, wrong, good, and bad. By contrast, second-order ethics or metaethics is concerned with the metaphysical nature of those categories. It asks what it means for something to be good, bad, right, wrong, and so on, and whether those categories exist in the first place.

That second question is the question of whether any (positive) moral claims are, or can be, true (and, implicitly, of whether anything exists in the world that can make them true). In other words, it is the question of whether there are any moral facts. The view that there are moral facts is called moral realism. It brings with it a commitment to cognitivism, the view that moral language is meant to express facts and so moral claims are truth evaluable. Moral realists are also usually objectivists, taking it that moral facts are not mind dependent, at least insofar as they are not made true or false by anyone’s believing them (other mental states, like desire or happiness, may well bear on what the actual moral facts are).

The contrasting view is moral antirealism, but in what follows we will be concerned with a particular version thereof, namely Moral Error Theory. This form of antirealism shares realism’s commitment to cognitivism, but all the same denies that there are any moral facts. As such, Moral Error Theorists maintain that positive moral claims are necessarily and systematically false.¹

¹By positive claims, I mean those that claim something *does* have a property like goodness

The findings of this thesis have a direct bearing on the realm of substantive ethics; if Moral Error Theorists are right, there is nothing for substantive ethicists to study. The claim that I will defend in this thesis is that Purely Evaluative Moral Realism is immune to the Argument from Queerness. Purely Evaluative Moral Realism is realism about a limited class of moral facts, while the Argument from Queerness is perhaps the most influential argument for Moral Error Theory. Purely Evaluative Moral Realism limits the range of possible moral facts where a more conventional unqualified moral realism would not, because it entails that nothing is right or wrong. But it preserves claims about goodness and badness, so it does not eliminate the domain completely as Moral Error Theory does.

More will be said about what exactly the Argument from Queerness and Purely Evaluative Moral Realism involve in Chapters 1 and 3 respectively, but hopefully the relevance of the one to the other is already clear, even if the details remain obscure at this juncture. If Purely Evaluative Moral Realism can avoid the Argument from Queerness, it will mean that this argument does not lead directly to Moral Error Theory as many have supposed (though it would still warrant the rejection of a great many first-order moral claims). In order for a view to avoid the whole argument, it must avoid each version thereof that poses some genuine threat to realism, so we can only answer in the affirmative if Purely Evaluative Moral Realism allows for further commitments that meet the conditions to avoid each incarnation of the Argument from Queerness without internal contradiction.

One other thing to note is that we can only claim that a view can achieve anything in particular if that view passes the test of basic coherence and plausibility. As such, I will offer a full defence of Purely Evaluative Moral Realism. What I will not do, however, is argue that it is the correct view of the metaphysics of morality. I have plenty to say about its virtues and implications, but the task of

or wrongness. As noted in Olson 2014, pp. 11-5, Moral Error Theorists cannot simply say that all moral claims are false without qualification because under their view, claims that something does not have some moral property (negative moral claims) are universally true. Moral Error Theorists can make this move so long as they deny the implication from ‘ ϕ is not wrong’ to ‘ ϕ is right or permissible’ as suggested by Pigden (2007, p. 453), and this seems reasonable since they take all of these categories to be empty.

comparing it to alternative views is largely ancillary to our discussion. So long as it is not obviously inferior to alternatives, it will remain a live possibility, which will be sufficient to undermine the Argument from Queerness as a direct route to Moral Error Theory.

The first chapter will be devoted to describing the Argument from Queerness and its various incarnations, and establishing the strongest forms thereof to more fully motivate the search for a form of realism that does not fall prey to those arguments. The Argument from Queerness concludes that there are no moral facts (or, more precisely, that we should not posit or believe in moral facts) because they do not figure into our best explanations of the world, and specifically of human moral thoughts, feelings, and behaviour. What makes explanations invoking moral facts worse than the alternatives is the queerness of moral facts, which is a particular form of strangeness sufficient to motivate ontological suspicion of those entities that possess it. Thus to accept the Argument from Queerness we must accept one or more queerness arguments, which nominate some feature of moral facts as the source of their putative queerness. I will conclude that two of these succeed, namely the ones concerning moral facts' normativity and their supervenience on natural facts, meaning that there are two successful versions of the Argument from Queerness to overcome.

The second chapter will examine a key premise from one of those successful versions. I mean to reject this premise, which I will label Strong Normativism, as it is directly incompatible with Purely Evaluative Moral Realism. It is the claim that the particular and unique normativity so often attributed to moral facts is conceptually necessary to moral facts, meaning that all such facts would be reason-giving in this queer way. Here, I will posit a distinction between normative facts and evaluative facts, the latter of which are not concerned with reasons or duties but with the degree of goodness or badness of a given thing. If evaluative facts can exist in the absence of normative facts, then Strong Normativism is false. I will go on to examine views that imply a necessary connection between normative facts and evaluative facts, foremost among which is the buck-passing account of value, and I will rebut arguments for these views. This opens the door for a form of moral realism that rejects Strong Normativism.

That form of realism will be the subject of the third chapter; Purely Evaluative

Moral Realism. This is the view that there are some evaluative moral facts, and no normative ones. Things can be good, but not strictly right or obligatory in the relevant moral sense. I will show how this view can overcome all the functional versions of the Argument from Queerness, and then address some objections to the view. While most of these can be addressed fairly quickly, one in particular will raise enough issues to warrant extended discussion. This objection arises from the idea that morality must be action-guiding and that no view of morality that fails to afford it this power can be sufficient. If there is some necessary connection between the guiding of action and normativity, then Purely Evaluative Moral Realism will fall short as a metaethical view.

In the fourth chapter, I will rebut this argument in two ways. First, I will show how a number of extant views in substantive ethics are, or can be made, compatible with Purely Evaluative Moral Realism. I will also argue that the moral facts posited by these views can guide our actions in the relevant way and so Purely Evaluative Moral Realism, insofar as it can account for all their claims, does not rule out the action-guidingness of moral facts. Secondly, from those examples I will abstract a general way of justifying certain normative claims—more limited in power and scope than is typically posited for those in the moral domain, and more akin to other norms generated through straightforward practical reasoning—with reference to evaluative facts and general claims about human psychology. Either of these alone should be sufficient to show that, even without a specially normative element, morality is very much able to guide action in the required way.

Having done all this, we will have a clear account of what Purely Evaluative Moral Realism is, and how at least one version of it falls outside of the scope of the Argument from Queerness. We will also have responses to many of the criticisms that such a view is likely to garner, so we be in a position to say that the view is defensible. Following these conclusions, there will be a brief appendix discussing some of the virtues of Purely Evaluative Moral Realism. While this is, as I have indicated, not the central concern of this thesis, they mitigate the worry that Purely Evaluative Moral Realism is an ad hoc response to the Argument from Queerness.²

²In point of fact, Purely Evaluative Moral Realism can be seen as a response to many

metaethical puzzles that, in one way or another, are owed to the normativity of moral facts. The Argument from Queerness became the focus of this project because it is, in my estimation, the most well-developed of these, and the most legitimate threat to moral realism.

Acknowledgements

This work was supported by the Arts & Humanities Research Council (grant number AH/L503848/1) through the White Rose College of the Arts & Humanities.

I would like to extend my gratitude to my supervisors, Johan Gustafsson and Christian Piller, whose insight and advice have been invaluable to me in assembling this thesis. I would also like to thank Mary Leng, Paul Noordhof, Tom Stoneham, and Richard Yetter Chappell for their valuable commentary on in-progress versions of the text, as well as to Jonas Olson and Charles Pigden for their remarks on some of the core ideas presented herein, and to the membership of my initial examination panel, whose remarks have shaped this revised version of the project.

Additional thanks are owed to all the participants in the University of York's postgraduate work-in-progress seminars, especially Celia Coll, Britt Harrison, and Bridger Landle. Finally, I am immensely grateful for the kindness and support of my mother and father, my brothers, and all of those in York's Department of Philosophy who have extended their friendship to me over the course of this project.

Author's Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Chapter 1

The Argument from Queerness

The Argument from Queerness expresses a general worry about moral realism: there is something strange about the idea of moral truth, which doesn't quite cohere with the rest of the world as we understand it. In this chapter, I will clarify this argument's assumptions, intended conclusions, and overall structure. The Argument from Queerness consists of two subordinate arguments for conclusions which, taken together, are meant to lead to moral error theory. One of these subordinate conclusions is that we can explain or account for our moral thoughts, feelings, and actions without invoking moral facts. The other is that if there were any moral facts, they would be metaphysically queer. It is this latter conclusion which is of greatest interest for our discussion, so the latter half of this chapter will be devoted to examining the many arguments for it.

1.1 Moral Error Theory

The Argument from Queerness is generally posed as an argument for the ontological aspect of Moral Error Theory. Moral Error Theory consists of two core claims: a semantic claim, and an ontological claim. The semantic claim is that moral language is truth-apt, in the sense that when we make moral claims we are trying to express facts about the world, and those claims are capable of being true or false. The Argument from Queerness takes this former claim as given (though Moral Error Theorists usually have independent arguments for it, which I discuss

briefly below). The ontological claim, which is the conclusion of the Argument from Queerness, is that there are no moral facts or properties.³

I should clarify here what I mean by ‘facts’ in general and ‘moral facts’ in particular. A fact is a state of affairs that obtains, and a fact exists whenever a particular entity has a property or stands in a relation.^{4,5} Facts are truthmakers in the sense that the truth values of propositions depends on what facts there are. Insofar as we can assume that truth involves a world-to-content relation, facts fall on the side of the world.⁶

All facts are about the distribution of properties (and relations, but I will discuss properties for brevity). Moral facts, then, are about the distribution of moral properties.⁷ In other words, they are about whether actions (and sometimes things closely associated with them, such as motives, character traits, and consequences among other things) are good or bad, or right or wrong in a particular way. The particular way in question—the moral—is distinct from other standards by which actions might be assessed, such as practical rationality, but is nevertheless such

³Error theories more generally are views such that all the claims within some domain of discourse are false. Moral Error Theory fits this schema in that it commits its followers to the falsehood of all positive moral claims, which follows from the ontological claim that there are no moral facts. There are error theories about other domains, to which occasional reference will be made in what follows, but this text is primarily concerned with Moral Error Theory so wherever not otherwise stated it is to this variety that I refer.

⁴Armstrong 1997, p. 1; Finlay 2014, p. 7. Facts might also usefully be described as actual distributions of properties and relations. While I have borrowed Armstrong’s definition of facts, I should make clear that I do not mean to commit myself to any part of the broader factualist ontology that he develops.

⁵This definition might be simplified if we can account for relations as relational properties, but that is not a question that we need to settle for the purposes of this project.

⁶Armstrong 1997, p. 13.

⁷Taking this in conjunction with the definition of facts given above, it follows that the claim that there are no moral facts is equivalent with the claim that there are no moral properties, which is why I spoke of them together when I first laid out the ontological claim of Moral Error Theory.

that all else is subject to it. Rather to characterize the morality any further, I take it that my readers have enough of a pretheoretical grasp of the moral domain to catch the intended meaning of ‘moral goodness’, ‘moral wrongness’, and ‘moral reasons’ well enough to be able to follow my investigations.

If there are no moral facts or properties, then no positive moral proposition is true.⁸ Positive moral propositions are those that attribute a substantive moral property (and so metaethical claims of the sort I am examining here are excluded from this class).⁹ From the claim that no moral proposition is true, it follows that no moral proposition could possibly be true either. This is because, as I discuss in §1.7, the moral status of a thing is a matter of necessity in the sense that if some action has a certain moral status, there is no possible world where that same action performed under all the same circumstances fails to have that moral status.

⁸It follows that beliefs with positive moral propositions as their content and sentences that express positive moral propositions are also necessarily false.

⁹I qualify the propositions in question as ‘positive’ because, if all claims of the type described above are false, this would include statements to the effect that such-and-such a thing is *not* good or bad or right or wrong. The falsehood of any claim implies the truth of its negation, so the Moral Error Theorist cannot simply say that all moral propositions are false without qualification, lest they place themselves in a contradiction (Olson, 2014, pp. 11-2).

This is not the end of trouble for the framing of Moral Error Theory, however. Thin moral concepts such as those mentioned above seem to come in opposed pairs (right and wrong, obligatory and forbidden, et cetera) whose members are often understood to be negations of one another, so that when you deny that something is right (as a Moral Error Theorist’s view obliges them to do) it seems that you claim that it is wrong. I owe my preferred solution to this problem to Pigden (2007, pp. 453-4), who denies the implication from ‘not right’ to ‘wrong’. Wherever we have two contradictory moral properties or categories like this, there is always the third possibility of having no moral status whatsoever, and Moral Error Theory is effectively the claim that all things fit into this category; nothing is good or bad or right or wrong or even morally neutral. Instead, those categories simply fail to apply in all cases. What this means, in effect, is that all *negative* moral propositions are true, because every action is both not right and not wrong, or not good and not bad. The idea of restricting Moral Error Theory to positive claims comes from Sinnott-Armstrong 2006, pp. 34-6, but it works only in light of Pigden’s observations.

As I have already stated above, the Argument from Queerness is not strictly an argument for the whole of Moral Error Theory. Rather it is an argument for the ontological thesis, and so against a particular variety of moral realism—non-naturalist realism. When J. L. Mackie advanced it in his book *Ethics: Inventing Right and Wrong*, it was alongside other arguments ruling out alternative versions of realism, making it the last of a series of arguments which, taken together, were meant to support Moral Error Theory.^{10,11}

Moral Error Theorists and non-naturalist moral realists share some key commitments about the nature of moral facts, which I will grant for the purposes of this chapter. These commitments are cognitivism and non-naturalism. Cognitivists hold that moral claims express a belief held by the speaker, in contrast to non-cognitivists who hold that moral language elliptically expresses emotions, prescriptions, or some other non-cognitive (read: non-belief) attitude. This is effectively the semantic claim of Moral Error Theory, as it makes the content of moral claims propositional and truth-apt; they are statements capable of being true or false.

The standard argument used by Moral Error Theorists (although not by them alone) to justify their commitment to cognitivism starts with the claim that the meaning of moral terms is determined by the intentions of ordinary speakers when they use them.¹² These intentions are revealed through ordinary moral discourse, which proceeds as though moral claims express exactly the kind of truth evaluable proposition that cognitivists hold them to express. This is what allows for full-blooded disagreements over questions of morality, and for moral claims to be embedded in formulae of propositional logic.¹³ I will not pretend that this alone

¹⁰Mackie 1977, pp. 23-4.

¹¹While Mackie's Argument from Queerness, and many that have followed after it, target non-naturalist realism, some versions of it may also target naturalist realism (Olson, 2017, pp. 64-5). This is true of those versions which rest on a queerness argument (see below) whose targeted feature of moral facts could be naturalized without undermining the claim that said feature is queer (see §1.4).

¹²Mackie 1977, pp. 32-3.

¹³If moral claims were not actually propositional in content, we would run into the Frege-

does away with non-cognitivism, even if it deals a significant blow to simpler forms thereof.¹⁴ I will, however, note that this project is foremost about what the Argument from Queerness proves. As such, it is the dialectic between a robust realism and Moral Error Theory which is of most interest to us. In addition, I will try to meet Moral Error Theorists on their own terms with regard to these questions. I will, therefore, treat cognitivism as given from this point.

Non-naturalists deny moral naturalism, which is the claim that moral facts are either identical to, reducible to, or constituted by natural facts.¹⁵ The issue of what is and is not natural is a knotty one, but it will come up often enough over the course of my project to make untangling it here worthwhile. For philosophers, the domain of the natural is that of the familiar temporal world, and we can follow many others in stipulating that our guide to what counts as natural is that which is posited by our best current science.¹⁶ General or metaphysical naturalism (as opposed to moral naturalism) is the view that the natural facts are exhaustive, that there are no facts beyond those countenanced by science.¹⁷ Moral Error Theorists and moral realists alike most often appeal to the normativity of moral facts (see §1.8) to justify their commitment to non-naturalism; there is, they claim, simply no way to account for moral facts' reason-giving properties in a strictly naturalistic framework.¹⁸ (I am not so strongly married to this position because of

Geach or embedding problem (discussed in Geach 1958; 1965.)

¹⁴In particular, moral quasi-realism as developed in Blackburn 1984; 1993 is notable in that it is meant to have the advantage of reflecting or vindicating ordinary moral language while still falling under the non-cognitivist banner. See also Hare 1952 and Gibbard 1990.

¹⁵This is a gloss, but a good one, which I have borrowed from Brink (1989, pp. 156-7). For a moral fact to be constituted by one or more natural facts is for it to be ontologically dependent thereupon in the sense that the moral fact is nothing over and above the natural fact or facts (as with reduction) while also being multiply realizable in the sense that the same moral fact or property have been grounded in different natural facts or properties.

¹⁶See Sturgeon 2006, p. 92, Kitcher 2011, pp. 3-4, and Olson 2014, p. 88. This is also one of the definitions discussed in Hampton 1998, p. 35. 'Best' here means best by science's own lights.

¹⁷This does not necessarily imply physicalism or materialism; see Chalmers 1996.

¹⁸See Mackie 1977, p. 32, Dancy 2006, p. 132, and Parfit 2011a, §134.

the conclusions I reach in Chapter 2, but I shall treat it as given for the purposes of examining the Argument from Queerness itself.)

Some of the various versions of the Argument from Queerness which I will examine in this chapter bring along further metaethical commitments, which I will examine as they become relevant. The above considerations will serve as an important baseline from which to make a broader assessment of these arguments. We also now have a clear intended conclusion in mind for the argument as a whole. This is the ontological claim of Moral Error Theory, that there are no moral facts, but since the Argument from Queerness takes the semantic claim as given I will speak of its supporting Moral Error Theory in general for brevity. Now that all of this is in place, we can move onto discussing the structure of the Argument itself.

1.2 The Structure of the Argument from Queerness

The Argument from Queerness is an argument from the best explanation, with its explanandum being morality itself. When I say ‘morality’ here, I do not mean the content of the correct system of substantive ethics. Rather, I mean morality as a phenomenon in the world, which is to say all of our moral thoughts, feelings, and actions. We can call these collectively ‘moral phenomena’. The subject of the debate between moral realists and moral antirealists is the ontological status of moral facts as defined in §1.1.

According to moral realists, our tendency to make moral judgements and respond to them with our actions is best explained by the existence of moral facts, but the Argument from Queerness is meant to refute this claim. Its conclusion is that not only are there other viable explanations for morality, but that there are also serious problems with the very concept of a moral fact, which together undermine the idea that moral realism is the best available explanation.

The Argument from Queerness has a bipartite structure; it comprises two subordinate arguments whose conclusions together imply that moral realism isn’t the best explanation for moral phenomena. The first of these is at least one debunking account of morality, which explains moral phenomena in natural terms without

reference to any special class of moral facts.

The second subordinate argument is at least one queerness argument, which is to say an argument to the effect that moral facts, if they existed, would be ‘queer’. This notion of queerness is a little obscure, and I examine it more thoroughly in §1.4, but for now suffice it to say that queer entities are ontologically or metaphysically suspicious in such a way that we shouldn’t admit them to our picture of the world without very good reasons. In other words, they are a cost or burden to any ontology that includes them. Typically queerness arguments nominate some feature of moral facts, and then argue that said feature is in some way problematic and so, therefore, are moral facts.¹⁹

So it seems that moral realism is under attack from two fronts; its commitments appear redundant in the face of debunking accounts, and inherently problematic thanks to queerness arguments. Neither of these considerations is a knockdown argument against realism on its own, but the threat they pose is significantly greater when they are given in combination. This is not just because of the mounting volume of considerations against it—though this is a kind of evidence in itself—but because each consideration undermines arguments against the other as evidence for the antirealist conclusion.²⁰

Consider debunking accounts. Taken on its own, a debunking account allows us to make something along the lines of Gilbert Harman’s redundancy argument against moral realism, where we reject moral facts on the basis of their redundancy alone.²¹ This involves an appeal to a simple quantitative version of Ockham’s razor, whereby we should not posit more entities (in this case facts) than are

¹⁹Olson 2014, p. 84.

²⁰The idea that the sheer volume of considerations for or against a view is evidence in itself will be particularly salient if we take seriously the view that little in philosophy is settled by decisive argument, and that we are better tallying the number and quality of arguments or evidence for a position as a metric of its plausibility. See Rawls 1971, p. 20; Lewis 1986, p. viii; and Enoch 2011, §1.4.

²¹Harman 1977, ch. 1. What debunking accounts actually show is that moral facts would be explanatorily redundant in the first place, and concluding directly from this that there are no moral facts is a further move we are not forced to make.

explanatorily required.²² The realist may wish to interject that all things are not equal because, they take it, the debunking account fails to completely explain moral phenomena, even if it tells a convincing story about the elements that it does explain. Such arguments often turn on the claim that whatever the debunking explanation misses can only be explained by a set of facts with some particular quality for which a naturalistic worldview has no room. However, when called to say what this quality is, realists run the risk of nominating the subject of one of the queerness arguments, in which case the debunking account seems all the more plausible for its lack of such troubling commitments. Of course, I cannot prove that every such suggestion would fall into this trap—it would be question-begging to assume that all non-natural properties must be queer—but the most popular and plausible candidate is the apparent normativity of moral facts, which is the subject of the queerness argument discussed in §1.8.

On the other side, one objection that could be fielded against the claim that we should reject moral facts simply because they are queer is that the ‘sheer queerness’ of moral facts does not seem to be enough to warrant rejecting them on its own. While it has come to be a philosophical term of art, when Mackie first used the word ‘queer’ it was simply a synonym for ‘strange’. As Platts points out, there are many strange things in the world such as aardvarks and neutrinos, yet we would look askance at anyone suggesting that because these things are so odd, we had best conclude that they don’t actually exist.²³ But even though these things might initially seem queer, when we place them into context as part of the world as a whole, they form part of our best explanations of things.²⁴ For example, though aardvarks exhibit many unusual traits, our sensory experiences of aardvarks are better explained by their actual presence than by hallucinations or some manner

²²In fact, at least some philosophers do not see the multiplication of particular entities as a theoretical vice (for example, see Lewis 1973, p. 87). There are arguments to the contrary as well (see Nolan 1997) but I am inclined to agree that more facts is not, in and of itself, much of a problem and where I appeal to Ockham’s razor myself it will be to the more standard qualitative version.

²³Platts 1980, p. 72.

²⁴Olson 2014, p. 87.

of elaborate conspiracy, precisely because we can also see the effect they have on their ecosystem, determine their evolutionary history from the fossil record, and otherwise resolve their existence with the rest of our understanding of the world without difficulty. So, while aardvarks are strange, debunking explanations of our beliefs in and about aardvarks (of which the alternatives I listed above are all examples) fail to be legitimate rival views in these cases. This is where the case of moral facts differ; if there *is* a plausible debunking account available, then the ontological problems raised by the queerness of moral facts begin to look much more serious.

Another realist response to queerness arguments is what are called ‘Moorean arguments’, not because G. E. Moore made such arguments, but because they mirror the structure an anti-idealist, anti-skeptical argument Moore made to defend common sense views about the external world.²⁵ Such arguments claim that there are some substantive moral propositions that are so obvious that we should have more confidence in them than in any of the premisses of an apparently successful queerness argument. The classic example would be the claim that it is wrong to torture an innocent child just for our own amusement. This would entail that there were at least one moral fact, and so long as that claim were indeed more plausible than any of the premisses for any of the queerness arguments, this would mean we have a very simple and plausible argument to the effect that we have sufficient reason to accept moral facts into our ontology despite their queerness. But we don’t need to take the time to compare this premiss to every premiss of every queerness argument, because here again the Moral Error Theorist can invoke a debunking account. If we have a convincing origin story for our moral beliefs (which doesn’t invoke moral facts), then we have an explanation not just of why we might have moral beliefs in general and some moral beliefs in particular, but also of why we hold them so strongly; whatever evolutionary or social purpose moral beliefs serve, they won’t serve it well if they are easily shaken. Particularly obvious examples (like torturing a child for fun) are particularly vulnerable to this, since the reason they are so obvious is likely that they rest on our most fundamental moral beliefs,

²⁵Olson 2014, pp. 139-40, commenting on arguments from Moore 1959 and Ewing 1947, pp. 30-2.

or on multiple beliefs, making them far harder to shake. So Moorean arguments fail to undermine queerness arguments, because debunking accounts give us reason to be suspicious of their key premisses.

So we have seen that the Argument from Queerness seriously undermines the idea that moral facts are the best explanation for moral phenomena, assuming both of its subordinate arguments are successful. We should not posit any fact without good cause, such as that it appears in our best explanations of the world, and since we have just established that moral facts do not form part of our best explanation of moral phenomena we are led straightforwardly to Moral Error Theory. The rest of this chapter will be devoted to determining whether the subordinate arguments are indeed successful. In §1.3, I will discuss debunking accounts in general, and describe what a successful one would look like. In §1.4, I will examine the notion of metaphysical queerness on which queerness arguments rely. The subsequent sections will examine the four queerness arguments identified by Jonas Olson to determine which are and are not successful, and thus to determine which version or versions of the Argument from Queerness warrant reply. Once we know this, we will be in a much better position to see whether Evaluative Realism can successfully answer the challenges raised. For clarity, I will call each queerness argument ‘the queerness argument from X’ where X is the feature of moral facts that it nominates as the queer element, and I will call the version of the Argument from Queerness that relies on the queerness argument from X ‘the Argument from the Queerness of X’.

1.3 Debunking Accounts

Before discussing debunking accounts, it is important to distinguish them from what are typically called ‘debunking arguments’. Such arguments begin with an origin story for a belief or beliefs of a given kind, such that the existence and content of those beliefs is in no way dependent on any corresponding facts. From the presence of such a story (and presumably some compelling reasons to believe it), they proceed to the conclusion that because the belief(s) in question were not formed through a truth-tracking process, we are not epistemically justified in

holding them.²⁶ They do not, therefore, warrant ontological conclusions of the kind at which the Argument from Queerness is aimed, and this is why I will not be making one in this section.

A debunking account of some belief is just the first step of a debunking argument: the origin story. A debunking account of beliefs about a given subject is an explanation of those beliefs that does not invoke any facts about that subject. For our purposes, a debunking account of morality is an effort to establish the starting claims about the origin of moral belief (and other moral phenomena, though these will simply be *explained* rather than debunked) which would be necessary for a debunking argument.

So a debunking account of morality explains moral phenomena without appealing to a special class of moral facts. As I have characterized them, moral phenomena might initially look to form a highly diverse class that cannot be explained in a unified way, encompassing as it does beliefs, emotions, and actions at the very least. But if each of these can be ultimately explained in terms one of the others, then really there is only one thing that must be explained. For example, a moral action might be undertaken in response to the belief that it is the right thing to do, and so explaining the belief would be sufficient for explaining the action.

Extant debunking accounts offered by Moral Error Theorists are typically projectivist ones. Under projectivism, agents project the moral emotions they feel in response to actions, people, and situations they encounter onto their causes and see them as objective properties of the world, namely moral properties. Put another way, they phenomenally experience the world as one with moral properties. Under this view, moral thoughts and the moral concepts employed in forming them are the result of these emotions, and explanatorily anterior to them.²⁷ It is to these emotions (although plausibly also to any other solely experiential element of the moral tapestry) that I refer when I speak of moral feelings. All of that said, there is no strict reason that there could not be a debunking account that treated these phenomena with a different explanatory priority, so long as it was still a unified

²⁶Kahane 2011, pp. 106, 111.

²⁷For a more comprehensive discussion of projectivism in Moral Error Theory, see Olson 2014, chs. 1-4.

explanation, and I will by and large speak of moral phenomena collectively rather than privileging one account.

The ways in which debunking accounts explain moral phenomena vary, but this explanatory power is what makes a debunking account a debunking account, and what enables them to play the role that they do within the Argument from Queerness. This is because simply by offering such an explanation for moral phenomena—by being a successful debunking account—they provide compelling evidence for the claims that there is another potential explanation for moral phenomena besides the existence of moral facts. This alone is sufficient to show, at the very least, that the question of the best explanation of moral phenomena is not settled in favour of moral realism. It also means that, though I have referred to debunking accounts as one of two subordinate arguments underlying the Argument from Queerness, they are not in and of themselves arguments in the strictest sense.

So what makes a debunking account successful? At minimum it must give us a plausible explanation for moral phenomena without invoking anything we do not already accept into our ontology. A very simple example might go something like this: a person thinks that some action is right or wrong because that is what their parents taught them when they were young.²⁸ We already accept the existence of parents, and that they influence their children, and we can see how that influence is sufficient to explain some beliefs including, plausibly, beliefs about morality. Moral facts play no part in this explanation and so need not be posited to explain this belief. There is a drive towards ontological minimalism such that debunking explanations of moral phenomena will usually be naturalistic explanations, but one needn't be committed to metaphysical naturalism to be persuaded by them; that they are more parsimonious than the alternative is enough.

²⁸This assumes that the content of morality is teachable, but this should be so just so long as it is propositional, which is entailed by cognitivism (to which Moral Error Theorists and moral realists are both committed). Even if some form of moral particularism is correct, and there are no generally applicable moral principles to be taught, simple debunking explanations can still explain some actual moral beliefs, as there is nothing philosophically mysterious about someone being 'taught' falsehoods.

The best debunking accounts will also be considerably more general than the example above. They will be *global* in the sense that they aim to explain moral phenomena as a whole, and in so doing will give a template for explaining individual moral beliefs, feelings, or actions. Ideally, this means explaining how moral phenomena originated and how they have persisted, as well as why our moral thought has the particular structure that it does. There have been a number of attempts at giving such systematic debunking accounts, and I will discuss some of these very briefly, but the amount of overlap between them is so extensive that it would be wrong to think of them as competing views, so I will not be placing them side by side for comparison. Nor will I be providing a debunking account of my own. Instead, I will try to present a characterisation of an archetypical debunking account that incorporates all the typical features thereof, so that it is clear what exactly debunking accounts achieve and what, if anything, we are granting when we admit that one is plausible. Thereafter, I will treat it as given that such an account is available, since I am trying to grant Moral Error Theorists as much as possible, and the one premiss that I mean to dispute doesn't feature in this side of the argument.

Debunking accounts generally explain moral phenomena with reference to their usefulness. Their usefulness, on such accounts, lies in the fact that it tips the scale in favour of altruistic behaviour and cooperation, and allows humans to flourish in a social environment that mixes collaboration and competition.²⁹ We would all be worse off outside of the context of a working society, and moral phenomena are a precondition for such a society's existence; cooperation allows humans to more easily meet many of their needs, but also often requires us to make sacrifices. The moral phenomena, then, are a set of behaviours, and a system of related emotional, conative, and cognitive states (that is, feelings, drives—inasmuch as these are distinct from feelings—and beliefs) that reinforce these behaviours. The behaviours I mean include paradigmatically moral behaviours—forbearing from theft and murder, helping others at a cost to oneself, and so on—as well as social behaviours that reinforce the whole system, such as the moral education of

²⁹Mackie 1977, pp. 107-15, 171 and Blackburn 1993, p. 168.

children, and praising certain actions and character traits.³⁰ No beliefs about the goodness or rightness of actions must be true, and none of the perceptions that lead to them must be veridical, for such these behaviours to be beneficial or for the whole system of phenomena to fulfil its supposed purpose.

On such a view, one of the functions of morality is to enable us to live in a cooperative social setting by counteracting traits that would otherwise stand in the way of (ultimately adaptive) moral behaviour, such as limited sympathies or weakness of the will.³¹ We couldn't reap the benefits of society if we favoured ourselves and our loved ones exclusively, or if we routinely flouted the conventions that make it possible.³² Next, we must wonder how it serves this function. Richard Joyce suggests that categorical imperatives are central to moral thought, and so to morality as a human practice.³³ A categorical imperative is an imperative that doesn't depend on an agent's particular ends. So if there are categorical imperatives, then there are inescapable reasons that apply to all of us regardless of our desires.³⁴ Joyce's claim is plausible enough; moral debates among ethicists and laypeople alike seem to be predicated on the inescapability (and, sometimes,

³⁰To be clear, all of this is true of the structure rather than the content of morality. It is entirely possible that the content that is 'plugged in' to this system in the case of a particular individual will be things we are disinclined to call good or right. For example, an agent motivated by virulent nationalism might still be thinking thoughts and feeling emotions that are characteristically moral in the relevant way.

³¹These examples are drawn from Mackie 1977, p. 107 and Joyce 2006, pp. 108-18, 208 respectively.

³²By 'we' here, I am referring to human beings generally. While individuals obviously do bend and break these rules, it is nevertheless the case that for a society to function in the desired way, most of its members must observe these principles most of the time.

³³Joyce 2001, p. 157 and 2006, pp. 60-1.

³⁴Joyce is not claiming that all of our moral judgements—much less every aspect of our morality—has this structure. Rather, his claim is that human morality necessarily involves judgements of this structure, but it needn't consist solely of such judgements. This will come up again in Chapters 2 and 4.

the overridingness) of moral reasons.³⁵ Making moral judgements of this form would counteract traits that might otherwise discourage moral behaviour.

So our debunking account of moral phenomena explains their persistence by the desirable outcomes they produce, and implicitly by the pernicious outcomes that would follow if we were to cease our moral practices. However, we might still be unsatisfied because we may find ourselves with an issue of regress. That is, we have an explanation of how humans are such that one generation that experiences and participates in moral phenomena can pass it onto the next, and of how they do so. What we do not yet have, though, is an account of how this cycle would begin. Those who suggest debunking explanations seem to agree that the answer lies in our evolutionary history.³⁶ As long as a case can be made that the thoughts and behaviours that make up morality are adaptive, as I have done above, then we can explain them through natural selection. One particularly salient detail for the debunker is that the utility of moral phenomena alone is sufficient for them to be selected for and so, as far as evolutionary processes are concerned, the truth or falsity of any beliefs included therein is a matter of indifference.³⁷ The story given above is vague but *prima facie* plausible, especially given the plausibility of evolutionary histories for our other psychological and physiological traits.

But it is also too simple. The explanation above appeals to what is in our interest now, but an actual evolutionary history must tell us why it was selected for when it evolved. Traits are selected when they are adaptive, and what makes a trait adaptive is that it increases an individual's reproductive fitness, which is to say it makes that individual more able to pass their genes onto the next generation.

³⁵We must bear in mind that what is under discussion here is the structure of morality as a feature of human thought and behaviour, not of moral facts. I will discuss the idea that moral facts have a categorically imperative structure, or something like it in §1.8 and Chapter 2.

³⁶See Mackie 1977, pp. 113-4; Blackburn 1993, pp. 168-9; Street 2006, pp. 113-21; and Joyce 2006, pp. 16-7 for examples.

³⁷This means that even if realists offer an account of moral facts such that they turn on what is useful in the way described above (for examples, see Gauthier 1986 and Scanlon 1998, §4.3), such claims would still be suspect in the eyes of the debunker insofar as they had been arrived at by a process which is not truth-tracking with regards to moral facts, namely natural selection.

An individual's wellbeing is only relevant to evolution insofar as it advances this adaptivity.³⁸ So it is beside the point whether an evolved trait serves our current purposes; what we really need is an account of why the elements of morality were adaptive when they were selected for.³⁹

Joyce offers an attempt at such an account, though his picture is too extensive to recount here. One thing that becomes apparent in it, and which is worth incorporating into our general debunking account, is that morality is a complex phenomenon that is likely to have arisen from multiple sources as diverse as prosocial emotions (such as love and friendship), reciprocity, and the 'co-opting' of non-moral motivational mechanisms such as disgust as the foundation for some moral emotions.⁴⁰ I will treat this as given going forward, but it is worth remembering that it does not matter for the purposes of the Argument from Queerness which particular evolutionary debunking account is true, so long as one is.⁴¹ This assumes that natural selection is not truth-tracking with regard to morality, but I have presented my reasons for thinking as much above.⁴²

This kind of origin story for morality is scientifically plausible, and it also has a couple of other advantages. It gives us a unified explanation of the many types

³⁸Joyce 2006, p. 15. Passing one's genes on need not be done directly by reproducing; protecting kin also increases the chance that one's genes will be passed along.

³⁹Joyce 2006, p. 107.

⁴⁰Joyce 2006, pp. 19-40, 140-1. Reciprocity is a particularly interesting example because we can see how a creature that is prone to rationalisation, as humans are, might idealise an inclination to treat others as those others have treated them as concepts like justice or moral desert.

⁴¹The Argument from Queerness can make the same use of a debunking account that is not grounded in our evolutionary history, but since evolutionary debunking arguments are the most widely endorsed as well as sufficiently global and plausible for our purposes, I will limit my discussion to them.

⁴²The observation that evolutionary debunking accounts only have these implications if natural selection is not truth-tracking with regard to morality comes from Kahane 2011, pp. 111-2. For defences of the claim that it is not truth-tracking in this way, see Joyce 2006, ch. 6 and Street 2006, §6.

of mental and behavioural phenomena that make up morality, as well the many different kinds of moral value people recognise, by telling a story about how they increase our reproductive fitness.⁴³ At the same time, it explains why there is a certain degree of agreement about the content and structure of morality, insofar as certain behaviours like ensuring the survival of oneself and one's kin, for example, are always adaptive.

Such accounts are not without limitations, of course. They are by their very nature speculative and subject to at least some degree of doubt. They are also incomplete; to fully justify the claim that moral phenomena can be explained without reference to moral facts, we also have to rule out the possibility of moral beliefs that actually are best explained by moral facts.⁴⁴ This is more likely to be achievable-in-principle than not, though, since even the most self-consistent and reflectively refined system of morality will have some bedrock intuitions for which we can provide a plausible evolutionary origin story.⁴⁵ Not to mention, any successful queerness argument would serve as evidence for this as well. And while debunking accounts are speculative, we need only one plausible candidate to provide a live alternative to realist explanations of moral phenomena and so serve the relevant role within the Argument from Queerness.

The role of a debunking account in the Argument from Queerness is to provide an alternate way of explaining moral phenomena that can take up the mantle of 'best explanation' once a queerness argument has shown moral realism to be problematic. But we might wonder at this point whether moral facts need to figure in our best explanation of moral phenomena for us to posit them. Christine Korsgaard points out explaining and describing phenomena is neither the primary occupation of humans, nor what we are typically doing when we make substantive claims about the reasons for and value of actions.⁴⁶ Instead we are trying to answer questions about how we should act, and she rightly observes that when we find answers to those questions, we find moral facts. If we take this as a straightforward

⁴³Street 2006, p. 134.

⁴⁴Kahane 2011, p. 119.

⁴⁵Kahane 2011, pp. 119-20. See also Mackie 1977, p. 41.

⁴⁶Korsgaard 1996, pp. 45-7.

argument for moral realism, it will be circular, since in making such an argument one would have to assume that these answers are available, which is to say that there are moral facts to be discovered.

If we instead take it as a response to arguments from the best explanation, then it still does not work against the kind of debunking account described above. The issue she identifies seems to be that the debunker is, as it were, looking for moral facts in the wrong place when they expect to see them in the explanations of moral phenomena. Where we should expect moral facts to present themselves is as answers to questions of how we should act, what we should value, and so on.⁴⁷ A response like this might work against debunking arguments for individual beliefs, but the kind of account that would effectively jeopardize all moral phenomena as evidence of moral facts does so by explaining away our tendency to employ moral concepts in general. If we can tell a story about not just the answers to moral questions but our very capacity to ask moral question, and that story does not require us to posit moral facts, we have no reason to suppose that those concepts have extensions or that satisfying answers are available to those questions, even if asking them is an essential feature of the kind of being we are.⁴⁸ This is not proof of the absence of moral facts, of course, but arguments from the best explanation are not in the business of offering definitive proof in any case.

These responses do not address the more fundamental point raised, which is that there might be ways to justify ontological commitments beyond explanatory power or indispensability. David Enoch offers an argument which derives from the same concern as Korsgaard's, but which offers a more complete explanation of how this might be so. He suggests that some facts are deliberately indispensable (in the sense that we are committed to them when we ask or try to determine what we

⁴⁷This argument leaves aside how we come to those answers. Appeal to moral phenomena would be circular, but it should suffice to say that the answers can be found at the end of some sort of appropriate chain of deliberation. If such a chain cannot be satisfactorily characterised, this will be a separate issue for this argument.

⁴⁸The fact that the best explanation (given queerness arguments) of our tendency to employ these concepts does not presuppose the existence of any moral facts places it in marked contrast to the analogous explanation for most of our other beliefs, discussed in more depth in §4.3.1.

should do), namely irreducibly normative facts (see §1.8).⁴⁹ According to Enoch, both explanation of the world and deliberation over how we ought to act are intrinsically indispensable—meaning that insofar as we are agents we cannot rationally give up on either—and so both can ground ontological commitments.⁵⁰ Putting aside questions of what deliberation actually commits us to, we can respond to Enoch’s argument by rejecting the claim that deliberative indispensability is sufficient to justify ontological commitments. As he rightly points out, though, this rejection must be backed up with reference to some non-arbitrary way of discriminating between kinds of indispensability.⁵¹ As Brendan Cline points out, only true posits (which is to say facts) can have explanatory power and to take a putative explanation as the actual (or true) explanation of some phenomenon is to take its posits to be true.⁵² Our best explanations are the ones that stand most chance of being actual explanations, and so those we have most reason to commit ourselves to (by the nature of the explanatory project). This will involve committing ourselves to those facts that they posit, and thus to the presence of certain entities and properties in our ontology. Deliberative commitments do not seem to close ontological questions in the same way. Consider free will; there is a plausible case that in deliberating we are committed to the idea that our options are open and that the decision over what to do is ultimately ours, yet the possibility that we are not free is no less open for it.⁵³ Commitment on either side is not what settles such questions, and the commitments in question arise from the phenomenology of deliberation (which feels like the search for genuine reasons for action, hence

⁴⁹Enoch’s argument is not directly concerned with moral facts, but moral deliberation is as respectable as any other kind, so while this argument doesn’t strictly aim to prove that there are moral facts among these irreducibly normative facts, its conclusion does create an ontological space for them, which makes it worth addressing here.

⁵⁰Enoch 2011, pp. 69-75.

⁵¹Enoch 2011, p. 55.

⁵²Cline 2016, p. 3234. This assumes scientific realism in the sense of taking scientific claims to be intended as statements of fact, but not more so than any of the other arguments discussed up to this point.

⁵³Cline 2016, p. 3245-6. For a brief case against the existence of free will, see §A.2.

the supposed indispensability of certain normative facts) so they are, like perceptions, preliminary data that we can subject to scrutiny.⁵⁴ We might even give an evolutionary account along the same lines as our debunking account.

One more difference between the forms of indispensability is that explanation by nature deals with facts. We explain one fact (the explanandum) by positing other facts, whereas deliberation starts not with a posit or a phenomenon but a question: ‘What should I do?’ This presupposes that there is an answer, but—as I have already noted—the fact that the question makes sense to us doesn’t entail that any satisfying answer exists. By contrast, the existence of any non-fundamental natural fact implies the existence of other facts that explain it. What Enoch’s argument actually proves (assuming we accept all of its assumptions) is that either he is correct and there are irreducibly normative facts, or that deliberation involves categorical error.⁵⁵ I do not mean to commit myself to Enoch’s claim that irreducibly normative reasons are deliberatively indispensable (in fact I am inclined to reject it), but whether or not this is so, we can conclude from the considerations above that deliberative indispensability is not a good guide for ontological commitments, which is sufficient to rebut this kind of argument.

To summarize, a successful debunking account will explain the origin and persistence of morality as a phenomenon in the world. In practice, we can expect it to be a general account that seeks to explain moral phenomena as the product of evolution by natural selection.⁵⁶ While these traits are not necessary for a debunking explanation, most take this form, and I have already stated that the basic idea is fairly plausible, and persuasive accounts that elaborate on the idea more fully have been given. What matters most for the purposes of the Argument from Queerness, though, is that we can give a debunking argument. If none can be given, all we have are sheer queerness arguments of the kind already described,

⁵⁴Cline 2016, p. 3246-8.

⁵⁵Cline 2016, p. 3250.

⁵⁶We may also see room made for sexual selection, where a trait is selected for because it is preferred by breeding partners for reasons that may not relate to its adaptivity, and cultural selection, where sets of ideas or practices can prove to be adaptive and so be passed down in the same way as genes.

which are not very persuasive, so without a debunking account the Argument from Queerness cannot get off the ground. To be clear, a queerness argument is the crux of the Argument from Queerness—a debunking account alone has even less dialectical power than a sheer queerness argument—but the queerness of moral facts will only count against moral realism if there is some alternative it can lead us to prefer. For the rest of this project I will assume that we can give such an account, and that it takes roughly the form that I have described above.

1.4 Queerness Arguments

We have seen that debunking accounts of moral phenomena give us reason to think that moral facts are redundant in explanations of these phenomena. The Argument from Queerness effectively takes the form of a two-pronged attack on the idea that moral facts are the best explanation of moral phenomena, and a debunking account forms one of its prongs by demonstrating that there is another explanation which is superior in at minimum one way or, if one is unmoved by appeals to ontological parsimony, is at least a legitimate rival.⁵⁷ The other prong, then, is the claim that moral facts are queer, which must be supported by a queerness argument.

Talk of explanatory redundancy will be fairly familiar to most philosophers, but it's much less clear what it means to claim that moral facts are 'queer'. Even committed Moral Error Theorists admit as much.⁵⁸ However, we don't need to know precisely what queerness is to understand the structure of queerness arguments, so I will explain said structure before moving onto the questions of what 'queerness' is and why it is a problem for moral realism.

Olson, who has done a lot of work in delineating and assessing the existing queerness arguments (which he draws from Mackie's presentation of the Argument from Queerness, on which modern Moral Error Theorists tend to rely), states that

⁵⁷Unlike other appeals made to ontological parsimony in this chapter, here we are employing the quantitative version of Ockham's razor; the realist model is worse by virtue of the fact that it posits more facts in general than the antirealist one. Such appeals have some power, but are less persuasive than the appeals to the qualitative version of the razor.

⁵⁸Olson 2014, p. 86.

each queerness argument nominates some feature of moral facts as the source of their queerness.⁵⁹ He lays each out separately, but they all ultimately share a structure that looks something like this:

(Q1) Some feature X is sufficient to make anything that possesses it queer.

(Q2) X is an essential feature of moral facts.

(Q3) Therefore, moral facts are queer.

Unlike the Argument from Queerness itself, this argument is meant to be deductively valid, and in fact it is. Hence, anyone who wants to counter any queerness argument must undermine (at least) one of its premisses.

In short, there are two questions that are indispensable in assessing any queerness argument. The first is: ‘Is X really sufficient for queerness?’ or ‘Does having X make something queer?’. The second is: ‘Is X really an essential feature of moral facts?’. If the answer to either is ‘no’, the queerness argument fails, and so will any version of the Argument from Queerness that depends on it.

Whatever feature we substitute in for X, (Q2) will be a conceptual claim about moral facts. It is the claim that all moral facts invariably have the feature in question. As I have mentioned, there are some assumptions about the nature of moral facts baked into the Argument from Queerness, and in some cases these can tell us whether X really is an essential feature of moral facts, but in some cases they cannot. It all depends on what X actually is. However, there is an underlying logic to these commitments, and it should be more generally applicable. Recall that moral facts are those facts that ordinary speakers express, or at least try to express, when they use moral language. This gives us a fairly straightforward criterion for determining whether moral facts necessarily have X: does the use of moral terms reflect this? Joyce suggests that we can answer this question via a translation test: if we can determine what about a term is non-negotiable and would have to be preserved in translating it to another language, then we will have determined what that term means.⁶⁰ Likewise, if we can use moral language in a

⁵⁹Olson 2014, p. 87.

⁶⁰Joyce 2001, p. 5.

way that does not imply that moral facts have X, then it seems that X is not an essential feature of moral facts. (This ordinary usage also tells us what counts as a moral fact, or else what the conceptual and structural boundaries of the moral domain—which I take it we can grasp pretheoretically—are.)

(Q1) is the claim that X makes its possessor queer, so in order to assess it we must finally pin down what it is for something to be queer. Ideally, this would amount to giving necessary or sufficient conditions for queerness, but even if this cannot be done I hope to give as clear an account as possible, so that we can more easily assess (Q1) for any given X. As I have already mentioned, Moral Error Theorists themselves have a hard time fully characterizing the notion of queerness. This ambiguity can make the Argument from Queerness appear weaker than it is; if the issue is with queerness and it's unclear precisely what queerness is, then it's unclear what the issue is. This is problematic not only for those who wish to defend the argument but also for those who want to attack it, though the move may still uncharitably be seen as avoiding criticism by glossing over vital details. As such, many previous accounts of what queerness amounts to come from critics who see queerness as precisely this kind of obfuscation (and I will argue below that many such readings are flawed) but they do draw on the remarks of Moral Error Theorists as the source of these interpretations, and this is where we should start as well.

As mentioned in §1.2, 'queer' is used in the sense of 'strange'. It is a term of art in metaethics, where it refers specifically to the way in which moral facts are allegedly strange and problematic.⁶¹ Mackie said that if moral facts did exist, they

⁶¹In modern parlance, 'queer' broadly denotes those who are nonconformist with regard to gender or sexuality. This usage began as derogatory, and debates about whether and when it still is so are ongoing. While these discussions clearly have a moral dimension they are nevertheless irrelevant to the metaethical investigation I am carrying out here.

What is of some interest is that the pejorative use of the word 'queer' predates Mackie's use of it significantly—see Foldy 1997, pp. 22-3 and Robb 2005, p. 262. It's also worthy of note that 'queer' has often been used to describe what is not merely unusual, but what is unusual in a way that suggests something is amiss or suspicious—see OED2, s.v. 'queer', sense 1. What this suggests is that howsoever moral facts are strange, we are meant to see

would have to be ‘of a very strange sort, utterly different from anything else in the universe.’⁶² So, on his view, at least part of what it is to be queer is relational; for a thing to be strange or unusual, it must be somehow different from other things.⁶³ A synonym that might get the idea across particularly well would be ‘peculiar’, since features can be peculiar in the general sense of being unusual, or in the sense of being peculiar *to* some limited class of entities. This reading of queerness as strangeness is the one Platts relies on in his objection to Mackie.⁶⁴ But I have already explained (in §1.2) how sheer queerness doesn’t necessarily raise ontological suspicions, so understanding queerness as nothing more than strangeness would not lead to a compelling version of the Argument from Queerness.

Platts’s examples fail to raise ontological suspicions because they can all be accounted for parsimoniously by fitting them into our pre-existing understanding of the world, which is to say by a process of naturalization. Perhaps queer entities are queer in virtue of the fact that they cannot be integrated in this way—remember that we have assumed along with Moral Error Theorists that moral facts are non-natural facts, so it would be contradictory to claim that they could be naturalized. On this understanding, queerness amounts to nothing more than non-natural status. This leads us to the understanding of the Argument from Queerness favoured by Jean Hampton, who takes the Argument from Queerness to be rooted in metaphysical naturalism, which is the view that natural facts are exhaustive and so there are no non-natural (or, for that matter, supernatural) facts.⁶⁵ Naturalism of this stripe is a view shared by Mackie and many other Moral Error Theorists,

that strangeness as a problem. This, as we shall see, gels very much with the way the term is employed within the Argument from Queerness.

⁶²Mackie 1977, p. 38.

⁶³In the case of metaphysics and ontology we can settle the question of what is strange purely on the basis of quantity or frequency, but there is obviously more to the question in other contexts. In social contexts, for example, what is normal (a term which tellingly contains the word ‘norm’) is at least partially constructed, shaping and being shaped by the expectations of members of a given society.

⁶⁴Platts (1980, p. 72).

⁶⁵Hampton (1998, pp. 21-2).

but this doesn't imply that it is one of the assumptions of the Argument from Queerness.⁶⁶ In fact, one may have all kinds of non-natural commitments and still see moral facts as queer.⁶⁷ Likewise, if a thing can be natural yet also queer by some other token, there must be more to queerness. Non-natural status might also plausibly be the feature nominated in a queerness argument. Just such a line of thought is present in Mackie's account.⁶⁸ But in nominating some feature as the subject of a queerness argument, one is suggesting only that it is sufficient for queerness, not that it is identical with queerness itself (any feature identical with queerness would need to be both necessary and sufficient for queerness).

Perhaps the most sophisticated attempt at isolating queerness comes from Lee Shepski, who draws several plausible meanings for the term out of Mackie's scant exposition. The first is the idea that an entity is queer if it is ontologically profligate, which is to say it is queer if it is not posited by our most ontologically parsimonious account of the world. If this account were correct, the Argument from Queerness wouldn't be a distinct argument unto itself but a simple appeal to ontological parsimony, and on these grounds Shepski rejects this interpretation. This last move is somewhat questionable, however. It is grounded in the assumption that the Argument from Queerness must be a novel argument in order to hold any rhetorical power or philosophical interest, but originality is not what determines how well an argument supports its conclusion, and in any case, wherever Mackie advanced the Argument from Queerness, he took himself to be expressing 'traditional' arguments against moral realism.⁶⁹ So we should not dismiss this reading of queerness out of hand, as Shepski does. I have already said that I take the Argument from Queerness to be an argument from the best explanation, and primarily motivated by a concern for ontological parsimony, though this is with

⁶⁶Insofar as the Argument from Queerness is meant to eliminate one category of non-natural fact, it can be seen as supporting metaphysical naturalism, rather than the reverse.

⁶⁷Olson 2014, p. 86 and 2017, pp.64-5.

⁶⁸Shepski (2008, p. 379), citing Mackie (1977, p. 39). If we are to say that non-natural status is *sufficient* for queerness without being identical to it, and without assuming general naturalism, I think we must give an argument like those explored in §1.7.

⁶⁹Mackie 1977, p. 35 and 1946, p. 77.

regard to kinds of entities rather than specific individuals, so more must yet be said to bring out its force. Otherwise, a debunking account may be enough on its own to motivate the rejection of moral facts, and we would not need to invoke queerness at all.

Shepski's second reading is queerness as mystery; to be queer is to resist explanation or comprehension. This hews fairly close to the ordinary meaning of the term, and reflects the remarks of modern Moral Error Theorists as well.⁷⁰ But it is also ambiguous between the different ways in which things can be inexplicable or incomprehensible, and none of these seem to yield a satisfactory account of the Argument from Queerness. Mackie explicitly rejects the idea that moral claims are outright nonsensical or incoherent, so this cannot be the source of their incomprehensibility.⁷¹ But if the inexplicability in question is just a matter of being beyond human understanding (as a matter of fact or in principle), this needn't have any ontological implications.⁷² Many theists, for example, believe both that God exists and that God is beyond human ability to fully comprehend. So, like queerness as simple strangeness, this reading should be rejected because it would not lead to a compelling version of the Argument from Queerness. One final way in which something could be inexplicable is if it is ontologically basic, which is just elliptical for Shepski's third interpretation.

The last reading of queerness that Shepski puts forth is that to be queer is simply to be *sui generis*; an entity of in a category all its own. But there are subtleties to the application of this term, regarding which philosophers have not been entirely consistent. There is one reading on which moral facts cannot be *sui generis* for the simple reason that moral facts are not a category *all their own* because they fit into the wider category called 'facts'.⁷³ The same can be

⁷⁰Olson (2014, p. 86) speaks of queer entities as 'giving rise to puzzlement'.

⁷¹Mackie 1977, pp. 21-1, 39-40 and Joyce 2001, ch. 2-3.

⁷²Shepski 2008, p. 376.

⁷³I owe this observation to Tom Stoneham. One might object to this reading that everything that exists belongs to the category of 'existent things', meaning that trivially nothing is *sui generis*, or that each thing belongs to its own category of things that have all and only exactly the properties possessed by that thing, meaning that trivially everything is *sui generis*. As I

said of moral properties in relation to properties in general. But we can determine Shepski's intended meaning through closer reading of his account of this issue. This reveals that what he actually takes Mackie (and Moral Error Theorists generally) to find problematic about queer entities is the sheer degree of difference between moral facts and anything else in our ontology.⁷⁴ Regardless of whether moral facts are strictly *sui generis*, they are not simply more of something with which we are already familiar. Rather, they are so different that admitting them to our ontology would require us to admit properties of a wholly unique kind. This unique kind would amount to a fundamental addition to our ontology, so what Shepski is really discussing is a reading of queerness as fundamental additivity.⁷⁵ To be clear, the fundamental addition that we are concerned with is not the category of the moral *per se*, but a property that belongs to moral facts or to moral properties, which we need not posit anywhere else in our account of the world. This is the feature nominated as the source of queerness by individual queerness arguments.

Shepski takes it that the reading of queerness as fundamental additivity is the best of the readings he examines, if only because it leads to an interpretation of the Argument from Queerness that is, by his lights, an original argument for moral antirealism, albeit one that fails. I too think this is the best reading of queerness, because I think it leads to a successful version of the Argument from Queerness. Before I explain why, however, we must respond to Shepski's contention that a version of the Argument from Queerness rooted in this reading would fail.

This conclusion is rooted in his general understanding of the Argument from Queerness as a whole, rather than the the ontological significance of status as a fundamental addition to our ontology.⁷⁶ Shepski's account of the argument is as

explain below, however, nothing I go on to say will turn on this reading, so I will not engage with this line of argument any further.

⁷⁴Shepski 2008, p. 377.

⁷⁵We also see this interpretation of queerness in Bedke 2009, p. 46.

⁷⁶Though he does not call them such, Shepski (2008, pp. 378-85) does give a number of queerness arguments nominating features in virtue of which moral facts might be *sui generis*, though by merit of what I say below I do not think he does them justice. I treat the candidates he raises in §1.5, §1.7, and §1.8.

follows:

- (S1) Some feature X is an essential feature of moral facts.
- (S2) Nothing has or can have feature X .
- (S3) Therefore, there are no moral facts.

Here, X is the feature in virtue of which moral facts are queer, the one that means they must be a fundamental addition to our ontology. It is a very simple deductive argument, so realism will be in trouble if the premises turn out to be true. But there are a number of ways in which this argument does not work.

First of all, nothing about this argument turns on the understanding of queerness as fundamental additivity. As written, this argument reads more as one where queerness is impossibility (yet another distinct reading). If S2 is true, then X will be sufficient to make its possessors fundamentally additive, but we do not need to make that inference in order to reach the conclusion of this argument. In fact, neither the term ‘queer’, nor the term ‘fundamental addition’ appears anywhere in this argument. Secondly, and perhaps more importantly, the argument *does* turn on S2, and S2 is beyond proof. Our failure to find anything possessed of X might be inductive evidence for S2, but this is far from strong enough to warrant the level of commitment we’d need to employ it as a premise here (not to mention, it will often be controversial whether we can find things possessed of X). We cannot treat S2 as a premise, because surely S2 is precisely the kind of claim the Argument from Queerness must substantiate, with the further step to S3—Moral Error Theory—being a mere formality.

Further, there are X s that will make this argument entirely circular. If the property nominated as X is definitional of moral facts—both necessary and sufficient to be one—then the claim that nothing can have X will just be the claim that there are and can be no moral facts.

In short, Shepski is expecting the fact that moral facts have a queer-making feature (and not, as I have noted, the fact that they are queer) to do far too much work in the argument. He takes it that X , whatever it is, must be enough to rule out moral facts all on its own. Queerness arguments should be sufficient to raise ontological suspicions, but they only warrant the conclusion that Moral

Error Theorists are chasing when taken in conjunction with a debunking account. The notion of queerness was only ever meant to do half a job, and the Shepskian account falls short because it fails to take this into consideration.

In fact, attending to that role and understanding the work that the notion of queerness does may be the best way to clarify it. First and most obviously, queerness alone is not grounds to reject something from our ontology. We have already admitted as much, and looking at the structure of the Argument from Queerness we can see that it doesn't need to be. Queerness is not meant to warrant much on its own; only a queerness argument and a debunking account together are meant to motivate Moral Error Theory. Each forms one half of a two-pronged attack on the idea that moral facts are the best explanation of moral phenomena.

A debunking account is an attack 'from the front' in that it presents a legitimate rival for the title of 'best explanation'. A queerness argument is an attack 'from the back' in that it undermines the idea that moral facts are a good explanation for moral phenomena in the first place. In order to see how it does so, we need to understand what separates a good explanation from a bad one.

An explanation is good insofar as it discharges an explanatory burden, which is to say that a good explanation answers questions about how or why something happens the way it does.⁷⁷ An explanation is bad, therefore, when it fails to discharge this burden or—worse—when it increases the burden.

An explanation will also be bad if invoking it worsens other explanations. For example, an explanation will be bad if it undermines the generality of an otherwise sufficient general account. Suppose I experience some eerie and uncanny phenomena in the dead of night, the activities of local pests and my own imagination—or even the activities of a large special effects team who have taken it upon themselves to deceive me—are better explanations of this than the actual presence of a ghost in my home. The existence of a ghost with all the archetypical spectral properties is incompatible with huge swathes of physics as we understand it, so

⁷⁷There is a looser sense of explanation that includes explaining what a thing is or does, but this might be more rightly called description or characterisation, and is not the sense that interests us when arguing from the best explanation.

by accepting an explanation that posited one, we would undermine our explanations of basically every other observable fact. In worsening other explanations, an explanation lowers or eliminates their ability to shoulder whatever explanatory burdens they were posited to deal with, so this too is a form of creating, rather than discharging, explanatory burdens. So there are two ways an explanation can be worse. And, as we shall see, both ways will be relevant for the Argument from Queerness.

This is the hurdle that an account of queerness must overcome: whatsoever is sufficient for queerness on the part of moral facts must also be sufficient in some way to undermine the power of moral facts to discharge the explanatory burden imposed by the existence of moral phenomena. Strange entities need not raise any burdens that cannot be discharged in the usual way, as we have already seen, and non-natural status will only do so against a background assumption of general metaphysical naturalism, which is why these readings are not congenial to the Argument from Queerness. Such entities can raise explanatory questions, but they can also fail to do so. This difference in the modal qualities—that queer entities will always saddle us with an explanatory burden, whereas strange or non-natural entities can do so but do not do so necessarily—implies that they are not one and the same property. Whether strange or non-natural entities do prompt such questions depends on whether we need to revise our ontology to account for them.

We cannot so quickly dismiss the other accounts we have examined. Shepski is right about the kind of features Mackie draws our attention to, but he makes two moves that rob the Argument from Queerness of its power. First, as already discussed, he expects X to do too much work. Second, he isolates the possible readings of queerness from one another and treats each separately. I contend that moral facts are mysterious and ontologically profligate precisely because they are a fundamental addition to our ontology, and this is why they impose the explanatory burdens that they do. But this will lead to a working argument only if being a fundamental addition actually creates the kind of burdens I have just claimed it does in the case of moral facts.

Any and all fundamental additions to our ontology are ontological primitives

that are not fit to be reduced to anything we already accept.⁷⁸ Such additions (as mentioned earlier) aren't necessarily a problem, but it should be clear that they can be. The fact that we ever appeal to Ockham's razor, or ontological parsimony more generally attests to this. But when is adding to our ontology a problem, and why is it so?

Shepski's second reading of queerness was Queerness as Mystery. Queer entities are generally understood to defy comprehension, or to be inexplicable. Ontological primitives are also inexplicable; by definition they brook no further analysis. Each ontological primitive we posit carries a cost in the form of truly intractable ontological questions so, as long explanations are in the business of eliminating questions, the fewer primitives we invoke the better we are doing.⁷⁹ This is part of what makes Ockham's razor a good principle of reasoning.

It follows that, if moral facts are fundamental additions to our ontology, then this alone is sufficient to decrease their explanatory power because they bring an additional mystery with them. Moral facts are mysterious and ontologically troubling because they are fundamental additions. They are shown to be ontologically profligate when a debunking account is added to the mix, since the debunking account shows that we are not forced to posit them to account for moral phenomena. In light of all this, moral realism is the worse explanation insofar it does a worse job of discharging explanatory burdens than the alternative, which it does by positing a fundamentally additive category of fact where the debunking account does not. Moral realism also suffers from the second potential flaw in explanations discussed above in that, by positing a fundamentally additive kind within the broader kind 'fact', they undermine the explanatory power of our otherwise unified account of what facts are and can be like. In these ways, queer entities can be said to raise the explanatory cost of the entire world. As such we should not commit ourselves

⁷⁸This reading also has fits well with the ordinary language reading of queerness; entities that amount to a fundamental addition to our ontology do so precisely because they are unlike other things, which is all it takes to be strange in the everyday sense.

⁷⁹The specific nature of these questions is ambiguous, but could be as simple as 'Why does this exist?' in the sense of interrogating either the origin or the specific nature of the thing in question.

to moral realism, and we should reject from our ontology anything it uniquely posits, namely moral facts.

There are some potential lines of objection to understanding queerness as fundamental additivity. One is to accuse the Error Theorist of misapplying Ockham's razor; that moral facts require a fundamental addition to our ontology is a problem only if we make that addition beyond necessity. Implicit in this is the claim that moral facts really are (or at least could for all we know be) explanatorily necessary, but this claim cannot be made if one also has access to a successful debunking account of moral phenomena, so there is really nothing to this argument. Another objection has it that the multiplication of entities has already happened insofar as we are discussing moral facts, and that Moral Error Theorists share the ontological commitments I have claimed are problematic with moral realists. I see no reason for simply talking about properties or kinds to have any ontological cost. We can posit as many properties as we like—the properties of being a dragon or being a ghost, for example—without issue. It is only once we insist that there is room for a given property in the actual world that we must pay the price for them. If we tried to insist that something actually was a dragon or a ghost, if we posited facts about the distribution of these properties, then and only then would we be on the hook for reconciling them with our existing explanations to develop a coherent picture of the world.

Some moral Error Theorists have urged a reading of queerness as something more exotic and pressing. Olson, for example, contends that the feeling of puzzlement that queer entities engender goes beyond the basic mystery or inexplicability of ontological primitives.⁸⁰ What he does not say is how anything could be more puzzling than that which is fundamentally and definitionally inexplicable, nor what it would even mean to be so. Richard Garner suggests that moral facts are 'unusual in an unusual way'.⁸¹ This is intended to highlight the strangeness of queer entities compared to less contentious ones, and so, I take it, to point to their status as fundamental additions. But we could pursue this even further to say that moral facts are somehow more different, standing apart from other ontologically

⁸⁰Olson 2014, p. 86.

⁸¹Garner 1990, p. 143.

basic entities in an even deeper way than those entities stand apart from each other. (After all, we already accept at least some ontologically basic entities, and so we might seek the best possible reasons to see these ones in particular as problematic.) We are then, however, faced with the question of what it would mean for one such class of entities to be more or less different from another—surely each ontological primitive is equally basic and unique. Claims like these, if they could be substantiated, might well make the Argument from Queerness more compelling, but reading queerness simply as fundamental additivity is still quite sufficient to motivate the Moral Error Theoretical conclusion, and has the advantage of not requiring us to commit ourselves to or argue for such troublesome notions.

Others still might insist that some kinds of property simply aren't fit to be admitted to our ontology. This I must grant—the Argument from Queerness concludes with such a claim—but if we take that as a premise we will effectively rehash Shepski's account and fall afoul of all the same criticisms. If there are some properties we cannot admit, it will take more than a queerness argument to tell us as much.

Finally, understanding queerness as fundamental additivity has one more important virtue beyond fitting well into the Argument from Queerness properly understood. This account of queerness explains why the companions-in-guilt response to the Argument from Queerness is so compelling. In the companions-in-guilt defence (also called, perhaps more accurately, 'the plea of innocence by association'), it is suggested that there are other entities—or even just one other—that share X, the quality in virtue of which moral facts must supposedly belong to a whole new category of facts. In other words, it is the denial that moral facts really do force us to posit a new category to account for them because the existence of companions in guilt means that the category in question is one that we already have in our ontology. There are still debunking accounts available to explain moral phenomena, but we don't have positive reasons to favour them over moral realism (at least, none that are rooted in ontological concerns) and so the Argument from Queerness would fail. There are limits to this; some Moral Error Theorists are willing to grant that their queerness arguments have implications beyond the moral

domain.⁸² So, for the companions-in-guilt defence to work, the companion or companions in guilt must be something we are unwilling to reject because, for example, it is highly uncontroversial or it bears some otherwise undischarged explanatory burden of its own. Moral Error Theorists can respond by trying to undermine the claim that the companion has the putatively queer feature in question, which would reassert the need to account for moral facts with a new ontological primitive. Whether there are any companions in guilt is something that must be considered case-by-case with respect to each queerness argument (each candidate for X) and so something I will address as I examine those arguments. Nevertheless, reading queerness as fundamental additivity explains why the existence of companions in guilt is a threat to a queerness argument.

This reading also means that both parts the Argument from Queerness depend on a commitment to ontological parsimony, though each invokes a different form thereof, and I think I have said enough to motivate the view that this is an important part of assessing the quality of an explanation. (As an aside, it seems to me that if my interpretation is right, and the Argument is based foremost in a concern for ontological parsimony, this may explain why error theorists are typically metaphysical naturalists. Few people deny the existence of the external world around them—the same one described by natural facts—and much that once looked inherently mysterious has come to be naturalistically explained by science, so if one seeks to give the most parsimonious account of the world possible, naturalism looks like a good candidate.)

So now we know what queerness is, and what it isn't, so we are in a position to make judgements as to whether the feature of moral facts that a queerness argument nominates as X would actually make them queer. We also know that one way of doing this is by giving a companions-in-guilt argument to show the queerness argument in question overgeneralizes. Whereas I have given a single account of roughly what debunking explanations look like, each queerness argument must be tackled individually. Since our debunking explanation is (by stipulation) plausible enough to do the work the Argument from Queerness requires of it, the question

⁸²Mackie 1977, p. 39. Olson (2014, §8.2) specifically argues that his own view may lead to an epistemic error theory, but that this is an acceptable consequence.

of whether the Argument succeeds overall will be settled by whether there are any convincing queerness arguments, and the remainder of this chapter will be devoted to tackling them.

1.5 The Queerness Argument from Moral Knowledge

The queerness argument from moral knowledge differs from the other three in that it isn't meant to undermine the ontological status of moral facts. Instead it is aimed at problematizing our moral knowledge, by showing that there would be something queer about the way in which we come to possess it.⁸³ This need not imply the queerness of moral facts, so this argument initially appears to be of no consequence for our purposes. However, Mackie proceeds as if it does have ontological implications, and it is at least plausible that we should be sceptical of a domain such that knowledge about it would constitute a fundamental addition to our ontology (or perhaps our epistemology, in the sense of amounting a wholly new kind of knowledge). Further, a number of thinkers have claimed that this queerness argument ultimately rests on a claim about the queerness of moral facts. As such, it is worthwhile to examine it more closely.

Mackie's version states that, given the nature of moral facts, we could only be aware of them by way of some special faculty or intuition. This could be either a special form of perception or some kind of intellectual faculty.⁸⁴ This is how Mackie understands moral intuitionism, and he takes it that all forms of moral realism must involve an intuitionist move, where we gain at least some of our moral knowledge through just such a special faculty. David Brink contends that we can interject here, and point out that this argument presupposes the queerness of moral facts, as the faculty by which we come to possess moral knowledge will be queer

⁸³Mackie 1977, p. 38.

⁸⁴See Mackie 1977, p. 41; 1980, p. 147. In later work, he leans more toward the understanding of intuition as an intellectual capacity, as a way of addressing his arguments to the most plausible intuitionist positions.

in the relevant sense only if moral facts themselves are queer (such that ways of knowing that we already accept will not suffice).⁸⁵ But this is too quick. It fails to take into account the core assumptions of the Argument from Queerness. We are, for now, proceeding under the assumption that moral facts are non-natural facts. Moral knowledge (or, more precisely, the belief necessary for moral knowledge) is a psychological phenomenon, happening firmly within the bounds of the natural world. Things within the natural world—such as beliefs—form a common causal order. Non-natural entities (and so non-natural facts) are usually understood to be acausal, which is to say to fall outside this order, and it seems that any human capacity that was in causal contact with something that is definitionally acausal would have to be different enough from everything else with which we are familiar to be called ‘queer’.⁸⁶ To be clear, the issue is not that moral knowledge cannot be accommodated by a causal theory of knowledge like Alvin Goldman’s (which I have nowhere assumed), but that the story above seems to involve causal relationships where we wouldn’t normally expect to see them given the nature of the entities involved.⁸⁷

But perhaps this is too hasty as well. We are not necessarily forced to agree that non-natural facts are acausal, or incapable of interacting causally with the natural world. Consider, for example, theistic explanations for the existence of the universe. That the universe exists is a natural fact, but theists’ explanation for this is that God—a supernatural entity—created it, and the act of creation seems paradigmatically causal. Even if you don’t buy into this particular account, it doesn’t seem to describe anything impossible, so we cannot reasonably rule out causal interaction in at least one direction (notably, the direction Mackie’s picture of moral intuitionism requires). At the same time, though, while they are non-natural, supernatural facts may not be the best analogy for moral facts. Parallels are more often drawn between moral facts and the truths of mathematics, and it would be quite odd to attribute causal powers to facts about quantities and numerical relationships. This analogy does pose a separate problem for the

⁸⁵Brink 1989, p. 180.

⁸⁶Brown 2013, p. 629.

⁸⁷See Goldman 1967 for his causal theory of knowledge.

queerness argument from moral knowledge, which I will discuss shortly.

On Mackie's straightforward interpretation of intuitionism, at least, this queerness argument succeeds, insofar as it shows moral knowledge to be queer. What are the ontological implications of this? Strictly speaking, the conclusion we should draw is that we should dismiss the possibility of moral knowledge, and this does give us some reason not to posit moral facts. After all, if we can't (and therefore don't) have any moral knowledge, then at least one of the things that moral facts could serve to explain doesn't exist at all. However, the epistemological conclusion here is suggestive rather than conclusive on the ontological question. It's perfectly coherent that completely unknowable moral facts exist, and that realists might still appeal to them in some explanatory context.

In any case, these conclusions only hold as long as we assume that the only way to know non-natural facts is through the kind of intuitive faculty Mackie describes. As other thinkers have pointed out, the most plausible versions of intuitionism don't claim that we gain moral knowledge this way. Rather, many intuitionists hold that moral facts can be known intuitively because they are self-evident in a way that our (ordinary, non-specialized) rational capacities can grasp.⁸⁸ Knowledge gained in this way would be synthetic a priori knowledge, and so we can construct a version of this queerness argument that holds that synthetic a priori knowledge is queer.⁸⁹ However, we could then make a very simple appeal to companions in guilt: the rules of logical inference and the conclusions of mathematics are things we generally think we know, and they are paradigm cases of synthetic a priori knowledge which are distinctly non-queer.⁹⁰ Hence this version of the argument fails straightforwardly.

Insofar as we are concerned with the ontological status of moral facts, the queerness argument from moral knowledge gives us, at best, a supplemental argument. The queerness of moral knowledge gives us a reason to be suspicious of moral facts, but it is insufficient grounds to say that they themselves are queer. Further, our ability to claim that moral knowledge is queer rests on a specific in-

⁸⁸See Brink 1989, p. 180 and Stratton-Lake 2002, p. 22.

⁸⁹Olson 2014, pp. 101-2.

⁹⁰Olson 2014, p. 103.

terpretation of moral intuitionism, which does not seem to be reflective of most seriously endorsed versions thereof. As such, we need not feel that this queerness argument contributes to the threat the Argument from Queerness poses to moral realism.

1.6 The Queerness Argument from Motivation

When Mackie first gave the queerness argument from motivation, he didn't distinguish it carefully from the queerness argument from special normativity, which will be discussed in §1.8.⁹¹ Both are concerned with the way in which moral facts are supposed to give us reasons for action, but the reasons in question are of different kinds. Since, as I will argue, one of them ultimately fails where the other succeeds, it is important that we draw this distinction.

The queerness argument from motivation locates the queerness of moral facts in their ability to motivate us to act in certain ways or, in other words, to give us motivating reasons. Motivating reasons are psychological phenomena. They are the considerations that explain the actions we take, or would take under certain circumstances, in that they play a role in the causal story we would tell if we were to give an account of why someone acted or would act in a certain way.⁹²

If we are to claim that this motivating power is the queer element of moral facts, we must insist that all moral facts have this power. In short, we must be moral motivational internalists. In the broadest terms, moral motivational internalism (henceforth simply 'internalism') is the view that moral facts always and necessarily provide agents with motivating reasons to act in certain ways, specifically the ways those facts indicate are good or right. Because this is meant to be a conceptual fact about morality, we can know this a priori, and it will be true whatever the actual content of morality turns out to be.⁹³ What it actually means to say that moral facts give agents motivating reasons depends on interpretation, but when we elaborate on these interpretations none of them turns out to be very

⁹¹Olson 2014, p. 104.

⁹²This means that they are primary reasons, as defined by Davidson (1963, p. 686).

⁹³Brink 1986, p. 28.

plausible. I will lay each out as fully as is necessary for our examination, but I will not ultimately endorse any of them, for reasons I discuss below.

Brink outlines three broad interpretations of internalism that he calls ‘agent internalism’, ‘appraiser internalism’, and ‘hybrid internalism’. For each interpretation, I will examine separately the questions of whether the described connection between moral facts and motivation itself would be queer, and of whether it is plausible that such a connection is an indispensable part of our conception of moral facts.

Agent internalism is the claim that, in virtue of the concept of morality, moral facts themselves give us motivating reasons.⁹⁴ That is to say, by merit of nothing more than their existence, moral facts would lead to us having at least some motivation to act in certain ways. Such a relation would be queer. It would involve a move straight from an ontological fact—that there are moral facts—to a psychological fact about the motives of agents, and since the relationship in question is supposed to be necessary it couldn’t involve processes like perception or reflection, which are liable to go awry. Such a relationship would be distinctly strange, and if it were a causal relationship we would also run into the further problem of non-natural moral facts having causal contact with natural psychological facts as described in §1.5. So a queerness argument from moral motivation based on agent internalism works in at least this regard.

But it’s implausible that this kind of internalism is conceptually true. If it were, then moral realists would be forced to reach some very odd conclusions. Supposing that there were objective moral facts with some direct sway over our motivations, we would see a lot of moral agreement, when in practice different people’s moral motivations can pull them in entirely different directions. Consider, for example, the debate about whether abortion should be legal; proponents of both sides seem to be motivated by their moral commitments and to give moral arguments for their positions. Of course, some will want to insist that such disagreements occur

⁹⁴Brink 1989, p. 40. Brink himself discusses ‘moral obligations’ giving us such reasons, but I think I am justified in switching the subject because it is the ontological status of moral facts that is in question. Also, it’s not clear that a deontic moral fact is in any way distinct from a moral obligation, since both exist in virtue of the distribution of moral properties.

because people disagree about the facts of the situation, and that their fundamental moral motivations are the same. This raises questions about whether the moral facts that exist (if there are any) are concerned with general principles, or the rightness of particular actions in particular cases, or both, which is a level of metaphysical inquiry that would draw us off the course of our main discussion.⁹⁵ In any case, such an objection exists here, but we are in no position to assume that disagreement over matters of natural or supernatural fact underlies every case of moral disagreement, though I have to acknowledge there may be no cases of unambiguous pure moral disagreement outside of the philosophy room, since there may be other kinds of fact at play in any situation.

A more promising line of response says that internalism only says that moral facts would give us some reasons, not that these reasons would be overriding, so perhaps this is a case of different people responding to different motivations. This is handily explained so long as we allow for the possibility of people's other motivations varying in strength and so overruling or conceding to moral motivations (the variation must be in these other motives because if moral facts have an unmediated effect on our motivations, then we cannot explain how moral motivations would vary from person to person). Attention to the world gives us every reason to grant this possibility, but it also highlights a problem already raised. Even correcting for other motivational factors, it seems like the content, and not just the strength, of people's moral motivations is strongly prone to vary. We might respond by suggesting that there are multiple moral facts at play in any situation, but we would then encounter the same problem described above of accounting for how they differ in strength between agents when their influence is meant to be immediate. The only way we might get around this problem would be to claim that moral facts are not just about what would be good or right to do, but what would be good or right for a specific person to do; that moral facts are agent relative. This wouldn't contradict any of the assumptions I made about morality at the beginning of this chapter, so let's suppose we can make this move. It still looks like there is a problem, because now we would just have to conclude that whatever

⁹⁵Given the definition of facts laid out in §1.1, however, any answer other than 'both' seems artificially limited.

motives we already have are reflective of the moral facts.⁹⁶ Further, we would all presumably reach the same conclusions about what a given person should do under given circumstances. So long as we think it is possible for someone to have false moral beliefs, this conclusion is absurd.

These might seem more like arguments that agent internalist moral realism isn't true than arguments that agent internalism isn't a conceptual claim about moral realism, and they certainly do support the former conclusion, but I think that it would be too uncharitable to attribute realists—and people in general—such an obviously flawed understanding of moral facts. We all acknowledge moral disagreement, and the possibility of moral error. At the very least, we see clear cases of this when theorists make contradictory claims about the most general principles of substantive ethics—those who endorse utilitarianism plainly disagree with those who reject it. But if the facts about these matters have an unmediated influence on our fundamental moral motives, we would expect to see much more agreement. And if the category of moral facts extends to facts about what is good or right in particular concrete scenarios, we would expect much more widespread agreement on what to do in any given case because all agents would have at least some motivation to take the morally right course of action. Given that, we cannot reasonably say that our ordinary concept of moral facts commits us to agent internalism simply because it so directly implies this kind of general agreement, which no one could reasonably endorse. Even if we take it that moral facts are capable of motivating us, I don't think anyone is claiming that they always do so in some direct and unmediated way, so it wouldn't be plausible to claim that this is conceptually essential to moral facts. As such, this version of the queerness argument from motivation fails, even though the relationship agent internalism describes would be queer.⁹⁷

⁹⁶Brink 1986, p. 29, interpreting Frankena (1958).

⁹⁷Brink doesn't identify anyone as an agent internalist, and given the obvious difficulties into which the view runs it may be that there are none. I do not think this matters, however, as Brink is laying out a possibility space so the view can be addressed systematically, and we have the same responsibility as he does to examine each part of that possibility space insofar as it is relevant to our purposes, hence the above observations.

The second version of internalism is appraiser internalism, the claim that moral beliefs necessarily motivate.⁹⁸ While we might plausibly claim that such a relationship between belief and motivation would be queer in the sense of being unusual (and even this claim would require us to presuppose or argue for a distinctly Humean model of motivation), and while this claim is about morality in the broadest sense, it is not a claim about moral facts.⁹⁹ As such, a queerness argument based on this form of internalism wouldn't have any implications for the debate about the ontological status of moral facts. Even if a successful version of this argument were to be constructed, it would show only that moral beliefs are queer, and so what? The existence of moral beliefs was never in doubt. In fact, they are examples of the moral phenomena that both moral facts and debunking accounts are attempts to explain. Furthermore, the queerness of moral beliefs wouldn't have any implications for moral facts; a fact is a fact precisely because it underlies a truth regardless of whether (or how) we believe that truth. So a queerness argument based in appraiser internalism cannot establish the queerness of moral facts.¹⁰⁰

Lastly, we have hybrid internalism, which splits the difference between the other views and is perhaps more promising than either. Hybrid internalism is the view that the recognition of a moral fact is necessarily motivating.¹⁰¹ 'Recognize' is a success verb; if you think you recognize a face in a crowd and call out to them, only to learn that you were mistaken and the person in question is not who

⁹⁸Brink 1989, p. 40.

⁹⁹By a 'Humean model' of motivation, I mean one wherein beliefs and desires are completely distinct kinds of mental state or event, and while they may interact there is never any necessary relationship between them, so no belief can entail a desire or vice versa.

¹⁰⁰For this reason we can forgo discussion of the 'moderate internalism' described by Dreier (1990) and van Roojen (2010, p. 499), which states only that moral judgements normally lead to motivations. In any case, such a view would not have sufficed for a queerness argument because it allows for exceptions like Brink's amoralist (discussed below), while for a queerness argument to work and problematise the entire moral domain the queer feature must be attributable for all moral facts.

¹⁰¹Brink 1989, p. 41.

you thought, you would have to admit that you had not recognized them at all. In short, you can only recognize something when you actually encounter it. So recognizing a moral fact is about being properly acquainted with it, which is to say having a moral belief that is both true and justified.¹⁰² It's not obvious what else there might be to 'proper acquaintance' with a moral fact, or how one can attain it, though we can at least say that proper acquaintance rules out disagreements about what the moral facts are (and indeed what counts as a moral fact) among those who share it. For the sake of argument, let's suppose that such an acquaintance is possible. We have, then, a view wherein proper acquaintance with a moral fact (of the kind which is fit to ground belief) entails that the acquainted agent has some motivating reason for action.¹⁰³ This seems to be the view of moral facts that Mackie held. He compares them to the Platonic Form of the Good, knowledge of which is meant to inherently motivate one to pursue goodness.¹⁰⁴ So at least some Moral Error Theorists have thought this way, and thought we could construct a working queerness argument on the back of hybrid internalism.

So we return to our two key questions. First: would facts that necessarily

¹⁰²It will occur to many readers that 'justified true belief' (JTB) is the paradigmatic rough-and-ready analysis of knowledge. To be clear, I am not saying that hybrid internalism is the view that moral knowledge is necessarily motivating. This is partly because the JTB analysis has been problematized in a number of ways (see Russell 1948 and Gettier 1963 for the most influential examples, and Ichikawa and Steup 2018 for an overview), and partly because internalists usually do not speak in terms of moral knowledge. Furthermore, I wish to avoid confusion between the queerness argument from moral knowledge, which effectively argues that it is the 'proper acquaintance' that is queer rather than moral facts themselves, and the version of the queerness argument from moral motivation I will go on to build, which is concerned with how moral facts affect agents assuming this kind of acquaintance is possible.

¹⁰³I say 'some' reason because the same fact could provide different agents with reasons to perform different actions. If it would be wrong for some agent A to ϕ , that is a reason for A to refrain from ϕ ing, and for some other agent B to prevent or dissuade A from ϕ ing.

¹⁰⁴Mackie 1977, p. 40. To be more precise, he held that this was a conceptual truth about moral facts, while also maintaining that nothing fit this description, hence his endorsement of Moral Error Theory.

provide motivating reasons to anyone who was properly acquainted with them be queer? It's plausible that they would be. For one thing, if proper acquaintance with a fact involves causal contact with it, we encounter the same problem we discussed in §1.5 of acausal non-natural facts affecting the psychology of agents in the natural world.¹⁰⁵ But, even if we suppose this relation is unproblematic, the facts in question would still be distinctly unusual, perhaps even unique. This is because they would be such that it would be impossible to be indifferent to them. We aren't really in a position to say whether there are any companions in guilt for such facts, because it's doubtful that we could ever verify with certainty that we were properly acquainted with the facts in question, even if we did have an account of what proper acquaintance amounts to. What is clear, though, is that there is an enormous number of facts that don't behave this way. We generally expect the kinds of fact science discovers, for example, to be motivationally inert, or to motivate us contingently when they are relevant to some desire or goal we already have. An internalist of straw might suggest that we simply aren't properly acquainted with them, and they would give us motivating reasons if we were, but this is plainly question begging. Plus, it's unclear what such facts would motivate us to actually do. Since there are so many examples of facts that don't necessarily motivate, and no way to confirm any companions in guilt, it seems fair to say that if proper acquaintance with moral facts did allow them to exert some mysterious 'motivational pull', this would be sufficient grounds to call them queer.¹⁰⁶

This turns out to be rather moot, however, because of our second key question: is hybrid internalism a conceptual truth about moral facts? The answer is fairly clearly 'no'. In order to prove this, Brink asks us to imagine a figure he calls the amoralist. This is an agent who wouldn't be motivated by moral facts, even if

¹⁰⁵While I noted before that internalists do not generally speak in terms of moral knowledge, and I don't intend to commit myself either way, it's still plausible that proper acquaintance amounts to knowledge of one kind or another. Olson (2014, p. 110-1) suggests that it consists in 'first-hand' knowledge, on the back of which he makes an argument like the above.

¹⁰⁶The phrase 'motivational pull' is borrowed from Olson (2014, p. 109), and is meant to convey the vaguely defined effect moral facts would have on an agent's mind under this account.

they were properly acquainted with them.¹⁰⁷ If we suppose that proper acquaintance amounts to knowledge, we can say such a person knows right from wrong, and is simply unmoved by such concerns. Fiction is replete with examples like Plato's Thrasymachus, who distinguishes between justice and injustice but does not favour the former. Certain sociopaths in real life are arguably such figures as well. Whether or not we find such people in the real world, the fact that we can conceive of them is quite telling, because if such a figure *could* exist this implies it is not part of the concept of a moral fact that recognizing it necessarily leads people to be motivated in certain ways. We could even go further and posit an agent who perversely seeks to do what is bad and wrong.¹⁰⁸ Again, we find examples in fiction, and even the devil as understood by some Christians could be characterized this way. One popular response is that such figures are not using moral terms sincerely, but as though they were in inverted commas.¹⁰⁹ That is, they do not really believe that the things in question matter, they are simply referring to what others generally take to be good, bad, right, and wrong. This objection might work against appraiser internalism, but *ex hypothesi* our amoralist's belief is genuine and well-founded. We might think an analogous argument can be made, claiming that they are not really properly acquainted, but this would again beg the question, and as I have argued there is likely no way to settle the debate of whether the figures in question really are properly acquainted with the moral facts they both avow and ignore. It seems that we have a case of conflicting intuitions between internalists and externalists (defined as those who deny internalism) but this in itself is telling.

One other avenue of objection to amoralists, which aims to more fully justify the claim that they are impossible, appeals to a particular notion of what it is to act for a reason. This is the view that whatever an agent's motivating reason is, that agent understands it to be a fact in virtue of which the action of which it stands

¹⁰⁷Brink 1986, p. 30. This is very different than the figure of the same name discussed by Raz (1999a, ch. 12), who is not properly acquainted with moral facts in that they are either wholly ignorant of morality or actively disbelieve it.

¹⁰⁸Joyce 2001, pp. 20-3.

¹⁰⁹Hare 1952, pp. 124-6, 163-5.

in favour would be good.¹¹⁰ This is more of a commitment about the nature of reasons than I intend to make at this juncture, but let's suppose that it is correct. The amoralist is still possible under such a view because, while it ties motivating reasons to judgements about value or goodness, it says nothing that indicates it is always and only moral goodness to which people respond when multiple values stand in competition (there are other kinds of value, such as aesthetic and practical value, as well as goodness for particular agents rather than in-and-of-itself). Nor does it imply that the recognition of value necessarily entails the recognition of a reason, leaving open the possibility of indifference to moral claims in particular.

What really matters to this debate is whether the amoralist is conceptually possible, whether the idea of such an agent is coherent at all, but if we assume that both sides are employing, and have a command of, the same concept of moral facts, then the fact that there is anyone that can make sense of the concept of the amoralist suggests that hybrid internalism is not a conceptual truth about morality. The internalist might insist that committed externalists and laypeople alike are not making sense of the amoralist because it is a nonsensical concept, yet it is not as obviously so as a circle with corners or any other patently self-contradictory concept, so this is something that must be argued for independently of their internalism, and no such argument seems to be forthcoming.

The forms of internalism described by Brink generally give us good starting points from which to construct a queerness argument, but it might be objected that they are not exactly representative. That is, they don't much resemble the best and most commonly endorsed forms of motivational internalism. As such, it would be irresponsible of me not to consider some more popular alternatives, though as we shall see the same factors that make them better accounts also make them less relevant to our discussion here.

Many of these views explicitly concern judgements rather than facts, meaning

¹¹⁰See Anscombe 1963, p. 70, Goldman 1970, p. 94, Davidson 1978, p. 55, and Stampe 1987, p. 355 for more or less explicit claims along these lines. Such a view might be grounded in the claim that reasons for action in the normative sense discussed in §1.8 are conceptually those facts in virtue of which an action is good. I discuss views like this in more depth in §2.2, but see also Raz 1999b for characterization and defence of such a view.

they are psychological claims rather than the kind of metaphysical claims we are concerned with here. In other words, they are forms of appraiser internalism, and so are irrelevant to our inquiry for reasons already discussed. One example that doesn't fall into this trap, though, is Korsgaard's internalism, which can apply to the acceptance of a moral judgement but also to the truth or knowledge of it.¹¹¹ Both of these latter options implies the existence of a moral fact, so there can be agent and hybrid versions of her internalism.

Her view, however, doesn't posit a strictly necessary connection between these facts and human motivation. She takes it that moral facts (or rather, the subject of internalism, which is moral facts for our purposes) are 'practical reasons', and in order to be such they must be at least capable of motivating action, specifically in rational persons.¹¹² Put another way, moral facts generally do give us motivating reasons, but their ability to do so is conditional on the rationality of agent to whom they apply. Therefore, it's possible for a person to remain unmoved by moral considerations, but only at the cost of being irrational. This model is plausible, especially compared to some of the versions Brink outlines, but that isn't really what we're discussing here. We are only interested in whether this is an essential feature of moral facts, and whether it's sufficient to make them queer. On the former count, while the overall picture is more plausible than the alternatives, we would still be hard pressed to call it an *essential* feature for the same reasons as hybrid internalism. On the latter count, the answer is quite complicated, because it may turn out that what I have just described isn't motivational internalism at all.

This model has it that, as long as an agent is rational (or acting rationally), moral facts (or accurate moral judgements) will give them motivating reasons. Korsgaard says true irrationality consists in 'failure to respond to an available reason.'¹¹³ She doesn't say what kind of reason, but it can't be a motivating one without the whole account being circular. Given her definition of irrationality, it seems like for Korsgaard—as for many philosophers—rationality consists in prop-

¹¹¹Korsgaard 1986, p. 8.

¹¹²Korsgaard 1986, p. 11.

¹¹³Korsgaard 1986, p. 12.

erly responding to reasons, specifically reasons for action (a response is something one does, even if passively or unconsciously) which are the subject of an entirely different queerness argument (an argument that is addressed in §1.8). So it seems that Korsgaard's internalism does give us fodder for a queerness argument, but not the queerness argument from motivation. There is nothing queer about rational people being motivated to respond to reasons, and in any case her model seems to hold that it is not the facts but the rationality of agents that is the source of this motivation.

So, however we cash it out, the queerness argument from motivation fails, because we have good reasons to think agent internalism and hybrid internalism are not conceptually true of moral facts, and because appraiser internalism only problematizes moral beliefs, whose existence is not in doubt.

1.7 The Queerness Argument from Supervenience

It is generally accepted that moral facts supervene on natural (or even supernatural) facts.¹¹⁴ I'll borrow a shorthand from Olson, and refer to natural and supernatural features collectively as *natural** features.¹¹⁵ One set, family, or—to stick to the parlance I have been using so far—domain of things supervenes on another when there can be no change in the former without a corresponding change

¹¹⁴The paradigmatic example of moral status depending on the supernatural is Divine Command Theory, which holds that actions are impermissible or obligatory only if they have been respectively forbidden or commanded by God. Other than pointing to such examples, I have no way of distinguishing the supernatural from other forms of the non-natural, but though this is rather vague it is still sufficient for our purposes.

¹¹⁵Olson 2014, p. 89.

in the latter.^{116,117}

So, if the moral supervenes on the natural*, two worlds that are identical in their natural* features will necessarily be identical in their moral features.¹¹⁸ For example, a utilitarian would say that the fact that some action is wrong supervenes on the facts about its consequences and those of available alternative actions, and if circumstances were to change such that all the alternatives caused more overall suffering, it would also cease to be wrong. This connection is, as noted above, necessary, but it is not conceptually necessary.¹¹⁹ We can say this because we are aware of much disagreement—between people who apparently use moral terms in the same way—with regard to what moral significance any particular natural* fact has.¹²⁰

This relationship is precisely what the queerness argument from supervenience nominates as the queer element of moral facts. So is supervenience on the natural* an essential feature of moral facts? I am inclined to say that it is. If it were not, we could claim that some pair of actions are exactly alike in all their (intrinsic and extrinsic) natural* properties—perhaps they are even happening in possible worlds which have hitherto been alike in all such properties—but differ in their moral properties. Such a claim, however, is likely to baffle hearers, and lead them

¹¹⁶It is more typical to speak of domains of properties standing in supervenience relations than to apply it to domains of facts, but there is precedent for discussion of the supervenience of facts, and in any case we can get the supervenience of facts out of the supervenience of properties given the definition of fact introduced in §1.1. See also Kim 1984, pp. 154-5, 169.

¹¹⁷The supervenience relation is fairly simple in and of itself, but it need not be governed by any simple laws. See Davidson 1963, pp. 3-4.

¹¹⁸Strictly speaking this is global, rather than individual or regional, supervenience, but nothing turns on this distinction for our purposes.

¹¹⁹Olson 2014, p. 91.

¹²⁰To be more precise, we would be unjustified in claiming that the necessity of moral supervenience is conceptual. We cannot dismiss the possibility that there actually is some conceptual connection between moral and natural* facts which is difficult to discover, and so has gone undiscovered so far. Either way, we cannot assume that a connection is conceptually necessary just because we have posited that it is metaphysically necessary.

to conclude that we don't really understand the moral concepts we are invoking.

We might submit that there are ways to resist such a claim. Imagine, for example, a cult whose members insist that whatsoever its leader does is perfectly good and right, even if it would have been a dire misdeed had it been perpetrated by anyone else. We would likely want to disagree with them, but we wouldn't want to say that they hadn't grasped the moral concepts they were invoking. This, initially, might seem to show that moral facts can differ even while natural* facts remain the same. However, though two actions may be the same in their natural properties (in the narrower sense that excludes the supernatural), there is still a relevant difference in the form of the supernatural fact that one act was performed by a messianic or prophetic figure, and the other wasn't. We might dispute both the truth and the moral relevance of this supernatural claim, but what matters is that the difference in moral assessment is still based in a perceived difference in natural* properties.

It seems that supervenience onto the natural* is indeed part of the concept of moral facts. To say anything else would be to ignore most (if not all) of both lay and philosophical moral discourse. If the queerness argument from supervenience is to succeed, though, it must be shown that such a relation is genuinely problematic in the relevant way.

Supervenience itself is a rather shallow relation; all it involves is covariance between two domains (in this case, of facts).¹²¹ It does not fundamentally explain anything (if anything it stands in need of explanation itself), but it does suggest the presence of some deeper metaphysical relation.¹²² Moral Error Theorists have appealed to these relations when making the case for the queerness of moral supervenience, typically a dependence relation. It's plausible to think that moral facts don't simply change together with natural* facts, but that they obtain 'because of' some other, 'lower-level' quality, and Mackie asks 'just *what in the*

¹²¹To be more precise, it is a two-place relation, but supervenience is reflexive, meaning that every domain trivially supervenes on itself. Hence, there are supervenience relationships involving only one domain.

¹²²Kim 1993, p. 167.

world is signified by this “because”?’¹²³ This might sound like an appeal to the kind of bafflement Olson discusses, and so not necessarily persuasive, but thinkers since Mackie’s time have begun to pick apart what might be troubling about this connection.

The ‘because’ in question is not a causal one. Rather, a thing is good ‘because’ it has certain natural* properties in the sense that having those properties ‘makes’ it good, in the same way we might say being short is part of what ‘makes’ it hard to reach objects on high shelves, or having a certain colour scheme is part of what ‘makes’ an item of clothing gaudy. So the natural* properties explain or determine the moral ones.¹²⁴ This is the more metaphysically robust notion of a grounding relation, and is the kind of connection that could explain moral supervenience.¹²⁵ One might be inclined to wonder, then, why we are discussing supervenience rather than grounding.

Moral Error Theorists who have made arguments along these lines have been concerned with the connection between the moral and the natural*, and have not been especially precise about its nature.¹²⁶ The argument I will give below only requires the weaker supervenience claim, so there is nothing to be gained by committing ourselves to the stronger grounding claim, even if it is almost, if not equally, as plausible as the supervenience claim.

With all that groundwork laid, we can turn our attention to the question of

¹²³Mackie 1977, p. 41. Olson’s (2014, §5.1) account of this argument also rests on such a dependency relation, although notably he takes his version, and the queerness argument from supervenience more generally, to fail to motivate Moral Error Theory.

¹²⁴This is a view to which at least some moral non-naturalists are explicitly committed. See, for example, Audi 1997 and FitzPatrick 2008.

¹²⁵See Chilovi 2021 for arguments that grounding entails supervenience, though we don’t necessarily need to grant that conclusion to see how grounding can explain a regular covariance, especially given how expansive the supervenience base—the natural*—is in the case of moral supervenience, allowing for fewer potential defeaters.

¹²⁶Mackie (1977, p. 41) did not name the relation he was talking about, seems to have been discussing grounding, and Olson (2014, pp. 88-9) invokes supervenience but his argument ultimately turns on both a supervenience claim and a grounding claim.

whether moral supervenience is, in fact, sufficient to make moral facts queer. As we have just noted, a supervenience relationship doesn't strictly mean that the supervening facts obtain 'because of' the supervenience base, and even if we posit a grounding relationship between the same sets of facts, the 'because' in question is not a causal one, and so the issues around non-natural facts entering the causal chain that were discussed in §1.5 do not come into play.

An alternative answer invokes the metaphysical principle *Hume's Dictum*, also called the coarse-grained criterion of property identity.¹²⁷ This principle holds that there are no metaphysically necessary connections between distinct entities—in this case properties (and, therefore, facts). This implies that two distinct properties cannot be necessarily coextensive, which is where two properties always and only occur together. Where there appeared to be such a connection, there would actually be just one property. That moral properties are distinct from natural properties is a core claim of moral non-naturalism, but we cannot reasonably deny the supervenience of the moral onto the natural*, so if Hume's Dictum is correct no moral property could actually exist.

One might at this point object that I am getting ahead of myself, for the discussion has moved from supervenience to co-extension, and these are not the same relation. However, the former does imply at least some instances of the latter. Consider the moral property of wrongness. Whatsoever is morally wrong also has some set of natural* properties, and some of those properties are morally relevant insofar as they would have to be different for any given wrong thing to have a different moral status. Let A_1 be the name of a specific wrong action, and M_1 be a predicate that means 'has the same set of morally relevant natural* properties as A_1 '. Let A_1, A_2, \dots be the set of all wrong actions across all possible worlds, with each A_i having its corresponding M_i . We can then construct a further predicate M' which means 'satisfies the conditions of either M_1 , or M_2 , or ...'. M' will be extensionally identical (in other words, coextensive with) the predicate 'is wrong', and since the conditions of M' are to possess one of a large number of sets

¹²⁷Moberger 2020, p. 328. In addition to the definition I am about to give, it can (for reasons explained below) also be characterised as the claim that two properties are identical if and only if they are coextensive.

of natural* properties, there is a natural* property (albeit a kluged together and highly disjunctive one) which is coextensive with wrongness.¹²⁸

If Hume's Dictum holds, then, moral facts that are distinct from the natural* facts on which they supervene are impossible. Impossibility is, of course, not the same thing as queerness, but it motivates the same conclusion; we should not admit impossible things into our ontology.

The natural question that follows this conclusion is whether or not Hume's Dictum is correct. It may have some intuitive appeal, but many thinkers have pointed out that it's open to a number of counterexamples. Some properties do seem to stand in a relation of necessary coextension. For example, all equilateral triangles are also equiangular triangles, and changing the length of any of the lines that makes one up also changes the angles between it and adjacent lines, so it looks like the property of being an equilateral triangle is coextensive with the property of being an equiangular triangle. Likewise, having a size and having a shape seem intuitively to be two different properties but in practice they are necessarily coextensive.¹²⁹ This view would further imply that impossible properties more generally, which necessarily have no extension, are all logically equivalent to one another.¹³⁰ Either that, or they are made somehow more impossible (whatever that means) by their shared extension.

In fact, it may be too quick to dismiss Hume's Dictum as grounded on intuition. The truth of Hume's Dictum, and the viability of the above counterexamples, may both turn on what we mean by 'property'. Bart Streumer suggests that properties are 'ways objects can be'.¹³¹ This is contrasted against an understanding of

¹²⁸This is paraphrased (and simplified) from Streumer 2017, pp. 9-11, which is itself derived from similar arguments in Kim 1993, pp. 69-71 and Jackson 1998, pp. 118-25. One possible line of objection is that there can be no disjunctive properties (see Armstrong 1997, Audi 2013 and Meinertsen 2021 for arguments to this effect) or that natural* properties in particular must be unified in a way M' clearly isn't.

¹²⁹Streumer 2017, p. 15.

¹³⁰Olson 2014, p. 93.

¹³¹Streumer 2017, p. 12. I take it that when Streumer says 'objects' he doesn't necessarily mean physical objects, but anything to which we can reasonably ascribe properties.

properties as ‘shadows of concepts’, which is to say as meeting the conditions for certain understandings or descriptions. This means that if the only way for an object to fit some particular description is for it also to fit some other description and vice versa, as in the triangle case outline above, then there is really only one thing happening—one *way things are*—to which those descriptions refer.¹³² Likewise, being water and being H₂O are conceptually different but in reality nothing but H₂O could have turned out to be water in the sense of being the thing that fills oceans and falls as rain and quenches thirst, so there is only one property here. In other words, Hume’s Dictum is true because when two concepts, descriptions, or predicates would necessarily have the same extension, they express a single property. This also dissolves the counterexample of impossible properties, because if nothing can be that way then the way in question is, under this definition, not a property at all.¹³³

However, while this understanding seems to resolve these issues, it faces a serious one of its own. I have here introduced it as a way of understanding, or perhaps defining, what a property is which supports Hume’s Dictum by way of entailment. However if coextension entails identity as Hume’s Dictum claims then it follows that each property is a unique way for objects to be. The entailment is biconditional. Hence any support this understanding lends is circular; the whole system is coherent but it lacks external backing.

And so we are back where we started, with an argument that relies on a highly contestable view about properties and how they are to be individuated, and short on independent arguments for that view. I won’t try to reach any conclusions about these issues here, as they fall well outside the scope of our current investigation. Suffice it to say, the number of provisos we need to attach to Hume’s Dictum make it rather shaky grounds for rejecting moral facts.

¹³²Sober (1982, p. 185) suggests these are distinct because we can imagine a machine that detects one and not the other, but since there is no way to present such a device with a triangle that is equilateral and not equiangular or vice versa, we have no way of knowing for sure to what it is responding. If Streumer is right, it would just be detecting the one way the figure in its totality actually is.

¹³³Streumer 2017, pp. 17-18.

We could submit that Hume's Dictum is actually an empirical claim and that we should believe it on the grounds that we have no good counterexamples to it, or that apparent counterexamples can be explained away. But then the claim that metaphysically necessary coextension between distinct properties is *impossible* is too strong; empirical enquiry can give us strong inductive grounds for claims about what is likely to be possible in light of our current understanding and thereby tell us which claims we should find *prima facie* suspicious, but it can never show in absolute terms that something cannot happen. So we are again at a point where it seems we cannot reasonably rely on Hume's Dictum as a grounds for a queerness argument. As the Dictum appears to be a dead end, if there is to be any plausible version of the queerness argument from supervenience, it must be backed up in some other way.

If Hume's Dictum is too strong, then perhaps what we need is a more modest claim, such as the *Modest Humean Dictum*, a principle introduced by Tristram McPherson. The Modest Humean Dictum goes as follows: 'Commitment to brute necessary connections between discontinuous properties counts significantly against a view.'¹³⁴ This wording suits our purposes particularly well. For one thing, we need to do a lot less to justify the claim that supervenience 'counts significantly against a view' than to justify the claim that it's simply impossible, and counting against a view is all queerness needs to do, as discussed in §1.2 and §1.4. For another, it speaks only of connection, for which supervenience alone is sufficient, and so we can leave aside the above argument for necessary coextension and address our arguments even to those who would be inclined to reject disjunctive properties. The third and final aspect of this principle that makes it suitable to our needs is that the way it characterizes the supervenience relation between the moral and the natural tells us *why* we should think that such relations are queer, rather than simply positing that there can be no such relations. Specifically, it is of particular interest that these relations are described as 'brute'. What does this

¹³⁴McPherson 2012, pp. 217-8. McPherson introduced this principle under the name Modest Humean, which I have tweaked into its current form for clarity. 'Discontinuous' properties here are those which do not stand in identity or reducibility relations with one another, those which are well and truly distinct.

mean?

In short, it means that the connection is inexplicable. Once you have accepted that the moral does indeed supervene on the natural*, there is an explanatory question to be answered: Why do these families of fact covary as they do? We could, as noted above, appeal to some other, more metaphysically robust relationship such as grounding to explain moral supervenience. But positing such a relationship—even as necessary—does not discharge the explanatory burden because the same questions arise around the grounding relationship: Why do these natural* properties make a thing good or bad or possessed of any other moral property?¹³⁵ Supervenience could also be demystified by positing an identity relation or a reduction relation between the supervening and subvening domains, but such relations are explicitly ruled out by the non-naturalist conception of moral facts with which we are working.

When McPherson calls this relation ‘brute’, he is saying that there is no answer to this question; that’s just the way things are. This kind of ground-level fact which doesn’t admit of any further understanding is an ontological primitive and so a *sui generis* entity. There is potential for disagreement as to whether the connection in which moral supervenience consists is an entity unto itself or a relational property of entities, but it does not matter for our purposes since if it is a property it is a property which itself is sufficient to make its putative possessor queer. So, in the former case moral facts imply the existence of a queer entity and in the latter case they entail the existence of an entity with a property that makes it queer, which is to say a queer entity.

As we discussed in §1.4, ontological primitives are necessary elements of any ontology and not inherently problematic.¹³⁶ Nevertheless, each such entity we invoke makes our ontology just that little bit more untidy and mysterious. In short, we consider brute posits to be a cost that a theory (in this case an explanation

¹³⁵McPherson 2012, pp. 221-2.

¹³⁶This is the case for *sui generis* entities, but moral supervenience will have even less of a leg to stand on if we take it to be a relational property, since it’s much less obvious that we must posit any *sui generis* connections to make sense of the world, as McPherson (2012, pp. 229-30) observes.

of moral phenomena) must justify. This, too, reflects a lot of what I said when we discussed queerness, which makes this promising as the basis for a queerness argument.

The supervenience of the moral onto the natural* is brute—it is hard to see how it could be anything else—and brute relations are paradigmatically queer. If all this is true, the last recourse by which we might try to undermine this queerness argument is with a companions-in-guilt defence. Such an argument, premised on the claim that brute supervenience relations are too common for us to reasonably do away with, has already been made by Brink. The reason it is hard to see how the supervenience of the moral onto the natural* could be anything but brute is because moral facts are meant to be *sui generis* in the Shepskian sense discussed in §1.4. In other words, they are a type unique unto themselves, which makes understanding how they would connect to other domains difficult. But Brink suggests that our reasons for seeing morality as *sui generis* are equally reasons to see other domains as *sui generis*, and that as such we see this kind of brute supervenience between, for example, different fields within science.¹³⁷

Certain natural facts do seem to be supervenient—for example, chemical processes supervene on the interaction of microphysical particles—but there is an important disanalogy here that causes this argument to fail, at least given our background assumptions. Here, Brink saying that this kind of supervenience is the same as the kind that holds between the moral and the natural*, specifically that it is brute and holds between ‘discontinuous properties’. However, it seems to be one of the primary virtues of science that the different domains within it are *not* discontinuous; the chemical processes are constituted by (or reducible to) microphysical ones, and in turn the events happening on that lower level explain the higher level events that occur on the observable scale.

There are some cases within the remit of science that are not so simple, such as the supervenience of the psychological onto the physical. It is widely accepted among both philosophers and scientists that some such relationship holds, but it is significantly more mysterious than the supervenience of the macrophysical onto the microphysical. However, this need not undermine the queerness argu-

¹³⁷Brink 1989, p. 174.

ment from supervenience. That this connection is unexplained doesn't necessarily imply that it is inexplicable in principle in the way a brute connection is, and it may be premature to declare that science will never be able to give us such an explanation. Even supposing that this connection really is brute, would-be Moral Error Theorists have the option of biting the bullet, accepting that the presence of this connection is a theoretical cost, and arguing that it is balanced out by the observations it explains. This latter strategy is not open in the parallel case of moral supervenience when there is a debunking account of moral phenomena to hand.

Supervenience that is unproblematic because it holds between continuous properties is, according to Moral Error Theorists, precisely not what we have in the case of moral supervenience. One of the core assumptions that the Argument from Queerness makes—and which we are accepting, for now—is moral non-naturalism, the claim that moral facts are of a fundamentally different kind than natural facts. This is often cashed out in terms of an 'is/ought gap' between the natural* and the normative (and so, implicitly, the moral). Brink thinks this gap is no worse than the 'is/is' gap between different domains of natural discourse because he is a moral naturalist. But it is a core assumption of the Argument from Queerness that, given the nature of moral language, 'normative facts are just too different from natural facts to be a subset thereof.'¹³⁸ As such there is no way for this kind of connection to be explicable.¹³⁹ Hence, under the assumption of non-naturalism Brink's companions-in-guilt argument fails because, although supervenience per se isn't unusual in any way, the kind exhibited by moral facts is exceptional relative to the examples Brink invokes, and would raise precisely the kind of explanatory questions that give the Argument from Queerness some of its force.

As an aside, there is one form of non-naturalism that would avoid at least the latter problem. A moral supernaturalism similar in structure to naturalism, that took moral facts to be constituted by or identical to some supernatural facts—

¹³⁸Enoch 2011, p. 4. The quote is explicitly concerned with the normative, but most thinkers subsume the moral under the normative so the intention is to cover moral facts with this claim. Whether this subsumption is warranted is discussed in Chapter 2.

¹³⁹McPherson 2012, p. 226.

like one wherein wrongness was identical with being forbidden by God—has no such issues because there is no problem for any fact's relationship to itself. Even Mackie admitted that there would be room for moral facts if there were a God issuing commandments.¹⁴⁰ However, such a view depends on the ideas that there is a God and that said God has issued certain edicts, neither of which we are in a position to assume, besides which we certainly cannot call such beliefs part of the standard conception of moral facts from which we are working, even if we take them to be compatible therewith.

Returning to the topic of companions-in-guilt responses, Brink's strategy is not the only one available to moral realists. If there are companions in guilt to be found for moral facts, they must actually share in the putatively queer feature, namely an apparently inexplicable yet necessary connection with certain natural facts. Brute connections like these may be more common than it initially appears. A moral claim is an assessment; it is concerned about values or reasons of a certain type. But the moral dimension is far from the only one along which things might be assessed. Other kinds of values or reasons (aesthetic, romantic, academic, et cetera) surely involve the same sort of relation, though, as they too turn on the natural features of that which is being assessed.¹⁴¹ For facts about other kinds of reasons and values to be workable companions in guilt, they must be relevantly similar to moral facts as we have understood them, which here means that the facts of the matter about these kinds of reasons and values must be non-natural facts. If values and reasons of these kinds can be fully naturalized, then there is nothing surprising or unusual about the supervenience relation between them and certain natural facts, for reasons discussed above. In order to resist this version of the queerness argument from supervenience, Moral Error Theorists would need to argue that moral facts alone have the kind of brute connections with which we are here concerned (or else that any actual companions-in-guilt still don't give the new kind of fact we have to posit sufficient explanatory force to be worthwhile). It seems

¹⁴⁰Mackie 1977, p. 48. Given that he was, in fact, a Moral Error Theorist, it should come as no surprise that Mackie (1982) rejected the existence of God and this kind of Divine Command Theory with it.

¹⁴¹I owe this companions-in-guilt argument to Tom Stoneham.

fairly likely to me that some such argument can be given in such cases, but it may be necessary to address each such potential companion in guilt individually, and I do not have the space here to do so, nor an exhaustive list of such candidates to be addressed. Suffice to say this companions-in-guilt response nominates much more plausible companions, and while it might still be defused, we are in no position to commit ourselves to the claims that we would need to make in order to do so without extensive further argument.

In conclusion, then, it seems like the queerness argument from supervenience has some merits, but it is not sufficiently strong to ground a full rejection of moral facts from our ontology. Moral facts do indeed supervene on natural* facts, and the way in which they do so raises intractable questions of the kind that motivate the Argument from Queerness, so long as we assume moral non-naturalism, which we are currently doing. However, the kind of supervenience is one we expect to see for any facts about non-natural values. Without ruling non-natural values out, we cannot say that the brute connections we would have to posit to account for moral facts cannot bear their own explanatory weight, as there might be any number of companions in guilt that call for the same move. The issues raised here may yet warrant a response (which I give in §3.2.2), but they don't raise such clear problems for moral realism that the error theoretical conclusion is irresistible.

1.8 The Queerness Argument from Special Normativity

The last queerness argument I shall examine is the queerness argument from special normativity. Like the queerness argument from motivation, this argument nominates moral facts' purported ability to give us reasons as queer. However, the reasons in question are of a distinctly different type; they are normative reasons.

Motivating reasons explain why an agent did, will, or would perform some particular action, and so they are facts that form part of the causal explanation of a matter of empirical fact (or, in the case of counterfactuals, would have been an empirical fact if the antecedent were true). By contrast, normative reasons are 'considerations in favour' of a particular course of action. They are the facts that

explain why we *should* follow a particular course of action, as distinct from why we would want to do so.¹⁴² Another way to understand them is as facts that make (in the sense discussed in §1.7) a specific course of action ‘worth doing’ or ‘a good idea’, though these are rough characterizations that won’t brook much analysis. We might also bear in mind how such reasons relate to ‘ought’ statements; one common analysis of the idea of ‘ought’—which I will be using—is as that which one has *most reason* to do.^{143,144}

The categories of motivating and normative reasons can overlap, but they don’t do so necessarily.¹⁴⁵ Suppose, for example, that I have a genuine normative reason to avoid catching scurvy. It follows that I have a reason to get plenty of vitamin C, and so a reason to eat more oranges.¹⁴⁶ I might add more oranges to my diet precisely because I don’t want to get scurvy, in which case my motivating reason

¹⁴²This distinction is a well-worn one, and is the same one described by Parfit (1997, p. 99), Dancy (2000, §1.1), and Olson (2014, pp. 105, 116), among others. Elsewhere, motivating reasons and normative reasons have been called explanatory and justificatory reasons respectively (Alvarez, 2017). Both, however, can be cashed out in terms of explanation; motivating reasons explain ‘is’ facts about human action while normative reasons explain ‘ought’ facts, so drawing the distinction in this way is better for our purposes.

¹⁴³This is the understanding employed by Parfit (2011b, pp. 1, 33) and I will take it as given. This is a ‘maximising’ conception of reasons, and rules out alternative view’s like Slote’s (1989, ch. 3) satisficing conception, which has it that one ought to do what is sufficient for one’s ends. There are positive reasons to favour Parfit’s account; if an agent, presented with two mutually exclusive and exhaustive options, performs the action that they have less reason to perform, then that agent has made a mistake. Also, see Rabinowicz 2008, p. 25 for independent arguments against the satisficing view of practical rationality.

¹⁴⁴Though I have called this an analysis, I am not committed to any particular relationship of ontological priority between reasons and oughts, just to their interdefinability.

¹⁴⁵Strictly speaking, when this happens the normative reason will be a fact, and the motivating reason will be a belief for which that fact is a truthmaker, but I take it as fairly intuitive in what sense these are the same consideration.

¹⁴⁶There may be questions about how normative reasons at different levels of specificity (like those just given) are related, but I will leave those to one side here.

is also a normative reason, but I might also do so because I discover that I'm much fonder of them than I remember, in which case I still have a normative reason, but it is not the motivating reason that explains my action. Likewise, I might have a false belief that raw meat is the best source of vitamin C, and so have a motivating reason to eat it, even though I might have normative reasons not to do so. This highlights one key difference between motivating and normative reasons. Normative reasons are facts, where motivating reasons are the contents of beliefs, taken as factual by the believer but only ever contingently corresponding to any arrangement of properties in the actual world. Whichever considerations are my motivating reasons are ones that I—rightly or wrongly—believe to be normative ones, and the ones to which I would appeal if I were asked to justify my actions (and inclined to answer honestly).

So we understand what normative reasons are. But the Moral Error Theorist does not simply claim that generating normative reasons makes moral facts queer. To do so they would have to endorse an error theory about normativity more generally, and claim that no one ever has a genuine reason to act in any particular way. This is a bold conclusion, though not one that has no supporters.¹⁴⁷ Generally, however, Moral Error Theorists don't make such strong claims. Rather, they say there is something queer about the normativity of moral facts in particular; that moral norms are strange or different in such a way as to make them ontologically suspicious. I call this putative feature of moral facts 'special normativity'. (Here and elsewhere in this project, by 'normativity' I mean status as a normative reason and by 'special normativity' I mean status as a specially normative reason.)

The claim that Moral Error Theorists are advancing here is that the kind of reason that moral facts are meant to be or provide are different enough from other normative reasons that it cannot be accounted for in the same way as less controversial norms. Thus we are forced to make an exception and posit a new quality to account for such norms. Moral Error Theorists disagree about exact nature of this difference—what makes this normativity special—but they are also, by my lights, all trying to characterise the same worrying feature, as should become apparent when we examine some of these accounts shortly. Between these examples

¹⁴⁷See Streumer 2017.

it should be apparent that there is indeed something idiosyncratic about moral normativity such that admitting them to our ontology will come at an explanatory cost, which is all we need for this queerness argument to succeed.

First, however, I should address the first of our two key questions: ‘Is special normativity an essential feature of moral facts?’ Or rather, I should address my failure to address it, as I will not be discussing it here. There is much to be said about the question of whether all moral facts must necessarily be specially normative reasons. So much, in fact, that we could devote an entire chapter to it, which is precisely what I intend to do. Because this pivotal matter will be covered in the next chapter, I will proceed here under the assumption that special normativity is an essential feature of moral facts, and the conclusions of our debate here will have to carry that proviso.

The answer to the second of our two key questions, whether the special normativity of moral facts is sufficient to make them queer, depends on what account we give of the difference between moral normativity and other forms of normativity. One such account, which we will treat as our main example because of the depth in which it has already been developed, comes from Jonas Olson. According to Olson, moral reasons are queer insofar as they are irreducibly normative. On this account, normativity in general is a reducible phenomenon. Typically, some agent’s having a reason to act in a certain way reduces to the fact that the action is likely to bring about the satisfaction of the agent’s desires.¹⁴⁸ For example, my reasons to eat more oranges as described above reduce straightforwardly to facts about my desires (for good health, or to experience citrusy flavours) and how they

¹⁴⁸Olson 2011, p. 78. This is offered as a reduction of hypothetical reasons in particular, and Olson stresses that it is not the same favouring relation that constitutes moral normativity. It is not intended as a universal recipe for the reduction of normativity, and in fact Olson offers no such thing (beyond the claim that certain conjunctions of non-normative fact are sufficient for the existence of a normative fact), as is made clear below. But his argument turns more on the claim that every normative fact is reducible rather than the claim that there is something *in particular* to which they all reduce, so this is not a problem. A more unified reductive account like that developed by Finlay (2014, chs. 1-3) may be preferable for independent reasons, but it is no better situated to support the queerness argument from special normativity.

can be fulfilled. Likewise, many of the reasons to which we are subject every day work in this way; they are nothing over and above facts about agents, their desires and the means that can be used to achieve them; the reduction relationship is a very straightforward one.¹⁴⁹ (This is not meant as a commitment to reasons internalism, the view that all normative reasons are ‘internal’ in that they depend on the desires of an agent.¹⁵⁰ There could be other plausible reduction bases for normative reasons. Olson gives the rules of a game or the demands of etiquette as examples for the reduction bases of conventional norms.)¹⁵¹ Under such a model the existence of norms imposes no ontological cost on our worldview because we already accept agents, means, ends, and other elements of the reduction base.

By contrast, irreducible normativity (as the name implies) cannot be reduced to any other domain of facts; an irreducibly normative reason is normative *simpliciter*; it’s normative all the way down. In other words, it just *is* normative. It is a common commitment of moral realists that moral norms are categorical, inescapable, or otherwise non-optional, which is why they must be accounted for with irreducible normativity. The reduction to which most norms are subject relativizes them to particular agents, goals, contexts, and so forth but a fact that just *is* a reason will be a reason from any perspective. Because this kind of normativity is not subject to any kind of reduction, it uniquely disrupts an otherwise adequate model of how normative reasons work, and forces us to posit a truly novel and distinct category of fact, thereby generating precisely the kind of burdens that are characteristic of queerness.

As I have already mentioned, Olson’s is not the only account of what is special about moral normativity. It is, I think, the best yet given for reasons that I will

¹⁴⁹Olson 2011, pp. 77-8. ‘Desires’ here is used in the broad philosophical sense to cover the objects of a wide variety of conative states as well as consciously adopted ends or goals.

¹⁵⁰See Williams 1979, pp. 25-8. It is worthy of note that this is not equivocation between motivating and normative reasons; some fact can be a normative reason even while failing to motivate on Williams’s view just so long as there is ‘a sound deliberative route’ from the agent’s motives to the action for which it is a reason (Scanlon, 2014, p. 11). This view is unrelated to Brink’s internalism discussed in §1.6.

¹⁵¹Olson 2014, p. 121.

make clear shortly. To do so, it will be helpful to cast our eyes over some other plausible candidates, including those put forth by earlier writers.

Mackie, on whose work Olson draws heavily, made queerness arguments targeting the ‘objective prescriptivity’ of moral facts.¹⁵² This is the way moral facts are inherently practical or make demands of us.¹⁵³ Neither objectivity nor prescriptivity is obviously problematic, so it must be only when prescriptivity itself is objective that we have something potentially queer.

If we want to know why this combination would be sufficient for queerness, the answers that appear in Mackie’s text are ones that should be familiar by this point. Firstly, he contends that moral facts (or ‘objective values’) are the only ones that would have this trait.¹⁵⁴ He also draws attention to the problems of grounding and supervenience (though not in so many words) and expresses none of the same concerns about hypothetical imperatives which, in contrast to categorical ones, are contingent on the desires or goals of the agents to whom they apply.¹⁵⁵ This issue, then, seems to be one of explicability; other kinds of prescription are unproblematic insofar as they can be explained while the prescriptivity of moral norms is a matter of brute fact.

Another thinker on similar lines is Joyce, who frames the issue with special normativity explicitly in terms of imperatives; moral judgements often take the form of categorical imperatives. Unlike other Moral Error Theorists, Joyce doesn’t claim that this is absolutely essential to moral facts, but he does think this structure is central to moral thought, which implicitly means it would be equally important for any real morality.¹⁵⁶ If categorical imperatives are a necessary condition for

¹⁵²To be more precise, Mackie attributed moral facts a queer power to provide reasons for action, but didn’t distinguish between normative and motivating reasons (Olson, 2014, p. 117). Those of his arguments that concern motivating reasons have already been discussed in §1.6, but there are others clearly targeting the normativity of moral facts, which are described below.

¹⁵³Mackie 1977, pp. 32, 40.

¹⁵⁴Mackie 1977, p. 38. A plausible interpretation of Mackie is that his use of the phrase ‘objective values’ includes both normative and evaluative facts.

¹⁵⁵Mackie 1977, pp. 41, 47.

¹⁵⁶Joyce 2006, p. 61.

morality, then showing them to be queer would ground Moral Error Theory quite well. When we look at the idea, though, we find ourselves in much the same position as with Mackie's ideas. It is meant to be the confluence of these factors—the fact that the imperatives they ground are categorical—which makes moral facts queer, for the same reasons Mackie invoked.

But Joyce actually discusses special normativity under a number of names; 'intrinsic action-guidingness', 'inescapable authority', and 'practical oomph' or 'clout' are all used more or less interchangeably.¹⁵⁷ All of these point vaguely in the direction of a normativity that applies universally and unconditionally to any and all agents, just as Mackie's account does, and just as many other Moral Error Theorists' versions of special normativity do.¹⁵⁸ The existence of a moral fact, then, implies a normative reason which applies to all agents without exception.¹⁵⁹ This is only sufficient for queerness if we take as given that there are no other reasons that work this way, and that this difference is sufficient to attach some explanatory cost to moral reasons in particular.

Both of these views are clearly reaching for the same idea, of moral normativity (which is to say special normativity) as something that genuinely applies to us in a way that is apparently unconnected to any goals or projects we might have, or to anything else to which we might normally appeal in order to ground claims about our reasons. Olson's account is doing this too, but its advantage over these earlier views is that it appeals to one feature of moral normativity—its irreducibility—to explain how moral reasons are supposed to have this power, and how moral facts are queer in the sense of generating explanatory burdens.

We can even use this model to account for other unique features of moral normativity that have not (to the best of my knowledge) been the locus of any version of the queerness argument from special normativity. Consider that almost any fact can be a normative reason, but its status as such is a matter of standing

¹⁵⁷Joyce 2006, pp. 57-64, 173, 192-9.

¹⁵⁸See, for example, Garner 1990, pp. 138-9, 143.

¹⁵⁹I take it that the relative strength of this reason can vary between agents based on their abilities and immediate circumstances, but in a way that is isomorphic with non-special normative reasons and so unfit to ground or undermine the attribution of queerness.

in a certain set of relations; some fact is a reason for a particular agent to perform a particular action. Since these are the variables involved in the relation, it is quite plausible that the relation holds in virtue of something about the agent and the action in question. Think back to the example of vitamin C: If I ask ‘Why should I eat more oranges?’, you might reply ‘Because they are a good source of vitamin C.’ This answer might satisfy me, but I think we would be hard pressed to call it complete. That is to say, there must be more to the explanation than this.

Imagine, for a moment, that I am a space alien for whom vitamin C is not an essential nutrient but a deadly poison. I think it is fair to say that, under such circumstances, the fact that oranges are rich in vitamin C is not a reason for me to eat them. So the fact that it is a reason in real life, where I’m an ordinary human being, is dependent on or explained by other facts, namely facts about how vitamin C interacts with ordinary human bodies.

Moral reasons are still relational in the sense that they attach to agents and actions, but facts about the agent are left out of the equation (or, at least, far fewer of them are relevant). That is, the relation does not seem to hold in virtue of any facts about at least one of the relata, but rather to attach to any and all things of the relevant kind (agents). This is another way in which moral norms are exceptional, and another one which we would expect to see on Olson’s model. This is because, if non-moral reasons hold in virtue of some pre-existing network of facts, we can account for this reductively, with the exception of special norms, which is to say moral ones.

Given all of this, one might start to wonder why I have chosen to frame this queerness argument around special normativity rather than irreducible normativity, since the latter can account for so many of the ways in which moral normativity is ostensibly queer. There are two related reasons to do this. Firstly, while Olson’s account effectively encapsulates the thoughts of some other Moral Error Theorists, and is the best version of the queerness argument from special normativity, it is still possible that a functioning version might be given based on some other feature of moral normativity. Secondly (and this matters because of the first reason), the response to the Argument from Queerness that I will go on to develop in the remainder of this project cuts across any such differences within the queerness argument from special normativity. That is to say, assuming that the queerness of

moral facts is located in their normativity because there is something unique and problematic thereabout, it does not matter for the purposes of my counterargument what that unique and problematic element is. It works equally well against all versions of the Argument from the Queerness of Special Normativity. Going forward, I will treat Olson's version as paradigmatic, but the fact remains that if we were to develop a version from another aspect of moral normativity, the view I go on to develop would be a counter to that as well.

For example, one might appeal to the strength of moral reasons, rather than (or in addition to) their scope. While it would be a stretch to say that moral reasons must always override other reasons, there is nevertheless an intuition that, all else being equal, moral considerations provide stronger reasons than aesthetic ones, for example. Even if we were to grant that their categorical scope were unproblematic, that wouldn't explain why they are to take precedence within the decision procedures of individual agents (who presumably have at least some competing reasons of various kinds).

This is not a complete account of special normativity. I have simply pointed to a feature of moral norms that is plausibly ubiquitous across moral facts and would raise some questions. It's entirely possible that these questions have answers such that there is nothing unusual or surprising about the strength of moral reasons, but by the same token it's possible that these observations could serve as the basis for a full account of the queerness of moral normativity.

That said, I will assume in future chapters that special norms fit into the broad paradigm established by Mackie, Joyce, Olson, and other Moral Error Theorists, since it is the most promising account available to us. A moral realist wishing to rebut the queerness argument from special normativity might at this juncture object to the claim that there is, in fact, anything special about moral normativity. Since we are treating Olson's account of irreducible norms as special, and thus queer, let us consider how it might potentially be undermined. One might, as many philosophers have done, call into question the connection Olson posits between desires and reasons. However, this particular element might easily be jettisoned, as long as we can offer another plausible reduction base for those reasons that we accept as genuine. Facts about what would promote an agent's objective wellbeing might be able to fill this role, for example. The view would then retain the virtue

of parsimony, in that it would still explain our non-moral reasons without invoking anything potentially queer.

A more promising strategy of attack is to call into question the difference that Olson postulates as the basis for his argument, namely that moral norms must be irreducible (to explain their suite of unique features) while other examples of normativity need not be. A way to do this would be to challenge the idea that there are unproblematic norms that can be accounted for reductively. One might be inclined to ask in what sense Olson's reducible norms are normative, when they are nothing over and above facts about agents, means, and ends. I do not think, though, there is any answer to give except that it is normative in a reducible way. What it is to be a normative reason of any kind on this account is to be a fact about how a given agent can satisfy certain ends or standards (and so they are contingent on the end being one that the agent has, or the standard being one to which the agent is subject). As far as an Olsonian Moral Error Theorist is concerned, they are also the only kind of norms that actually exist.

There is a separate question of whether we should, indeed, admit that any of our reasons might be subject to this sort of reduction. When Olson says that a reason is reducible, what he means is that the favouring relation that makes it a reason—it's standing in favour of a specific action—can be reduced, which is specifically denied by thinkers like Scanlon, who claims that this relation is unanalysable.¹⁶⁰ Olson claims that the question of whether there is anything queer about irreducible norms is 'metaphysical bedrock'.¹⁶¹ But why should we accept this appeal to intuition?

Certainly there are independent reasons to favour a reductive approach to normativity. By reducing normative properties and facts to non-normative ones, it is more parsimonious not just with regards to the total number of kinds of thing we admit to our ontology, but also the number of sources of knowledge and kinds of logic we need to accept, which allows it to avoid a wide array of puzzles and controversies.¹⁶² Of course, we can rely on parsimony as a principle of theory

¹⁶⁰Scanlon 2014, p. 44.

¹⁶¹Olson 2014, p. 136.

¹⁶²Finlay 2014, p. 16.

choice only when all other things are equal, and the realist may simply insist that they are not because reductive accounts always leave out something essential.¹⁶³ This is the bedrock issue to which Olson referred, and it really might be one that leaves us at a stalemate of intuitions, unless the realist (or, specifically here, the primitivist about normativity) can give a strong argument for their position.

Primitivists could point to a reason that they take to lack a plausible reduction base, but this strategy would rely on their intuitions that those considerations are reasons, which Moral Error Theorists are unlikely to share. Historically, primitivists have appealed to arguments along the lines of Moore's open-question argument.¹⁶⁴ This, however, also falls short in a couple of key ways. Firstly, the intended conclusion of that argument is only that reductive accounts fall short, which means it underwrites antirealism about norms just as much as primitivism (though Moral Error Theorists share realists' commitment to cognitivism, so this is not so serious a flaw for our purposes).¹⁶⁵ Secondly, this argument relies on the apparent nonequivalence of moral properties and any prospective reduction base, but has nothing to say about synthetic or non-obvious reductions.¹⁶⁶ This latter point is particularly salient, and its force will be made much clearer with a little more detail about the open-question argument.

The open-question argument starts with the observation that for any supposed naturalistic reduction of a normative property, we can find something that has all of the required natural properties, and still reasonably ask whether it also has the normative property. For example, it does not seem to be a settled question whether pleasure is good. From there, Moore argues that—regardless of whether utilitarianism is true—goodness and pleasure are not one and the same, because if

¹⁶³Note that reductive accounts of normativity needn't necessarily be naturalistic. Raz's (1999b, p. 22) account of reasons is that they are ways in which an action would be good, which is a reduction of sorts, and I see no reason why one couldn't accept this along with the non-naturalism brought in at the start of this chapter.

¹⁶⁴Scanlon 1998, pp. 58, 60.

¹⁶⁵Finlay 2014, pp. 11-2.

¹⁶⁶See Durrant 1970, p. 360; Boyd 1988; Brink 1989, pp. 144-70; Schroeder 2005, p. 23; and Finlay 2014, p. 12.

they were this question would be closed. This argument works off of an assumption that, given the necessary nature of identity relationships, the open question should not seem legitimate if goodness were identical with some particular natural property. But just because identity claims, when true, are necessarily true, this doesn't mean they are obviously true or knowable a priori.¹⁶⁷ To borrow a famous example, the planet Venus is identical with the (apparent) stars Phosphorus and Hesperus, but this is something that we had to learn empirically. Nevertheless, there is no possible world in which the planet Venus can have failed to be the planet Venus.¹⁶⁸ If necessary identity relationships can be non-obvious, then pointing out that the question of whether goodness is identical to any natural property are open tells us very little, and certainly not that they cannot possibly be identical.¹⁶⁹ Thus, reliance on this kind of argument cannot justify commitment to primitivism.

So, all else being equal, we have some reason to favour reductive accounts over alternatives. The natural question is then whether all else actually is equal, whether a reductive analysis can capture everything about normativity. The dialectic on this so far seems to come down to a battle of intuitions, about which all I can say is that I share Olson's.¹⁷⁰ (Though it is also perhaps worthy of observation that, in contrast to reductive explanations, intuitions of irreducibility do not have a sterling track record in fields where final consensus is more common than in philosophy.)¹⁷¹ The balance of reasons, then, seems to be in favour of a reductive account.

¹⁶⁷Kripke 1980.

¹⁶⁸There are, of course, possible worlds where the three astral bodies astronomers observed turned out to be three distinct objects, but none of those are identical with—or even really a counterpart to—our Venus.

¹⁶⁹This observation comes from Putnam (1981, p. 208), couched in terms of words and how their non-synonymy implies nothing about the properties they denote.

¹⁷⁰We might cash out the difference in intuition in terms of whether one thinks there is an explanatory gap in a reductive account of norms, though one might grant that such a gap exists and still be a reductionist; see Mehta 2017.

¹⁷¹Finlay (2014, p. 13) cites both classical elements and the view that atoms are fundamental and indivisible as examples of this.

Companions-in-guilt arguments of the kind discussed in §1.4 are a good candidate for something else to which a would-be moral realist might appeal. If we already admit of other irreducible norms, then there is no new kind of fact or exception to our account of reasons to impose explanatory costs on the admission of moral facts and so they will not be queer. The question, then, is whether there are other forms of normativity that have the unconditional scope of moral reasons (which is suggestive of their irreducibility) that are widely accepted.

One strong contender is epistemology.¹⁷² Reasons for belief are often seen as universal in the same way as moral reasons; everyone has a reason to believe whatever is suggested by evidence. Norms of this kind are less controversial than moral ones, at least among non-philosophers, and they seem to possess the same kind of unconditional reach. Whether this is actually the same kind of normativity possessed by moral facts will depend on how we cash out special normativity, though it certainly appears to exhibit the unconditionality or inescapability that would be explained by the irreducibility of moral facts. But just because some features of epistemic norms would be explained by their being irreducible, it does not follow that we should consider them to be such. If a reductive account is available, it is to be preferred for reasons of qualitative parsimony and to unify our understanding of normative reasons. So the real question is whether we can give any such account. I think we can.

True beliefs are very useful in that whatever reasons an agent otherwise possesses, that agent will be better positioned to act on those reasons if their beliefs about relevant matters are true rather than false. Hence one has reason to believe in a truth sensitive way (that is, to follow epistemic norms) at least insofar as one has any other normative reasons; epistemic reasons reduce to other reasons and, in turn, to the reduction bases of those reasons.¹⁷³ One might even argue that insofar

¹⁷²Olson 2014, pp. 138, 155. See Scanlon 1998, ch. 1, Cuneo 2007 and Bedke 2009, pp. 41, 55-7 for arguments to the effect that epistemic norms are companions in guilt with moral norms.

¹⁷³There are a couple of views for which this strategy is ineffective, namely a general error theory about normative truth (in which case explaining epistemic normativity is no longer a concern) and a view whereby all of the reasons possessed by actual agents relate solely to their inner mental life and none relate to the external world in any way (which is implausible for

as an agent's reasons are subject to change, they have reason to believe according to epistemic norms in general such that they will always be positioned to act on whatever reasons they come to possess. Another popular candidate are the rules of deductive logic, but the very same reasons to seek truth over falsity can explain their universality as a standard of reasoning.¹⁷⁴ As long as we are not forced to account for putative companions in guilt as specially normative, these arguments are insufficient to undermine the queerness argument from special normativity.

We might think that appealing to such arguments opens up the possibility of making a similar claim about morality, that it too rests on some universal or near universal desire. I'm actually open to such a picture, but invoking it here is tantamount to rejecting one of our starting assumptions in this section; that the normativity of moral facts is special normativity, which applies to all agents no matter what. I will ultimately reject this assumption in Chapter 2, and discuss views like this in §4.2 and §4.3, once I have put aside the assumptions that make them untenable.

One final point on companions in guilt is that finding one, or even several, wouldn't necessarily guarantee that special normativity isn't queer. In at least some cases, Moral Error Theorists are willing to discard companions in guilt along with moral facts.¹⁷⁵ Only if a putative companion in guilt cannot be explained with reference to ontological categories that we already accept and is itself such that rejecting it from our ontology would carry a substantial theoretical cost, will a companions-in-guilt argument be successful. Appealing to special normativity will violate principles of parsimony, unless attempts to account for normativity without a special irreducibly or inherently normative category are themselves ontologically profligate. In practice, however, this isn't what we see. Thinkers like Olson and Bernard Williams give fairly unified accounts of normativity, wherein only a small

independent reasons).

¹⁷⁴The example of deductive logic as a candidate companion in guilt is a move by Olson (2014, p. 137) but he dissolves the threat by relativizing it to the standard of correctness in reasoning. Relativizing the normative power of such rules to the same thing as more general epistemic rules is my own move, and a more parsimonious one.

¹⁷⁵See Mackie 1977, p. 39 and Olson 2014, §8.2.

number of putative outliers call for us to invoke special normativity.

So it looks like an appeal to companions-in-guilt arguments cannot be counted on to save moral facts from the queerness argument from special normativity. The logic of the argument itself is solid, but there may be no one surefire version of it because each version depends on an understanding of when and how we come to have normative reasons, none of which can be definitively proven over the others. Even if we, say, use thought experiments to test our own intuitions, this only refines our own application of the concept, and someone who had contrary intuitions about which reasons people have could never be strictly proven wrong. Still, even understood in all these different ways, the special normativity of moral facts, in its specialness alone, gives us genuine reason to believe that moral facts are queer. Olson offers us a broadly applicable account of special normativity as irreducible which encompasses many of the the features other Moral Error Theorists have highlighted, so I will treat his version as paradigmatic and the remainder of this project will be concerned with what one must be committed to having accepted it.

What we have discovered in this chapter is that, taken together with the assumption of cognitivism and non-naturalism, a debunking explanation of moral phenomena and a successful queerness argument give us grounds to reject moral facts from our ontology. We have also learned that the queerness argument that poses the most serious threat to moral realism is the queerness argument from special normativity. Whether or not one finds this argument persuasive, it is the most widely endorsed, and the remainder of this project will have important implications regarding what Moral Error Theorists can safely conclude off the back of it. As such, this project will henceforth be focused primarily on developing a response to the Argument from the Queerness of Special Normativity. This version of the argument depends on the assumption that special normativity is an essential feature of moral facts. In the next chapter, I will examine and challenge this claim, which will in turn set the groundwork for challenging the Moral Error Theorist's understanding of moral facts in general.

Chapter 2

Normativism and Non-Normativism

In the last chapter, I concluded that the Argument from Queerness gives us a strong case for rejecting moral realism, just so long as we accept certain assumptions. Two of these are common across all variants of the argument. These are cognitivism, which we shall accept, and non-naturalism, which I shall address briefly in Chapter 3. Other assumptions only come into play in some of versions of the Argument from Queerness. Much of the last chapter consisted in identifying these assumptions and subjecting them to scrutiny. One prominent example of an assumption that survived this examination is that moral facts supervene on natural (or supernatural) facts about the world, and that this relationship is necessary so that nothing that is wrong in this world is right in any other possible world.

In one case, however, I left a key assumption unexamined. This was one of the assumptions underlying the Argument from the Queerness of Special Normativity, to the effect that all moral facts are (or provide) specially normative reasons for action. We shall examine that assumption in this chapter.

2.1 The Normativity of Moral Facts

Olson calls the claim that all moral facts are specially normative reasons ‘the conceptual claim’.¹⁷⁶ So it’s meant to be conceptually true that moral facts have special normativity as a feature; doing so is just part of what it is to be a moral fact and nothing that isn’t a specially normative reason could rightly be called a moral fact. I will be discussing a number of rival conceptual claims in this chapter, so it will be helpful to give each its own name. We can call this one:

Strong Normativism All kinds of moral facts entail specially normative facts.¹⁷⁷

I should specify for clarity that when I speak of normative facts here, I mean facts that attribute some agent a normative reason to act. When expressed propositionally, they take forms such as ‘Some agent *A* *ought/oughtn’t* to ϕ ’, ‘*A* has *normative reason(s)* to ϕ ’ or ‘ ϕ -ing is *right/wrong*’. Specially normative facts, then, attribute specially normative reasons that are meant to apply objectively to the agents in question without depending on facts about those agents’ desires or any conventions that apply to them (or any other plausible reduction base).

Strong Normativism is the view to which most Moral Error Theorists are committed. Strictly speaking, this view is Normativism about morality. There could be normativist views about any putatively normative domain of discourse (such as epistemology), and we could describe them by swapping out ‘moral’ for the appropriate term. However, our concern is with moral facts, and so when I use the word ‘Normativism’ without qualification, I mean Normativism about morality.

The reason I have called this view ‘Strong Normativism’ is to distinguish it from the following weaker view:

¹⁷⁶Olson 2014, p. 124.

¹⁷⁷Brink (1989, p. 40) calls Normativism ‘internalism about reasons’ (and its denial ‘externalism about reasons’). I have avoided this terminology to avoid confusion with other internalism/externalism distinctions within philosophy, and especially with Williams’s reasons internalism, which I discussed briefly in the last chapter. Brink also offers a strong/weak distinction within these views, but does so along different lines, so using different names clearly distinguishes between views we discuss.

Weak Normativism Some kinds of moral facts entail specially normative facts.¹⁷⁸

This is the view that special normativity is a conceptual feature of morality, but not necessarily a feature of every individual moral fact.

Unlike Strong Normativism, Weak Normativism is insufficient to ground the Argument from the Queerness of Special Normativity. Weak Normativism amounts to the claim that, as a matter of conceptual necessity, there is at least one kind of moral facts whose members each entail at least one specially normative fact. It remains neutral on whether there is any kind whose members lack this feature. It also remains neutral on whether the specially normative kind(s) it posits actually have any members.¹⁷⁹ That is to say, they might be empty even if Weak Normativism is true. As such, this view doesn't necessitate the actual existence of any specially normative facts, which are necessary for the Argument from the Queerness of Special Normativity to get off the ground.

I freely grant that both these views have a certain level of intuitive appeal. After all, thinking of moral facts as categorical imperatives (which aren't contingent on our own projects or desires or goals) dates back at least to Kant, and deontological accounts of ethics are framed entirely in terms of what we ought—what we have most reason—to do. Many philosophers even think that ethics is founded on questions like 'How should I live?' and 'What ought I to do?', which are fundamentally questions about the reasons we have. Nevertheless, given that I have already argued that special normativity is indeed queer, anyone who wishes to resist the Argument from Queerness must deny Normativism. Denying Normativism could

¹⁷⁸Here and in other definitions, I am using the reading of the particular quantifier 'some' favoured by Priest (2008, pp. 42-3) where it is not taken to have existential implications. As such, we can talk about kinds of moral facts without presupposing that there actually are any moral facts of the kinds described, or getting into debates about whether empty kinds are possible. Some kinds of dragons have four legs, and some have two, but to say as much does not commit us to the existence of any dragons.

¹⁷⁹It is in this way different from the view expressed by Joyce (2006, pp. 60-1), who takes it that if there are any moral facts, at least some of the ones that actually exist entail normative facts.

lead to one of two distinct views. Firstly there is:

Weak Non-Normativism Some kinds of moral facts do not entail specially normative facts.

This is, more specifically, the denial of Strong Normativism. At the same time, it's compatible with Weak Normativism, because the existence of at least one class of moral facts that are specially normative reasons is in no way contradictory to the existence of at least one class of moral facts that are not. If we instead deny Weak Normativism we get:

Strong Non-Normativism No kinds of moral facts entail specially normative facts.

Whereas the two weak views are compatible, the strong version of each view is incompatible with either version of the opposite view. Strong Normativism is incompatible with any form of Non-Normativism, and Strong Non-Normativism is incompatible with any form of Normativism.

Strong Non-Normativism is a strong claim both in the logical sense of committing one who makes it to more than a corresponding weak claim, and also in the sense that it is a bold and surprising contention. It says that it is conceptually true of moral facts that they never have special normativity as a feature. This is counterintuitive, especially given that the debunking argument I sketched out in Chapter 1 explicitly claims that we do much of our moral thinking in terms of categorical imperatives, which is to say in terms of special norms. Since conceptual analysis is done by examining how we think about the concept in question, the idea that moral facts *don't* have a structure that we frequently think of them as having isn't very plausible as a conceptual claim. In Chapter 3, I will argue that *as a matter of fact* there might be no moral facts that entail specially normative reasons, but this is an ontological claim, not a conceptual one. If we had to endorse Strong Non-Normativism to overcome the Argument from Queerness, we might be up the philosophical creek without an argumentative paddle.

Fortunately for those of us who wish to resist the Argument from Queerness, we don't need to endorse Strong Non-Normativism to undermine the queerness argument from special normativity. We need only make a much more modest

commitment to Weak Non-Normativism. This is because, while any amount of special normativity would render the moral domain suspect, the claim that (conceptually) at least some moral facts do not entail specially normative reasons is totally compatible with the claim that (actually) none of them do.

But we cannot just commit ourselves to Weak Non-Normativism tout court. Normativism clearly has many adherents, and much of the work done in substantive ethics seems to presuppose it.¹⁸⁰ As such I will need to argue that Weak Non-Normativism is itself plausible. I will not go so far as to argue that it is the only live option, but I will offer a robust defence.

One common argument against Weak Non-Normativism (and so in favour of its negation, Strong Normativism) is the normativity objection. This is simply the claim that any account of moral facts that doesn't characterize them as normative in some way has missed something essential about moral facts.¹⁸¹ The objection is that normativity is essential to moral facts. But any normativity that inheres only in one fact and is in no way contingent on, for example, facts about the agents to whom moral reasons apply would be special normativity. In other words, the normativity objection is just a restatement of Strong Normativism, meaning that this response assumes the point at issue against Non-Normativism.

Nevertheless, some people will feel that Strong Normativism and the normativity objection have some intuitive force. So I should say something in support of Non-Normativism. The first step toward this must be to provide a viable candidate for a class of moral facts that doesn't entail specially normative reasons. So

¹⁸⁰The fact that substantive ethics is often called 'normative ethics' attests to this. Since that name seems to presuppose some form of Normativism, I will always use the terms 'substantive ethics' or 'first-order ethics' to speak of philosophy concerning what moral facts there are, or which things have which moral properties.

¹⁸¹See Parfit 2011a, pp. 324-7 and Olson 2014, pp. 82-3. What this argument strictly shows is that no purely descriptive natural or supernatural fact (about, say, an action's propensity to cause pain or accordance with the will of God) is an explanation of any normative fact. As such this is more directly a response to moral (super)naturalism, and founded on the assumption that there must be a normative element to any moral fact, which is a problem for the reasons I detail below.

far, I have spoken fairly abstractly about ‘kinds of moral facts’ but all this talk of Normativism and Non-Normativism will turn out quite moot if I can give no account of what I mean when I talk about moral facts that don’t entail specially normative facts. It turns out we don’t need to look far, though, as the normative domain is already half of a well established dichotomy within ethics.

Moral philosophers often contrast the normative (which we have already discussed) with the evaluative. Evaluative claims are claims about whether certain things (actions, states of affairs, etc.) are good or bad, and in what ways and to what extent they are so. Likewise, evaluative facts are about the distribution of goodness, badness, and other evaluative properties. Consider the following two claims:

Murder is (morally) bad.

I ought (morally) not to commit murder.

Philosophers and laypeople alike will say that these claims don’t mean the same thing; they express two different ideas. There has been much debate over whether and how these domains are connected, which we will go into later, but what matters for our current discussion is this: if these sentences genuinely express two different claims, then it is possible in principle that their truth values could differ, which is to say that one could be true while the other is false. I don’t think this is a strange claim to make. Consider a utilitarian faced with a choice between one action which will generate a single unit of happiness, and another action which will generate two units of happiness. Based on their axiology, they would be right to call both courses of action ‘good’ but only the latter is the one they ought to pursue.¹⁸²

If an evaluative claim can be true while an apparently corresponding normative claim is false, this implies that whatever relationship might hold between these categories, it is not one of necessary entailment. That is to say, evaluative moral facts are a good candidate for moral facts that aren’t specially normative reasons

¹⁸²Depending on how the debate between holism and atomism about reasons (see §2.3.6), we might even say that there is not just more reason to pursue the latter, but no reason to pursue the former.

just so long as they are really separable rather than one being entailed by or contained in the other.¹⁸³

In other words, if we can have an evaluative claim that is true while its corresponding normative claim is false, then we must grant some form of Non-Normativism.¹⁸⁴ This is because this combination of truth conditions entails the existence of a class of moral facts that do not entail specially normative facts.¹⁸⁵ We have already seen examples where there seem to be evaluative facts at play without our needing to posit any normative ones, so we also have reason to think just such a combination is possible.

I have said much about what will happen if values and norms are not connected by necessary entailment, and this must at least plausibly be the case if we are to construct a version of moral realism which escapes the Argument from the Queerness of Special Normativity. In §2.2 and §2.3, I will address various arguments to the effect that values do, in fact, necessarily entail norms. The failure of these arguments provides a sort of negative case for the plausibility of the claim that values need not entail norms, but perhaps a negative case alone is not enough. As such, it will be worthwhile for me to provide some positive reason to buy into said claim. I do not pretend that any of these arguments are incontrovertible, but they are strongly suggestive—enough that they constitute evidence that values do not necessarily entail norms.

These arguments are all based in linguistic intuitions, which should be a rea-

¹⁸³There are philosophers who subsume the evaluative under the normative (in a general sense, which includes but needn't be limited to the special sense with which I am concerned) such as Thompson (2008). Likewise, there are philosophers who subsume the normative under the evaluative, such as Huemer (2009). On the former case, for example, evaluative facts are just a special case of normative facts, so the existence of an evaluative fact entails the existence of a normative fact via an identity relation. Not all views that posit an entailment relation between value and normativity predicate it on such a subsumption, but they represent a significant subclass of such views.

¹⁸⁴Normativists can and will differ on what they take the corresponding normative claim or claims to be, but this does not matter for our purposes.

¹⁸⁵As noted previously, this holds whether or not there are any such facts. See note 178 above.

sonable enough starting point, at least by a typical Moral Error Theorist's lights. As has already been mentioned, Moral Error Theorists appeal to the intentions behind ordinary speakers' use of moral language to justify some of the assumptions of the Argument from Queerness, since they take it that the meaning of moral terms is determined by those intentions. We will use the same method to call the assumption of Strong Normativism into question.

First, consider what seems to follow from the sample claims above. The latter claim is a deontic claim—a claim about what someone ought to do—and so a claim about the balance of all their moral reasons and so its truth would directly entail the existence of some reasons for action. The same doesn't hold for the former claim, which is evaluative, in that it appears to simply assign a moral property to a class of actions.¹⁸⁶ (Even a Normativist should grant that any connection we suppose to be present is less *obvious* in the former evaluative claim. The question of whether evaluative truths are just normative truths under another guise will be addressed in §2.2 and §2.3.) If normative facts necessarily entail the existence of reasons and evaluative facts do not necessarily entail the existence of reasons—both plausible claims—then it follows straightforwardly that there is a substantive difference between them.

Secondly, just so long as ordinary speakers would assent to the claim that my example sentences above express two different thoughts—and I believe that they would—then the distinction between norms and values is part of ordinary moral language. If it is part of moral language then it is also part our moral concepts. This is less a point in favour of Non-Normativism per se and more a point in favour of the weak views discussed above over the strong ones, since it would mean that Weak Normativism and Weak Non-Normativism (each of which can accommodate facts of both kinds) are more compatible with the meaning of moral terms than either of the strong views.

Finally, we can illustrate the difference between these categories by way of

¹⁸⁶Kirchin 2013, p. 180-3. Kirchin draws the distinction in question between evaluative and directive claims, both of which he subsumes under the normative to match the terms employed by those to whom he is responding, but draws it along the same lines as I have employed, so this is just a terminological difference.

another example where they seem not to coincide. Suppose we have a pair of universes, called Inferno and Paradiso, each of which contains a single mind and nothing else. Each is causally isolated, like a Lewisian possible world.¹⁸⁷ These minds are such that they will have a determinate series of experiences for the duration of their respective universe's existence and will do so entirely passively; they are subjects, but not agents. The subject in Paradiso will experience a life of joy and fulfilment (the specifics don't matter here), while the subject in Inferno will endure a wide variety of horrible torments and privations and nothing else. I think it is very intuitive to say that Paradiso is good and Inferno is bad (and that, consequently, Paradiso is better than Inferno). These are all evaluative facts. It is not nearly so intuitive to think that we have any reasons in relation to these universes. Since we cannot affect either of these agents in any way (this is true *ex hypothesi*), there can be nothing we ought to do to or for them, which is to say no overall reason, and so no contributing reasons. Hence, we have a case where evaluative facts fail to coincide with any related normative facts, which again points us at the cogency of this distinction. A Normativist might counter that we have reasons to hold certain attitudes toward each universe, but I think it would be premature to ascribe a rational failure to someone who had no attitude toward something they cannot possibly affect in even a minute or indirect way.¹⁸⁸

These are fairly basic arguments, and the force of each is dependent to a greater or lesser degree on whether one shares my linguistic intuitions, but they do provide enough positive support for the lack of entailment between value and normativity to make it plausible in the absence of defeaters.

I have focused exclusively on these two categories because, at time of writing, I can think of no other candidates for a non-normative kind of moral fact besides evaluative facts. A few candidates present themselves, but all of them seem to me

¹⁸⁷See Lewis 1986, §1.6.

¹⁸⁸Another example that makes the same point is Bykvist's (2009, p. 5) example of a world where there are happy egrets but no past, present, or future agents, which was originally used in an argument against buck-passing accounts of value of the kind discussed in §2.2.4 and §2.3. Such a state of affairs would be to some extent good, but no-one can have reasons to bring it about since it is incompatible with the existence of agents, and thus reasons.

to be such that we could account for them without going beyond the established categories of normative and evaluative facts. For instance, facts about what is just (or unjust) or what is sacred (or profane) might each initially look like a distinct and additive category of moral facts, but upon consideration both tie into the normative in direct and intuitive ways; we speak often of the demands of justice and say that injustices ought to be redressed. Sacredness too might be thought to give us normative reasons, or to change the evaluative status of actions involving that which is sacred (an otherwise neutral or very slightly bad act may become an enormity if a sacred object is involved, for example). Another possibility is facts concerning virtue, but these can either be understood as encompassing both the normative and the evaluative or as a subtype of the latter (see §4.1.2). One final candidate is the category of the supererogative, which is to say actions above or beyond the call of duty. Not all substantive ethical systems admit of supererogation (maximizing utilitarianism, for example, can make no sense of something better than what is morally required) but the fact that any can makes it worthy of consideration. That said, the supererogatory too is easily accounted for in terms of values and norms. A course of action that goes beyond one's duties is simply one that includes doing whatever one has most or overall reason to do (in other words, what one ought to do) and is also good in some other way not necessary in living up to one's duties. In brief, it holds a certain normative status and a certain evaluative status at the same time.

None of this implies that there definitely are no other non-normative categories of moral facts. Lacking any actual candidates for a third kind of moral fact, however, we can proceed as though normative facts and evaluative facts exhaust the possibility space of the moral domain. Once we have done so, the difference between values and norms seems to be the best grounds on which to argue for any form of Non-Normativism.

This brings the task of the remainder of this chapter into focus. When we began, I had already conceded that special normativity was genuinely queer. The only way to circumvent the Argument from the Queerness of Special Normativity, then, is to embrace a realist metaethics which doesn't admit of special normativity. I will say more about what such a view would look like in Chapter 3, but whatever form it takes, such a view is incompatible with any form of Normativism. As such,

if we are to endorse it, we must also endorse Non-Normativism, which itself is only plausible if there is a non-normative kind of moral fact. I have offered evaluative facts as a candidate. However, the sharp distinction between evaluative facts and normative ones may be challenged.

Many philosophers argue that there is some kind of necessary connection between these kinds of facts. If, as I am assuming, these two kinds are exhaustive of moral facts, then that position entails that Strong Normativism holds and neither form of Non-Normativism can be true. As such, the rest of the chapter will examine views about what manner of connection, if any, holds between evaluative and normative moral facts, and what implications this has for the debate between Normativists and Non-Normativists of various stripes.

2.2 The Connection between Norms and Values

If Normativists want to maintain that there is some entailment relationship from values to norms, then they must be able to describe the nature of that connection. They cannot simply declare that it is a necessary truth. Or rather, they can, but this would seem like an ad hoc response to the Non-Normativist, and leave the connection itself utterly mysterious. Such a claim would run afoul of the Modest Humean Dictum, since it posits a brute necessary connection between two discontinuous entities.¹⁸⁹ Such connections are a cost or downside to a theory for reasons discussed in §1.7.

Rather than a flat declaration, then, Normativists would be well-served by appealing to some existing metaethical view which holds that value necessarily entails normativity. We can call such views entailment views. The bald statement that I just discussed is an example of such a view, albeit not a very good one (as I hope I made clear). But there are other understandings of normativity and value that involve such an entailment relationship, yet are significantly more sophisticated in that they tell us what kind of relationship underlies that entailment. In short, they tell us *why* values entail norms. So far I have spoken only of value entailing norma-

¹⁸⁹This claim would also contravene Hume's Dictum, but in light of the discussion in §1.7, that would be shaky grounds to reject it.

tivity and not the reverse, and this is because neither form of Non-Normativism needs to posit any norms whatsoever, meaning that whatever they might entail is irrelevant. By contrast, both forms of Non-Normativism claim that it's possible for there to be values but no norms, which would be impossible if values themselves implied the existence of norms. The rest of this section will address some entailment views and whether we should be moved by a Normativist's appeal to them.

2.2.1 Moorean Consequentialism as a Conceptual Truth

One well-known example of an entailment view is the Moorean view that moral consequentialism is a conceptual truth. He claimed that statements like 'I ought to ϕ ' and ' ϕ -ing is the action I can perform which is most conducive to goodness' were synonymous, that they expressed the same proposition and their truth or falsity depended on the same facts.¹⁹⁰ If this relationship were anything but conceptual, it would involve brute supervenience of the kind that the Argument from the Queerness of Supervenience attacks.

Many have pointed out that Moore's suggestion doesn't seem to reflect the actual use of moral terms by normal speakers.¹⁹¹ By the standards of many Moral Error Theorists, this view has therefore already failed, but the view I will ultimately defend also demands an understanding of moral facts that doesn't straightforwardly reflect all of our moral language, so I cannot dismiss this view wholly on those grounds. Perhaps I could if it were a *tout court* declaration, but Moore offers an argument for his position, which we must examine.

An action that is uniquely what we ought to do is so, he asserts, because it is unique in value. This couldn't plausibly mean that it is the only valuable thing, nor that it is the most valuable, so he says that it must mean that the world as a whole would be better for its being performed. But this argument falls flat for two reasons.¹⁹² Firstly, the initial premise is one that non-consequentialists will deny;

¹⁹⁰Moore 1903, p. 147.

¹⁹¹See Brandt 1959, pp. 357-8 and Williams 1985, p. 16.

¹⁹²Carlson 1995, pp. 44-5.

it is central to such views that what we ought to do isn't necessarily the act that promotes most value, which would imply that at least some value facts don't have the normative implications Moore is positing. (This is not to say deontologists cannot hold entailment views, just not this one; see §2.2.2.) But there's a case to be made that insofar as it *is* what we ought to do, it is better than the alternatives. Even if we accept that, though, there's no reason we must think that these three interpretations of uniqueness in value that I have just described are exhaustive. We can if we assume that an action's moral status is determined by how it affects the goodness of the world as a whole, but to do so assumes the point at issue against non-consequentialists.

Later, Moore advanced the view that claims about duties and about value aren't identical, but are logically equivalent.¹⁹³ This is still an entailment view, since any proposition entails everything that is logically equivalent to it. Hence it still undermines Non-Normativism if true. Moore justifies this view by claiming that deontologists, for the most part, fail to consider the possibility that following their principles could lead to a suboptimal outcome.¹⁹⁴ Despite this uncharitable understanding of non-consequentialist thought, he also recognizes that some deontologists will see that this could happen, and in those cases he can only appeal to the self-evidence of consequentialism to refute them.¹⁹⁵ But this, of course, is circular. The intuitive force of moral consequentialism's most basic and central claims should be apparent to anyone.¹⁹⁶ But its more counterintuitive implications have led some philosophers to reject it all the same. So we can hardly claim that consequentialism (as a comprehensive moral view) is self-evident. If the conceptual entailment from values to norms fails this test, though, the Moorean no longer has any argument. Their claim becomes a brute posit, and one that begs the question against deontologists and—more importantly, for our purposes—the view that values needn't entail moral (and thus special) norms. Hence, the Moorean entailment

¹⁹³See Moore 1942, pp. 558-9, 600-1 for both the advancement of this thesis, and a defence of the idea that two things can be logically equivalent without being identical.

¹⁹⁴Moore 1912, pp. 176-7.

¹⁹⁵Carlson 1995, p. 46.

¹⁹⁶Greene 2013, pp. 107, 166, 170, 194.

view is not a compelling challenge for Weak Non-Normativism.

2.2.2 Values to Be Respected

Consequentialism seems like a natural fit when trying to give an entailment view, because it ties the rightness or wrongness of an action directly to the relative value it produces. But a deontologist view may also present us with such a picture, depending on how it understands value. Many deontologists do, after all, make evaluative claims even while their view focuses on claims of reasons and obligations.

We may take each of these broad ethical paradigms as a theory of how we are to respond to value. Maximizing consequentialism is the view that values are to be promoted; we ought to act in such a way as brings about the greatest amount of value possible.¹⁹⁷ By contrast, one way to understand deontology is as the view that (at least some) values are to be respected, which is done by instantiating them in one's own actions.¹⁹⁸ This latter view is as much an entailment view as Moorean consequentialism, since it too posits a necessary move from values to reasons for action (even if the specific actions it recommends are different).

However, this view needs to be motivated, or else it will be just another case of brute supervenience with all of the problems discussed in §1.7. As I alluded to in §2.2.1, I suspect that the only way to avoid these problems is for the necessity introduced to be conceptual. But if we claim that kind of deontological entailment view is a conceptual truth, we run afoul of many of the same objections that Moore did; in doing so we would blatantly assume the point at issue against consequentialists and against the view that there is no necessary entailment between norms and values.

Normativists might point out that even if the debate between these positions is irresolvable in principle, it is still a debate between *these positions*, and so we end up with entailment, and thus Normativism, either way. This is correct on

¹⁹⁷There are other forms of consequentialism, like the satisficing consequentialism discussed by Slote and Pettit (1984), but they are all concerned with the totality of value produced by an action.

¹⁹⁸Pettit 1989, p. 117.

one level; if the Moorean and deontologist entailment views both threaten Non-Normativism, then the disjunction of them will also threaten Non-Normativism. However, this presupposes that these views exhaust the possibility space regarding entailment relations between values and norms. Put another way, any Normativist who made this argument would be offering us a false dichotomy. We know that these are not the only possibilities, since we have been discussing these views primarily in contrast to a third option whereby no such relationship is posited. Moreover, we have seen that we do not have sufficiently strong reasons to endorse either of these positions, so the idea that we must choose between them exclusively when an alternative exists will not hold much water. Hence neither the Moorean nor the deontologist entailment views is sufficient to ground Normativism.

2.2.3 From Personal Goodness to Moral Norms

Some philosophers posit a direct move from personal goodness both to moral goodness and to moral norms. If moral goodness and moral reasons both reduce to, or are at least entailed by, personal goodness then it looks like wherever there is an evaluative moral fact there will be a normative one as well. (At least, this will hold just as long as facts about personal goodness are the only basis for moral facts of either kind, but since views like this are often explicitly intended to provide necessary and sufficient conditions for moral facts this is not an unreasonable interpretation.)

We can use a view put forth by Peter Railton as an exemplar for positions of this kind. He suggests that both being good and being a reason reduce to what he calls ‘objectified subjective interest’. This is what a fully informed and rational version of some agent would want the non-idealized version of that agent to want in a given set of circumstances.¹⁹⁹ This is an account of ‘goodness for’ (rather than ‘goodness *simpliciter*’)—that is, goodness from the perspective of a particular individual.

¹⁹⁹Railton 2003b, p. 11. He also admits (2003a, pp. 57, 62-63) that this is a theoretical notion to which we have limited epistemic access but also notes that what actually makes claims about this kind of goodness true isn’t the conditional claim about what someone would want under circumstances of ideal rationality but the reductive base of that conditional claim itself, that is to say whichever facts make the conditional true.

Railton takes it as given that a person has, at least, reason to do whatever will promote their own ends. (This need not commit anyone who accepts it to reasons internalism; this claim implies nothing about what other reasons agents may or may not have). And, by combining this instrumentalist theory of reasons with his account of objectified subjective interest, he comes up with a basic standard of rational action: acting in a way that promotes one's objectified subjective interest. Hence the most basic and uncontroversial kinds of reasons are tied inextricably to (what is likely to be) the most basic and uncontroversial—or at least most innocuous from the standpoint of ontological parsimony—kind of good in a fairly straightforward way.

He goes on to suggest that moral goodness is the objectified subjective interest of society as a whole, and rightness or moral reasonhood is whatever is practically rational by that metric.²⁰⁰ By society, he means nothing more than the aggregate of its members, with each one's interests being weighed equally. There are, of course, important differences between the individual and social cases, but there are also some striking similarities.²⁰¹ If we follow the analogy and individuals have a reason to promote or pursue that which is good for them, then societies have reason to promote or pursue what is good for them (or, more accurately, their members). There are puzzles here about whether and in what way groups can actually have reasons, but they do not bear directly on our enquiry. Let us assume that that individual agents are subject to reasons grounded in the social point of view so long as they are members of society, which all agents are because we have no choice but to live alongside others.²⁰² This turns on the assumption that things like stability and the maintenance of moral rules as part of public life is in each of our objectified subjective interests, but this seems plausible enough, especially if one has accepted the Argument from Queerness and thus also accepted a debunking account of moral phenomena. However, no part of my project turns

²⁰⁰Railton 2003b, pp. 21-2.

²⁰¹Railton 2003b, pp. 23-7.

²⁰²Even a hermit cannot escape the effect of other people on their lives, especially now that humans can affect the environment on a grand scale. The whole world may well be a single society in the intended sense.

on the truth of any of Railton's claims, so this is really of no concern for us. We need only worry about whether views like this, if true, will imply a necessary connection between that which is morally good (that is, goodness *simpliciter* or as close to it as exists in Railton's account) and our normative reasons in a way that would endanger Non-Normativism. And, as it stands, Railton's view does seem to entail a move from value to normativity.

The characterization I have just given is a very thin sketch of just one member of a class of views, each member of which is much too complex, nuanced, and comprehensive to cover properly here, and a Non-Normativist could respond by interrogating some of the specifics thereof. However, as much as could potentially be said about in this regard, I don't believe it necessary to say any of it here. This is because, while views like this guarantee an implication from value to normativity, none of them are sufficient to ground Strong Normativism. They may not even be sufficient to ground Weak Normativism.

Remember that Normativism is the view that all moral facts entail not just normative facts, but specially normative facts, whose normative 'reach' extends to all rational agents (and is, on many accounts of its specialness, brute or inherent). But the type of good discussed in this entailment theory just doesn't seem to be up to the task of grounding that kind of normativity. The moral goodness discussed is ultimately grounded in a distinctly personal notion of goodness, and the same is true of the form of normativity with which we are supposed to associate it. Railton admits straightforwardly that, under the view he constructs, the normativity of morality doesn't have the categorical scope many would want to insist is characteristic thereof.²⁰³ I am basically amenable to such a view—this should become abundantly clear in Chapters 3 and 4—but it flies directly in the face of any Normativist convictions.

The same objections apply to similar views. For example, Ralph Wedgwood argues that every kind of normative 'ought' is interdefinable with 'better' insofar as as whatever one ought to do is better than alternatives according to some system or set of preferences.²⁰⁴ But tying reasons to preferences in this way makes them

²⁰³Railton 2003b, p. 30.

²⁰⁴Wedgwood 2009a, pp. 511-2.

agent-relative in a way that does not fit well with special normativity (Wedgwood acknowledges this, as it is the route by which he avoids an implicit commitment to consequentialism). We could postulate a set of preferences that was in some sense distinctly moral, but it would still only have normative force over agents who actually did prefer in this way. These views ultimately generate their moral norms through instrumental reasoning, and while I don't think this move is illegitimate (in fact I suggest something similar in §4.2), I doubt that instrumental reasoning can ever give us a route to special normativity.

More than Normativism of any stripe, these views seem to imply that the act of valuing something on the part of an agent has some normative implications for that agent, but not that the actual value of any external object should be reason-giving in the same way. (That is, this view does not entail that actual matters of evaluative fact would be specially normative in the sense of being reason-giving for any and every agent.) This is an interesting implication, and one that seems to gel with our observations of human behaviour, but it hardly undermines Non-Normativism. Views that reduce both normative and evaluative properties to some third category (or at least extant examples thereof) cannot ground Normativism because they aren't entailment views of the kind Normativists need to invoke.

2.2.4 The Buck-Passing Account of Value

A more promising way that philosophers have tried to account for the connection they posit between norms and values is to try to analyse one in terms of the other, or make one reducible to the other.²⁰⁵ This claim has none of the caveats of those from §2.2.3 and can much more easily ground Strong Normativism. Certain versions of the Fitting-Attitude Analysis of Value fit this mould. The Fitting-Attitude Analysis is the claim that something is valuable if and only if it is the fitting or appropriate object of a 'pro-attitude', which here means one of broad range of attitudes that can be loosely defined as those that we would typically think of as positive, such as desire, admiration, or approval.²⁰⁶ Likewise for a thing to

²⁰⁵When one property is reducible to another, the same facts will underlie the truth value of propositions given in terms of either property.

²⁰⁶Rabinowicz and Rønnow-Rasmussen 2004, pp. 391-6.

be disvaluable is, under this view, for it to be the fitting object of a ‘contra-attitude’, which is a broadly negative or disfavoured attitude.²⁰⁷ The versions of this analysis with which we are concerned are those which take the notion of fittingness to be a normative one, which is to say those that hold that a response to something is fitting when there are normative reasons to have that response to it.²⁰⁸ There are other interpretations of fittingness, which lead to very different understandings of the Fitting-Attitude Analysis, but the normative reading is of most interest to us, since we are concerned with the relationship between reasons and value.²⁰⁹

The most widely endorsed normative interpretation of the Fitting-Attitude Analysis (and perhaps the most widely endorsed version more generally) is what T. M. Scanlon called the buck-passing account of value.²¹⁰ Richard Rowland states the account as follows:

The Buck-Passing Account of Good *Simpliciter* What it is for X to be non-instrumentally good *simpliciter* or of final value is just for X to have properties that provide reasons for us to have non-instrumental pro-attitudes in response to X (such as to desire or admire X for its own sake).²¹¹

²⁰⁷The name ‘contra-attitude’ is drawn from Rabinowicz and Rønnow-Rasmussen 2004, p. 391, but other terms have been used to denote the same category, namely ‘con-attitudes’ by Nowell-Smith (1954, p. 112) and ‘anti-attitudes’ by Ewing (1947, p. 150). I have simply gone with the one I most liked the sound of.

²⁰⁸See Danielsson and Olson 2007, p. 512 for an explanation of the move from fittingness to normativity.

²⁰⁹For an example of a non-normative understanding of fittingness, see Gert 2016.

²¹⁰Scanlon (1998, p. 97) calls it this because it involves passing the normative buck—the job of explaining or grounding normativity—from goodness to other properties, as we shall see in the explanation below. It’s worth noting that the exact relationship between the Fitting-Attitude Analysis and the buck-passing account is subject to debate (Olson, 2013, p. 626) and some might object to my characterizing it as a version of the Fitting-Attitude Analysis, but I think the structural similarities are sufficiently apparent that wherever I equivocate between them I do so innocently.

²¹¹Rowland 2019, p. 9. It’s worth noting that this view goes beyond just being the Fitting-

This needs some explanation. Goodness *simpliciter*, is goodness in and of itself, as opposed to goodness relative to some context, such as goodness qua some particular kind or goodness for an individual.²¹² Of course, our discussion is concerned with value, but Rowland identifies goodness *simpliciter* with final value, and I take it that when people speak of ‘goodness’ in an unqualified way they are speaking about value in the most general sense.

If the buck-passing account is the correct analysis of value, then evaluative moral facts are ultimately facts about what reasons we have, in which case value entails normativity and the arguments I made in §2.1 fail. As such, any view based on Non-Normativism (and hence any view that is immune to the Argument from Queerness) would be untenable.

We can contrast the buck-passing account with what Rowland calls the value-first account.²¹³ This is the reverse of the buck-passing account, holding that the same reducibility relation exists between values and norms in the opposite direction.²¹⁴ This isn’t necessarily a problem for Non-Normativism, since the value-first account is strictly the claim that reasons entail values. That said, if the entailment were biconditional—as it would be if the value-first account were the identity claim that to provide a reason just is to be good, which is a plausible interpretation—it would follow that value entails normativity.²¹⁵ A paradigmatic

Attitude Analysis with a normative interpretation of fittingness insofar as it takes a stand on the ontological question of whether normativity or value is more fundamental, siding with normativity. This question is largely orthogonal to our investigation (though see §2.3.1) but the buck-passing account is still most worth discussing because it has been so widely defended.

²¹²Rowland does offer buck-passing analyses of these other forms of goodness (2019, p. 91, 81, respectively), but since we are concerned with the relationship between value and normativity in general, it suffices for our purposes to limit our discussion to this more inclusive formulation.

²¹³Rowland 2019, p. 9.

²¹⁴It’s worthy of note that one disanalogy between these views is that while the buck-passing account has a canonical form in terms of reasons for pro-attitudes, most extant versions of the value-first account are concerned with reasons for action. There has been little discussion of those variants that are strict inversions of the buck-passing account (Way, 2013, p. 28)

²¹⁵In fact, if both the value-first account and the buck-passing account are seen as an identity

example of the value-first account would be a Moorean view wherein value (in the form of goodness) is irreducibly simple, and what it is for us to have reasons to promote or desire something is just for that thing to be of value.²¹⁶ Our discussion need not concern such views directly, since they don't necessarily endanger Non-Normativism, but it is worth noting their existence, both as a point of contrast for the buck-passing account and because some other entailment views are examples of value-first accounts.

The last alternative in Rowland's taxonomy is the no-priority view. This view is simply the denial of both the buck-passing and value-first accounts, which is to say the claim that no reducibility relation exists between reasons and values.²¹⁷ Rowland takes it that the lack of a priority relationship between value and normativity would mean that neither category is more fundamental.²¹⁸ This suggests either that both values and norms are ontologically basic, or that both are actually reducible. An example of the latter view might reduce each to a distinct set of natural facts. I would like to suggest that the label 'no-priority view' should also apply to views where one of the two is fundamental and the other reducible, but not to the first. For instance, a view that takes value to be fundamental but has an instrumentalist view of reasons, where they are naturalistically reducible to facts about an agent's ends. This exhausts the possibility space for views about the fundamentality and reducibility of values and norms.²¹⁹

claims, the only difference between them is what they take to be ontologically basic.

²¹⁶Moore 1903, pp. 24-6. See Smith 1994, ch. 5, Raz 2001, pp. 164-6, Olson 2006, and Wedgwood 2009b for other articulations. As Olson points out, it's questionable whether Moore himself would have endorsed the value-first account later in his career, but the view certainly seems to be implicit in the *Principia* and it's easy to see why someone influenced by Moore might adopt it.

²¹⁷This isn't to say that the no-priority view is incompatible with all entailment theories—it isn't—but it does mean that it doesn't imply an entailment from value to normativity in the way that the buck-passing account does.

²¹⁸Rowland 2019, p. 9. See Wedgwood 2009a for an example of such a view, where value and normativity stand in a biconditional entailment relationship, but no priority relationship.

²¹⁹This more inclusive understanding of the no-priority view is closer to the account thereof

As I have already stated, the buck-passing account and Non-Normativism are mutually exclusive. Therefore, in order for us to have license to endorse Non-Normativism or any view that entails it, we must argue that it is no less plausible than the buck-passing account. As such, we must scrutinise the case for the buck-passing account to determine whether we have good reasons to buy into it.

2.3 The Case for and against the Buck-Passing Account

Given the incompatibility of the buck-passing account and Non-Normativism, our first instinct as would-be Non-Normativists might well be to offer criticisms of the buck-passing account to show that it cannot be true. However, while it does have its share of potential problems, I'm not confident that any of them is sufficient to motivate the full rejection thereof.²²⁰ In any case, I do not need to discuss any of the arguments against the buck-passing account.

This is because I have reasons for rejecting the buck-passing view, both as a general view and as a grounds for Non-Normativism, which are independent of any such concerns. These can be brought to light by instead examining the arguments given in favour of the buck-passing view. After all, that one view is incompatible with another is only a point against that first view if the latter well motivated. While the buck-passing account has many adherents, I mean to show that it isn't nearly so well motivated as many suppose. Additionally, any and all arguments for it are, in so being, also arguments against Non-Normativism because the two views are incompatible. This makes discussing them more directly relevant to our investigation than the arguments against the buck-passing account discussed

described (but not defended) by Way (2013, p. 29).

²²⁰For a discussion of the problems for the buck-passing account, see Olson 2013, pp. 630-3. I have nothing to add to the arguments given there, nor any objections to them, which is why I have chosen not to reproduce them here. For more in-depth discussion of perhaps the most widely endorsed such argument, the wrong kind of reason problem, see Rabinowicz and Rønnow-Rasmussen 2004 and Rowland 2019, §6.6.

above, which strictly have nothing to say about Non-Normativism except that it cannot be ruled out on grounds of the buck-passing view.

Bearing in mind that our aim here is, and remains, showing that Normativism cannot be assumed *tout court*, in examining each of these arguments I will consider their merit in supporting the buck-passing account as a theory in its own right, but I will also be concerned with their merit in supporting it *as evidence for Normativism*. Likewise, potential counterarguments will be subject to the same treatment; if they have normativistic implications then the argument—or the concerns that motivate it—may not stand in support of the buck-passing account but it will stand in support of Normativism, which is what we must ultimately avoid if we are to build a view that doesn't fall victim to the Argument from Queerness.

One final thing to bear in mind about the arguments below is that, with the possible exception of the redundancy argument (discussed in §2.3.5) they all conclude that we have a *ceteris paribus* reason to accept the buck-passing view, which is to say a defeasible one. As such they are most persuasive when taken as a cumulative case, and some of them might be valid even if we turn out not to have overall reason to endorse the buck-passing account. As it happens, I don't think any of these arguments are successful. Many of them fail to support the buck-passing account over rival views, and some of them are even circular in the context in which we are examining them.

2.3.1 The Ontological Parsimony Argument

A. C. Ewing points out that the buck-passing account gives us the most ontologically parsimonious version of moral non-naturalism.²²¹ This is because it reduces all moral concepts to 'ought'.²²² I have already acknowledged that concerns of parsimony give us reasons to favour some views over others, all other things being equal.

²²¹Ewing 1939, p. 14.

²²²The buck-passing account is usually given in terms of reasons, which Ewing (1959, p. 126) later did, and the ideas of 'ought' and 'reason' are so intimately tied together that the difference shouldn't matter here. Which, if either, of these is more fundamental has been hotly debated, but is likewise irrelevant for our purposes.

If we grant that we should favour more parsimonious accounts over less parsimonious ones, and that the buck-passing account is more parsimonious than either the value-first account or the no-priority view, then we should favour the buck-passing account over any of the alternatives. The success of this argument turns, then, on whether the buck-passing view really is the most parsimonious option available to us.

There is no reason for us to suppose that the value-first account must be less parsimonious than the buck-passing account.²²³ The buck-passing view has it that reasons are fundamental and that values are reducible to reasons, while the value-first account has it that values are fundamental and reasons are reducible to values. Each of these pictures posits one ontological primitive and one reducible entity; they are equally parsimonious.

However, some versions of the value-first account undermine Non-Normativism just as much as the buck-passing account does, so if we are to be secure in endorsing a Non-Normativist viewpoint we shall be better off asking whether the buck-passing view is more parsimonious than the no-priority view. Rowland argues that it is, because where the buck-passing account posits only one ontological primitive, the no-priority view posits two; buck-passers hold that facts about value just are facts about reasons whereas proponents of the no-priority view must insist that the two are quite separate.²²⁴ But this betrays the assumption I mentioned in §2.2. It assumes that just because the no-priority view takes it that there is no relationship of priority between normativity and value, that it must therefore also take both to be ontologically fundamental, which isn't so. Take the example I gave earlier of a view that considers value to be primitive, and reasons to be reducible, but reducible to facts about human desires and how they can be satisfied.

²²³Olson 2006, p. 532 and 2013, p. 630. Olson's discussion is concerned with the Moorean view, but the logic of the argument that follows could apply to most, if not all, versions of the value-first account. That isn't to say that there couldn't possibly be some form thereof that is less parsimonious than some form of the buck-passing account, but I can think of no examples. In any case, the argument will stand so long as there is even one version of the value-first account which is not less parsimonious than any version of the buck-passing account.

²²⁴Rowland 2019, pp. 70-1.

Like both the value-first account and the buck-passing account, it makes one out of normativity and value ontologically basic and the other reducible, so by this metric it's just as parsimonious.

Proponents of the buck-passing account might object at this juncture that the model I just suggested is still less parsimonious because it introduces another set of fundamental entities beyond values for reasons to reduce to. But what of it? Unless they want to insist that their ontology doesn't make room for facts about, for example, human psychology and causation, this move is illegitimate. We aren't concerned with the number of entities the theory posits so much as the number it requires us to add to our ontologies, and most of us already grant the existence of natural facts like those mentioned above.

It's also notable that the buck-passing view is touted only as the most parsimonious form of non-naturalism, since it posits a *sui generis* non-natural property of being a reason. The no-priority view isn't bound in this way; a view whereby both value and normativity were reducible to distinct sets of natural facts would be more parsimonious still, as it would posit no such entities. The Argument from Queerness as I gave it in Chapter 1 assumes non-naturalism, but all the variants of Normativism and Non-Normativism are strictly neutral on that question, as is the Purely Evaluative Moral Realism that I describe in Chapter 3, although I favour a naturalistic interpretation thereof.

So it seems like ontological parsimony doesn't give the buck-passing account an edge over the no-priority view. If we assume moral non-naturalism then the most parsimonious versions of each are tied. And, if we don't make that assumption then the no-priority view may well be more parsimonious.

2.3.2 The Heterogeneity Argument

The heterogeneity argument is based on the premiss that moral status is 'irreducibly heterogeneous', which is to say that any two good things needn't have anything in common except their goodness. This might seem like an oddly controversial starting point, since certain substantive ethical theories (such as hedonism) do necessitate that good things all have certain natural features in common. We can at least say that moral status is not necessarily homogeneous, though; even a

committed hedonist can make sense of the claim that something besides pleasure—knowledge, for example—is intrinsically or non-instrumentally good, even while disagreeing with it. So, while that initial contention seems very strong, it looks much more plausible if we consider it as the rejection of certain restrictions on which moral claims can be made together. Either way, however, I think this argument has deeper problems that render it untenable even if one accepts this premiss.

Were we to encounter two good but otherwise very different objects (or events, or people, or states of affairs), we might be tempted to ask for an explanation of how they share their goodness when moral qualities supervene on natural* ones and their natural* qualities are so radically divergent.²²⁵ The buck-passing account answers this by saying that what they have in common is that their natural* features give us reasons to hold pro-attitudes toward them. The pro-attitudes in question needn't be the same—in fact, we shouldn't expect them to be—because this is, in effect, an account of how things can be valuable in different ways.²²⁶

This is meant to be one of the major virtues of the buck-passing account (and by extension the Fitting-Attitude Analysis); it can elegantly account for a plurality of values.²²⁷ Of course, to argue that this is a reason to accept it one must first accept that this kind of pluralism about values (and everything it entails, including potentially the incomparability or incommensurability of values) is conceptually possible. Again, these positions are contentious, and though they might seem *prima facie* plausible, that doesn't necessarily mean that they will hold up to closer

²²⁵As in Chapter 1, by 'natural*' I mean to encompass both natural properties, and supernatural ones such as being willed by God.

²²⁶A potential problem for buck-passing accounts is the question of what unifies pro-attitudes—what makes them all 'pro'. See Rabinowicz and Rønnow-Rasmussen 2004, pp. 141-42 for some discussion of this. There is also the question of whether an account of this unifying property that does not somewhere invoke value, making the Fitting-Attitude Analysis circular, can be given. I owe this worry to Stephen Everson. Since I am not defending the buck-passing account, and I am more concerned with undermining positive arguments for it than directly arguing against it for reasons already given, I will not attempt to give an answer here.

²²⁷Rabinowicz and Rønnow-Rasmussen 2004, pp. 139-40.

scrutiny. But let's suppose that we do accept values pluralism. The heterogeneity argument still fails to support the buck-passing account over the no-priority view.

The argument basically says that we have (at least some) reason to favour the buck-passing account over rival theories of the relationship between normativity and value because it accounts for value pluralism—or the possibility thereof—in a way that those other theories cannot. There are two ways we can understand this. One is that the actual plurality of values is explained by each value having some corresponding pro-attitude, and are unified insofar as pro-attitudes are unified by their shared 'pro-ness'.²²⁸ The other is that there is a plurality of values to which a single particular pro-attitude is warranted, which explains the unity of morality.²²⁹ The second interpretation has a pretty major flaw if we take the buck-passing account as a reductive analysis of value, which seems to be how it is intended. If being such that a pro-attitude is a fitting response is all that there is to being a value, then the pro-attitude in question seems to be the only thing that could differ to make one value distinct from another, so this reading is incompatible with any real pluralism about values. If, however, one is invoking the buck-passing view in defence of Normativism, as we are, it is enough for that purpose if the account only sets necessary conditions for being a value. This leaves room for expansions of the account wherein other points of variation are possible. As such, I will respond to this interpretation as well, for the sake of completeness. This second reading of the argument doesn't explain the plurality of values so much as posit one and use it to explain both how morality can be a unified domain of discourse, and why we have the intuition that good things needn't have anything in common but their goodness, but this still gives the view some explanatory power, which serves as a point in its favour. We might also combine both of these interpretations so that distinct types of value—moral, aesthetic, epistemic, etc.—each have a number of 'sub-values' under them and are each unified by a characteristic fitting attitude. However, I think that the heterogeneity argument fails for both interpretations (for reasons I explain below) and that this third option would simply inherit both

²²⁸See Rabinowicz and Rønnow-Rasmussen 2004, p. 401.

²²⁹This interpretation was suggested to me independently by Paul Noordhof and Christian Piller.

of the parent views' flaws.

The first view does successfully account for pluralism in a way that is not possible under a Moorean model because it takes value to be goodness, which is a single unified property. However, the point of comparison we are interested in is the kind of model presupposed by Non-Normativism. While I tend to speak of goodness or value as something unified for simplicity, there's no reason that a Non-Normativist has to be committed to this kind of model; they might just as well posit multiple values which may well be incomparable or incommensurable or both. They can just say that there are many ways for things to be valuable—many sets of evaluative facts.²³⁰ They would even face fewer problems in doing so than Normativists, including buck-passers, because they don't need to offer a solution to the normative puzzles that arise from incomparable values. (If such values also raise evaluative puzzles—itsself doubtful since it seems nonsensical to ask which of two genuinely incomparable things is better—they are, at least, less pressing since since we need not suppose any further practical or deontic question turns directly on the answer.)

Proponents of the buck-passing account might counter that while both views are compatible with a plurality of values, only their view can explain said plurality because under the buck-passing account each value corresponds to a different pro-attitude, whereas for those who endorse the no-priority view the plurality of values must be brute.²³¹ Such a model would be both ad hoc and ontologically profligate. But we can turn this argument back on itself and ask how the buck-passer is to explain the plurality of pro-attitudes. That there is such a plurality is uncontroversial; attitudes are a part of nature and we are all perfectly familiar with the likes of admiration and approval, but this does not mean that there is no cause to explain them.

One possible answer is that they are a response to the values that actually—and

²³⁰An account is then needed of how values are non-trivially unified as values, and distinct from one another as moral values and aesthetic values and so on. Non-Normativism itself does not commit us to any one answer, and I do not have one to offer here, but one is owed as part of a more complete metaphysics of value.

²³¹I owe this observation to Mary Leng.

somehow independently—exist in the world, but this response has two major flaws. First, it puts the buck-passing account back on the same level as the no-priority view in that it requires us to posit the same number of ontological primitives. Second, since the buck-passer analyses values as reasons for pro-attitudes, this response is actually the claim that we developed the capacity for these attitudes in response to the reasons to hold these attitudes. This seems to be circular, as our reasons to have some attitude surely cannot logically (or chronologically) precede our ability to have it. There is a plausible case that one cannot have a reason to do anything that one cannot do, and it seems quite reasonable to say that a fitting or required response is never also an impossible one since ‘ought’ implies ‘can’.^{232,233}

The alternative is to assert the primacy of pro-attitudes and say that the values (reasons for pro-attitudes) that we can acknowledge and respond to are determined by the kinds of pro-attitudes we are able to hold. Values would then exist in the same way as colours and other secondary qualities, in that to be valuable is to have properties such that they give rise to a certain response from creatures like us, but have these properties whether or not any responders of the right kind exist. This neatly explains the posited correspondence between values and pro-attitudes, but it also runs the risk of causing this argument to fail to support Normativism even as it supports the buck-passing account. Claims like this may undermine the idea that the normativity of value is special normativity, because it makes that normativity contingent on the psychology of actual agents. A non-human agent (or just an atypical human one) might lack certain pro-attitudes and so have no reason to respond to certain values, when one of the hallmarks of special normativity—the kind Normativism assigns to moral value—is that it applies to

²³²For one case against reasons to do the impossible see Streumer 2017, §65-8. Consider also that we could not reasonably criticize someone for failing to do an impossible thing, whatever benefits it might have had.

²³³We might, at best, be said to have a reason to develop the capacity for some attitude that engendered certain responses for which we had independent reasons (insofar as one can have reason for such a fundamental kind of change; the actual explanation is probably an evolutionary one not underlaid by any form of normativity) but it is more expedient just to say we have reasons for those responses themselves.

all agents regardless of these kinds of facts about them. Hence, a version of the heterogeneity argument based on this understanding of the normativity of value is no good to the normativist.

The second interpretation likewise fails to explain the plurality of values better than a Non-Normativist model; both simply take it as given (though in the latter case doing so at all is inessential to the view). But, if it is correct, it may make the buck-passing account superior insofar as it can explain why a certain array of values comes together to form the moral domain. To do so, it needs to nominate one attitude that is fitting to hold towards all forms of moral value. If such an attitude could be identified, this interpretation would lead to one point in the buck-passing view's favour, but this is another heavy burden for those making this argument to bear on top of a commitment to pluralism about values and the claim that this interpretation should be favoured over the first. As such, while this argument doesn't appear to have any fatal flaws, it is still not very persuasive.

These problems need not be irresolvable, but they at least show that the buck-passing account may be less versatile than many have thought, and that many philosophers will have further work to do to make it cohere with their other commitments.

The argument from heterogeneity claims that the buck-passing account is better than the value-first account or the no-priority view insofar as it can better account for value pluralism. While the value of doing so is an open question, we have seen that it is no better at accounting for a plurality of values than the no-priority view, and so this argument fails.

2.3.3 The Open-Feel Argument

One argument for the buck-passing account draws on another influential argument within metaethics, Moore's open-question argument, in a somewhat unconventional way. This is Scanlon's Open-Feel Argument. Before I explain it, I should say a little about the open-question argument to make clear how the one draws upon the other.

The open question-argument is an argument for moral non-naturalism. The open question is '... but is it good?', where the blank is filled with a statement that

something is, for example, pleasurable or in accordance with human nature. When Moore calls this question ‘open’ he’s saying that it’s legitimate or reasonable to ask. He asserts that we can always legitimately or reasonably ask such questions, and that the reason that this is so must be because no natural property is identical with goodness (because the question of whether a thing is identical with itself is neither a legitimate nor a reasonable one; everything is necessarily identical with itself).²³⁴

This argument depends crucially on an intuition: we conclude that the open question is open because every example we can think of *feels* open. This feeling, the open feel of the open question, is what the Open-Feel argument draws on. The argument says that one of the virtues of the buck-passing account is that it’s able to explain this open feel. This is because it analyses value, and so goodness, in terms of reasons for pro-attitudes, and we can always ask whether something’s natural properties give us reason to favour, admire, desire, et cetera that something. This, in turn, is because identifying natural properties such that facts about them entail some further facts about what one ought to do isn’t the same thing as actually reasoning through to those practical implications.²³⁵

Olson objects by pointing out that this logic can easily be flipped on its head. If we ask whether we have reasons to hold a pro-attitude toward something, we may simply be asking whether it is good.²³⁶ But this kind of reversal does nothing for the Non-Normativist, because it still allows for the kind of necessary connection they wish to deny. A more promising tactic is to question whether the open feel needs any further explaining at all. The reason I say ‘any further explaining’ is that it seems to me that Moore’s open-question argument relies on the same kind of logic as the open feel argument. By Moore’s lights the non-identity of goodness with any natural property is what explains the open feel of the question ‘...but is it good?’. Just as identifying properties whose instantiation entails a practical conclusion isn’t the same as drawing that practical conclusion, identifying one property that is coextensive with another (as goodness might be with some

²³⁴Moore 1903, pp. 16-7.

²³⁵Scanlon 1998, p. 96.

²³⁶Olson 2013, p. 628.

particular natural property) isn't the same as identifying that second property.

If the open-question argument works, this is how it works, and so the open feel demands no further explanation. But who says it works? In §1.8, we discussed a counterargument based on non-obvious identities. It may seem that said counterargument should lead us to favour the buck-passing account. The open feel of the open-question need not be explained by the non-identity of goodness and any natural property. The buck-passing view offers us another explanation. But the demand for an explanation of the open feel still seems rather odd. If the truth of identity claims doesn't need to be obvious, it's not clear why we ought to be surprised that the open question feels open. The open feel doesn't seem to demand explanation in such a way that it's any great virtue of the buck-passing account that it offers some explanation thereof.

As such it seems that if we accept the Kripkean idea that identity relationships are necessary but not necessarily knowable a priori, then the open question fails to be mysterious. Some thinkers may be inclined to reject this contention about identity, but they will find themselves back where we began, where the non-identity of goodness and any natural property explains the openness of the open question. Either way, we don't seem to have any especially strong reasons to accept the buck-passing account.

2.3.4 The Demystification Argument

Another argument for the buck-passing account which is grounded in the view's supposed explanatory power is the demystification argument. What it aims to explain is the apparent normativity of values. That is to say, there is a compelling intuition that an ascription of value brings with it normative reasons to think, feel, or act a certain way, and this view explains this by giving an explicit characterization of that connection.²³⁷ I grant that an explicable connection is much preferable to a brute one for the reasons discussed in §1.4 and §1.7. So, insofar as the buck-passing account is the best (or only) explanation for some connection, we have some reason to accept it.

²³⁷Rabinowicz and Rønnow-Rasmussen 2004, p. 391.

However, we can immediately see that this argument goes awry. Firstly, this virtue doesn't belong exclusively to the buck-passing account; the value-first account can explain the same connection in the same way—through reduction of one category to the other—so the buck-passing account may be neither the only, nor the best explanation.²³⁸ As I've already stated, the value-first view needn't be incompatible with Non-Normativism (though it can be), which would rob this argument of some of its power to support Normativism. That said, even if the explanations these two views offer of this particular connection are equally elegant, the buck-passing account may be a better explanation if we have other reasons to accept it (such as the other arguments I discuss in this section) or independent reasons to reject the value-first account.²³⁹ I don't think that any of the arguments for the buck-passing account work, but it's worth acknowledging the possibility that it might be superior to the value-first account in ways that fall outside of the scope of any of the arguments I discuss here. Suppose that it is. This argument still fails in the context in which we are examining it.

Recall that we are examining the demystification argument in support of the buck-passing account both in its own right and as evidence for Normativism. In that light this argument is flawed from its first step. The demystification argument is circular when given in support of Normativism. The buck-passing view may well give us the best explanation of a connection between normativity and value, but the idea that this is a theoretical virtue relies on the assumption that there is such a connection to be explained in the first place. But we only began discussing this account because it offered theoretical support to the idea that there is a necessary connection between norms and values, which is to say Normativism. We cannot, therefore, invoke that idea to argue for the buck-passing account. It plainly assumes the point at issue.

It may be objected that the very fact that we feel a need to explain such connections is evidence that such connections holds. But this does not stand up to scrutiny either. As noted in our characterization of the the demystification

²³⁸Olson 2013, p. 629.

²³⁹For examples of independent reasons to reject the value first-account, see Rowland 2019, chs. 2, 3.

argument, there are apparent connections between value and normativity. That is to say, it *seems* that such connections hold. But apparent connection may be appearance only; we can explain the appearance of something without positing its actual existence.

As it happens, we already have some good candidates for an explanation of why there might appear to be some kind of necessary connection between norms and values. And these candidates do not require that any such connection actually holds. Think back to the discussion of debunking accounts of morality in §1.3. We noted that categorical imperatives are key to our moral thoughts (without making any commitments as to the structure of moral facts). And it is certainly plausible that people might draw connections between those imperatives and value as part of normal moral thought, whether in an attempt to systematize their ideas or just because they are getting their wires crossed. Or we might appeal to a model of human psychology where perceptions of value (the act of valuing, rather than actual instantiations of value as a property) go hand in hand with normativity of the unproblematic non-special kind, as discussed in §2.2.3. None of these are fully developed accounts, but what matters is that it's almost certainly possible to construct and defend accounts like these that don't invoke Normativism to defend the common presumption of Normativism. These would be, therefore, debunking accounts of Normativism.

So we have seen that the demystification argument's theoretical virtues are dubious, and that it fails to establish the buck-passing account as evidence for Normativism because it rests on an assumption of Normativism. As such this argument gives us no reason to accept the buck-passing view.

2.3.5 The Redundancy Argument

One of the most compelling arguments for the buck-passing account is what has come to be known as the redundancy argument.²⁴⁰ It begins with the observation that normative reasons for some course of action always boil down to the natural* features of that action (or its consequences). That is to say, we should perform some action because it will, for example, save a life or instantiate a (naturalistically

²⁴⁰This name comes from Crisp 2005.

described) virtue or maximize utility. This would make goodness explanatorily redundant as a further property on top of the natural* ones in virtue of which our reasons arise.²⁴¹ Scanlon suggests that we should therefore understand goodness as the higher-order property of having lower-order natural* properties that give us reasons of a certain kind, which is precisely the kind of analysis the buck-passing account offers us. In other words, goodness offers us no *non-derivative* reasons for action.²⁴²

Rather, we have reason to respond to good things in certain ways because we have reasons to respond to those natural* properties in virtue of which they are good. The observation that goodness ‘does’ nothing over and above the properties that ground it is one we can’t reasonably contradict in light of the supervenience of the moral upon the natural*, to which I’ve already assented. It also gels with ordinary language; when someone tells us something is good we might ask why or how it is so, and we would expect an answer in terms of its natural* features. (Unless the thing in question is meant to possess final value, where any answer would need to be metaethical in nature.) As such, the redundancy argument initially seems to have legs.

But cracks begin to show in this account once we scrutinize it more closely. For example, we have already discussed how familiar natural facts always have the potential to ground norms of the familiar kind that seem to be rooted in (or at least intimately connected to) human desires. Given that, the fact that something is good (assuming it is good *simpliciter*) might still have work to do making the jump from non-special to special normativity. I don’t mean to endorse such a model (in fact it would be disastrous for my project if I did, as it rests on a

²⁴¹Scanlon 1998, p. 96. It’s worth noting, though, that Scanlon (2014, pp. 37-8) later says that thick evaluative concepts can be reason-giving in the same way as natural* ones and that the evaluative and descriptive parts of such concepts cannot be analysed apart (these claims need not bother Non-Normativists unless it is assumed that the reasons connected to thick concepts are specially normative). This later view contradicts the model I have just described, so it seems Scanlon has moved away from this view, but the redundancy argument is still seriously advanced by other philosophers like Rowland (2019, §4.2).

²⁴²Rowland 2019, p. 61.

distinctly Normativistic understanding of goodness), but it does show that even by Normativist lights there is something missing from the picture proponents of this argument suggest.

Thin evaluative notions like goodness and badness may also have some utility in expressing our overall reasons. Where something is good in virtue of one or more of its natural* features, and yet at the same time bad in virtue of one or more of its other features, our final response is grounded in an all-things-considered judgement that is most easily given in these terms. That said, it is questionable whether we tend to think or speak of all-things-considered reasons as something beyond the aggregation of all of the other reasons—grounded in natural* facts—at play. If it is so much as possible for such judgements to be accounted for in this way, it will weaken this kind of response.

But there is a deeper problem with this argument, which it shares with the demystification argument; it's circular. Again, this is a criticism we can only levy against it in the context of grounding the buck-passing account for the further purpose of grounding Normativism because the redundancy argument takes Normativism as an implicit premiss. It assumes that the natural* properties that ground a thing's goodness are the same natural* properties that ground our reasons to respond to that thing in specific ways. But there's no *prima facie* reason that we can't say that while both are grounded in natural* properties, they might well be grounded in two wholly distinct sets of such properties, and that any overlap between those sets is contingent. We don't need to assume any kind of entailment relationship—much less an identity relationship—between moral facts and reason-giving concerns. The redundancy argument assumes the point at issue against Non-Normativism and so fails to ground Normativism.

2.3.6 The Similarity Argument

The final argument for the buck-passing account is one I shall call the similarity argument. Like many of the other arguments discussed in this section, it tries to support the buck-passing view by showing that it can explain something we observe about our moral thought and language. The explanandum in this case is a number of apparent similarities between the realms of the normative and the

evaluative.

The go-to example for those who advance this argument is in terms of the debate around their structure.²⁴³ In both cases we have the idea of a *pro tanto* or contributory instantiation. For example, something might have features which make it *pro tanto* good, which is to say good in some way or to a certain extent, without that thing necessarily being good overall. Likewise we might have *pro tanto* reason—a reason of a certain weight—not to perform an action because it causes some harm, but overall reason to perform it because it prevents some greater harm. And in both cases these instantiations have given rise to a debate between atomists and holists. Atomists believe that when something is *pro tanto* good (or a reason) then it is so in all circumstances under which it obtains. So if pleasure is *pro tanto* good, even sadistic pleasure is good insofar as it's pleasure, though it might be bad overall. By contrast, holists would say the fact that sadistic pleasure is bad 'erases' or 'wipes out' that *pro tanto* value, and an analogous view exists for *pro tanto* reasons.

We may find one or the other of these positions more plausible, but what matters for the buck-passer and the Normativist is that we find the same dichotomy across both domains. From there they can offer a simple argument from the best explanation; if this similarity and others like it are best explained by the buck-passing account then we do have (appropriately enough) *pro tanto* reason to accept the buck-passing account. The obvious question is whether that account is in fact the best explanation, which it must be exclusively if this argument is to give us reason to accept it exclusively.

Rowland discusses one alternate explanation; a version of the no-priority view that posits a necessary connection between normativity and value that isn't explained by a dependence or reduction relation.²⁴⁴ Such a view posits a brute Normativism of the kind I argued against at the beginning of §2.2. If it were the best explanation of this observed similarity, however, we would have some reason

²⁴³Regarding the relevance of the debates I describe hereafter to the buck-passing account, see Way 2013, p. 31 and Rowland 2019, §4.5. For the debates themselves, see Hurka 1998, Dancy 2004, and Fletcher 2009.

²⁴⁴Rowland 2019, p. 74.

to accept it. And, though such a connection wouldn't give us an argument for the buck-passing account, it would give us one for Normativism, which is our primary object of inquiry. As Rowland points out, though, an entailment relation between two things doesn't entail structural similarity between the relata; 'ought' implies 'can' yet debates about what we can do and what we ought to do take very different forms.

The other obvious alternative is the value-first account. The similarity between normativity and value extends to all norms and values—at least, no counterexamples come to mind—so the only version of the value-first account that would be at least as good an explanation as the buck-passing account would be a biconditional one, since norms and values would always coincide under such a view. This, too, would entail Strong Normativism so if the value-first account and the buck-passing account explain the similarity equally well (or there were some argument that this biconditional version of the value-first account explained it even better) then this argument would not support the buck-passing view but it would support Normativism.

If I am to undermine the similarity argument's power to motivate Normativism, I must offer an alternative explanation that isn't grounded in an entailment view and is, at least arguably, superior to the buck-passing account. But I believe that I can. It's possible that proponents of this argument simply have the order of causation backwards. That is to say, rather than the similarity being evidence that we should accept a necessary connection between value and normativity, it might be that it is evidence that many thinkers already have. If one has Normativism or something like it (for example, the belief that all moral properties are of a single unified type) as an implicit and unexamined background assumption, then one is likely to treat the relationship as a simple equivalence. And if the properties are equivalent, one would naturally expect to see this kind of similarity in their structure.

Note that I do not say it would be well-motivated; as Rowland notes, the former doesn't strictly entail the latter. Rather, I take it that these assumptions, if they are such, lurk in the background and have thus far gone without the kind of close scrutiny that would reveal these flaws.

I have already made a case for why someone might implicitly accept Norma-

tivism when I discussed debunking accounts of moral phenomena. If we accept a debunking account—which I think we have reason to do independently of any metaethical debates—then we are actually one step ahead because it, along with the argument above, explains the similarity between debates about values and reasons without making any commitments as to the structure of either. Insofar as potentially controversial commitments (which is to say, most commitments) are a cost for a theory and neutrality a virtue, this gives us a better explanation than the buck-passing account offers.²⁴⁵

That being the case, this argument fails to ground the buck-passing account. The considerations that motivate the similarity argument might still have grounded Normativism if the actual best explanation for the observed similarity were one that entailed Normativism, but as we have seen that isn't the case either. As such, the similarity argument fails on both counts. It's in good company, though, as none of the arguments we have considered give us reason to endorse the buck-passing account. So the buck-passing view does not support Normativism.

In this chapter I have discussed only one assumption of one of the two apparently successful queerness arguments. As I shall argue in the next chapter, however, it is a pivotal one. The assumption, Strong Normativism, is that moral facts must always entail special norms.

We already know that any form of moral realism that is to withstand the Argument from Queerness must deny this assumption. To do so is to endorse some form of Non-Normativism, and I have argued that the most plausible such views rest on the denial that values necessarily entail norms. We have seen that attempts to undermine Non-Normativism of this kind by invoking various entailment views fail, because these views themselves are not sufficiently motivated.

To be clear, I haven't shown that Normativism is implausible or that Non-Normativism is superior to it. Rather, it seems to me that both positions gel with some of our moral intuitions, and are *prima facie* plausible. What I have argued in

²⁴⁵We might also appeal to Finlay's (2014) reductionist view, which ties values and norms together but does not assign special normativity to moral values, to explain this observation. I do not wish to rely on this view, however, as it is still more committal than the one described above.

this chapter is that attempts to rule out Non-Normativism fail. It remains a live option, so any more comprehensive views that assume or entail Non-Normativism cannot be dismissed purely on the grounds that they do so. This also means that the Argument from Queerness is not on such secure ground as it might initially have seemed to be, as the denial of one of the vital premises of the Argument from the Queerness of Special Normativity is no less well grounded than the affirmation thereof. The fact that Non-Normativism remains a live possibility opens up some possibilities which will be explored in the next chapter.

Chapter 3

Purely Evaluative Moral Realism

In the last chapter, we discovered that the version of the Argument from Queerness which poses the largest threat to moral realism assumes Strong Normativism. But we also discovered that this assumption can be called into question. I therefore suggested that for any metaethical view to be immune to the Argument from Queerness, it must at the minimum deny Strong Normativism and affirm some form of Non-Normativism. This is the view that it isn't conceptually necessary that all moral facts entail specially normative facts, or, put another way, that there could be a class of facts which are moral (in the sense of being about the content of morality) but are not specially normative reasons for action. I suggested that the best candidate for such a class is the class of evaluative moral facts. Such facts are concerned with goodness and badness without necessarily saying anything about rightness, duty, or reasons for action. Examples include those facts expressed by claims like 'stealing is bad' (as opposed to 'stealing is wrong' or 'one ought not to steal') or 'pleasure is good'.²⁴⁶

The main question this thesis aims to answer is whether Purely Evaluative Moral Realism can avoid the Argument from Queerness, and with the groundwork

²⁴⁶These facts could be facts about action types, or states of affairs, or personality traits, or mental states, or anything else. Which things can actually be the bearers of moral value is a question that goes far beyond the scope of our inquiries here, but nothing in the view I will go on to describe in this chapter is incompatible with any particular answer.

of the last two chapters in place we are finally in a position to begin tackling that question in earnest. In this chapter, I will lay out what I mean by ‘Purely Evaluative Moral Realism’. This will involve specifying what commitments it involves, and what implications it has.

I will then go on to show why we might conclude that such a view is immune to the Argument from the Queerness of Special Normativity. I will also discuss how this view might elude the Argument from the Queerness of Supervenience since, while it is not as strong as the Argument from the Queerness of Special Normativity, the claims that would make it a legitimate threat to moral realism might still be argued for. Once this is done, I will anticipate and respond to objections to Purely Evaluative Moral Realism. Since the hurdle with which we are concerned is the test of basic plausibility, I will not devote the same kind of space to extolling the virtues of the view. This is not to say that there are no arguments in favour of Purely Evaluative Moral Realism. There is at least the obvious one; it avoids the Argument from Queerness while preserving some highly intuitive moral claims, such as that certain actions or states of affairs are bad. There are also further arguments that I explore in Appendix A. Here, though, I will limit myself to defending the view, as it only needs to be a live possibility to undermine the conclusion of the Argument from Queerness.

3.1 What Purely Evaluative Moral Realism Is and What It Isn’t

Purely Evaluative Moral Realism involves both positive and negative commitments about moral facts and what they are like. Some of these commitments it shares with existing views, while also differing from those views in key ways. I will outline these commitments, commonalities, and differences in turn.

3.1.1 Commitments of Purely Evaluative Moral Realism

As argued in Chapter 2, traditional ethical theory recognizes two kinds of moral fact: specially normative facts and evaluative facts. And, while our assessments of

these categories are often connected, they do not stand in any kind of relationship of necessary entailment. So we need to consider the following claim:

Weak Non-Normativism Some kinds of moral facts do not entail specially normative facts.²⁴⁷

This is a conceptual claim, and one that is compatible with any number of other metaethical commitments. Weak Non-Normativism has no implications for questions of moral realism versus antirealism, nor for naturalism versus non-naturalism. But it contradicts Strong Normativism.

I also argued that any form of realism that was immune to the Argument from Queerness from Special Normativity would need to be committed to Weak Non-Normativism. But it would need to make further claims as well—if only to qualify as a form of realism at all. We could do the bare minimum to characterize such a realism, and simply add the claim, firstly, that there are some moral facts and, secondly, that they don't necessarily imply any specially normative facts. This rather thin sketch still allows for the possibility of specially normative facts, though, and thus still has the potential to fall foul of the Argument from Queerness. If we are to be sure our realism can avoid the Argument from Queerness, it must be free of this possibility. It must be purely evaluative.

Purely Evaluative Moral Realism consists in the following three (or two) claims:

There are moral facts.

Moral facts are *all* evaluative in nature.

No moral facts entail specially normative facts.

If we understand the distinction between moral values and special norms to be a metaphysically deep one, then the latter two of these claims are logically equivalent.

Since these claims are what the purely evaluative view amounts to, a commitment to Purely Evaluative Moral Realism brings with it a commitment to Weak Non-Normativism, at the very least. Purely Evaluative Moral Realism is also compatible with:

²⁴⁷As stated in §2.1, a normative fact is a fact about normative reasons, and a specially normative fact is a fact about specially normative reasons.

Strong Non-Normativism No kinds of moral facts entail specially normative facts.

This would entail Weak Non-Normativism. In fact, if one assumed Strong Non-Normativism (and could not suggest a third category of moral facts, see §2.1), then Purely Evaluative Moral Realism would be the only kind of moral realism that could possibly be true. However, as we have already discussed, the various forms of Normativism and Non-Normativism are meant to be conceptual claims, and it simply isn't plausible that the very concept of a moral fact excludes the possibility of special norms.^{248,249}

Purely Evaluative Moral Realism goes beyond making conceptual claims to making ontological claims. It is committed to the modest thesis that it's possible for there to be moral facts that don't entail special norms, and then goes on to claim that *as a matter of fact* all moral facts are of such a type. A Purely Evaluative Moral Realist could even explain why they exclude specially normative moral facts from their ontology; they need only point at the Argument from the Queerness of Special Normativity. These two ontological claims—that there are moral facts, and that they are all evaluative rather than (specially) normative—are the core of Purely Evaluative Moral Realism, and since their truth entails the possibility of moral facts that aren't specially normative they entail Weak Non-Normativism, but since they don't entail the logical impossibility (but rather the mere absence) of specially normative moral facts, Purely Evaluative Moral Realism doesn't stand in the same relation to Strong Non-Normativism.

This makes Purely Evaluative Moral Realism, at its core, a rather thin view. It is neutral with regard to the actual nature of moral facts, except insofar as it maintains that they are evaluative, and that being an evaluative moral fact doesn't

²⁴⁸One might claim that some non-cognitivists are committed to Strong Non-Normativism, since their view too does away with special normativity and is presumably meant to be taken as a conceptual truth, but this is not so. Non-cognitivists take it that moral language isn't in the business of making factive claims in the first place, and need have no commitments whatsoever about the nature of moral facts.

²⁴⁹As stated in Ch. 2, the kinds mentioned are conceptual entities and need not have any actual members.

mean being (specially) normative or entailing anything (specially) normative. It is even, for the most part, neutral on what it is for a fact or property to be evaluative. A Purely Evaluative Moral Realist must reject any view of value generally or moral value in particular that posits irreducibly or otherwise specially normative reasons, but there are a number of accounts of what value is that are amenable to the purely evaluative view.

For example, one simple answer available to Purely Evaluative Moral Realists is that value is a non-natural primitive that won't brook any analysis. To be sure, this claim needs to be about value in general and not just moral value, lest it become a suitable subject for a queerness argument (see §1.4 and §3.3.2), but it is nonetheless open to an adherent of the evaluative view.

They might equally well go the opposite route and attempt some kind of naturalistic reduction. One such view reduces and relativizes all instances of goodness as goodness for the realisation of a specific end. Stephen Finlay holds that attributions of goodness are often logically incomplete. What is missing—but usually supplied by context—is the end-state to which the goodness in question is relativized. To say that something is good is then to say that the probability of that end-state coming to obtain is greater if that thing exists (or is performed, if it is an action).²⁵⁰ Elsewhere in this book, I have relied on the view that normative reasons—ontologically unproblematic ones, at least—are also reducible to some kind of means-ends relationship of this stripe. One might object, then, that insofar as they share a reduction base, this means that values and norms stand in a relationship of necessary biconditional entailment, which directly contradicts the central claims of Purely Evaluative Moral Realism. The reasons in question, however, can all be understood as non-specially normative since they all reduce to these means-end relationships. This would even go for moral reasons, which would share their reduction base with moral values. There is no room in this model for a form of value with the kind of unconditional reach that Normativists assign to

²⁵⁰Finlay 2014, §§2.1, 2.6. This is a paraphrase; Finlay's final analysis is the result of extensive argument and includes modal elements that I have omitted, since all that matters for our discussion is whether the kind of reduction involved in the analysis is compatible with Purely Evaluative Moral Realism.

morality and so no way for it to fall victim to the Argument from the Queerness of Special Normativity. Whether morality can be properly understood in terms of some characteristic end is a separate question, and one that is closer to substantive ethics than to metaethics, so I will not engage with it here.²⁵¹

Yet another option draws on the examples from §2.1 of *Paradiso* and the world with happy egrets but no past, present, or future agents. These worlds or the states of affairs that they comprise are good where no reasons can exist. This serves as evidence against views like the buck-passing account, which analyse value in relation to reasons and our status as agents, but it also suggests that we might have more success going in the other direction and analysing it in terms of our status as patients or subjects. I do not have a detailed picture to offer, but we can make a start. We would need to start with characteristic features of patients, such as affective responses or just the experiential character of that to which they are subject. We could not appeal to normative reasons for such things without contradicting our existing commitments, but we could appeal to the responses or experiences that actually do or would obtain. This would give goodness for particular individuals ontological priority over other kinds of goodness, and would need to attach particular responses to particular values. A very general version of this would probably lead to something like a utilitarian axiology. None of these consequences need to be objectionable to the Purely Evaluative Moral Realist, however, so this looks like another live option, albeit a thinly sketched one.

I do not mean to commit myself to any of these views, nor do I strongly favour any of them over the others. The point I hope to make by introducing them is to show that there are answers available to Purely Evaluative Moral Realists when the question arises as to what the evaluative facts they are willing to admit to their ontology are ultimately facts about. (There may be independent reasons to favour or reject any of these views, but here I am speaking about those considerations which are reasons for Purely Evaluative Moral Realists in particular.)

²⁵¹That said, if we were to pursue such an investigation, it would be worth bearing in mind that the answer would not, on this view, need to warrant any kind of special regard. Rather, it would just need to be distinctly moral in its areas of interest.

3.1.2 Contrasting Purely Evaluative Moral Realism with Similar Views

Purely Evaluative Moral realism is not the first metaethical theory to deny special normativity to moral facts, but the commitments mentioned above set it apart from previous attempts at such a model.²⁵² I will examine some of those attempts here for contrast.

One example would be what has been called moral realism *without clout*. This view, or perhaps class of views, is a kind of naturalist realism which holds that moral norms (that is, normative moral facts) exist, but the normativity they possess isn't the special normativity posited by Normativists but the same contingent kind possessed by many other, less controversial, kinds of fact.²⁵³ Such views are actually immune to the Argument from Queerness in the same ways as Purely Evaluative Moral Realism (as discussed in §3.2), but they differ in that they are still concerned with reasons and obligation. That is, they don't deny the existence of normative moral facts—in fact, they may take morality to be purely normative—just so long as the normativity in question isn't special.²⁵⁴ This normativity can

²⁵²These views, discussed below, have not used the term 'special normativity' but they have tried to deny the distinct kind of normativity so often attributed to moral facts with the usual suite of features—universality, independence from agents' desires or opinions, etc.—which is what I mean to capture when I speak of special normativity.

²⁵³See Tresan 2010 and Baeten 2012 for examples. To be clear, the normativity in question could be non-special even if moral norms are applicable to, say, all humans, just so long as it is contingent on something that, as a matter of fact, is true of all humans.

²⁵⁴Observant readers may note that these commitments leave advocates of such views free to remain neutral on the relationship between norms and values. This might make realism without clout preferable to the purely evaluative view—assuming, at least, that neutrality on all these questions is a virtue—and lead readers to wonder why I have devoted this project to Purely Evaluative Moral Realism. However, the overview I have given here is deliberately thin and shows only what these views have in common and how it stands in contrast with my Purely Evaluative Moral Realism. Each comes with a defence of the possibility of realism without clout—effectively an alternate set of arguments for Non-Normativism along a very different line

still be a feature of the facts themselves at the most basic level, which Purely Evaluative Moral Realism pointedly denies. For the Purely Evaluative Moral Realist the moral domain consists of bare matters of fact whose normative significance, if any, must be imposed post hoc by beings of the kind to whom reasons can apply, which is to say by agents.

That leads me to my next point of contrast. Given the connection—or, according to some thinkers, the identity relationship—between special normativity and categorical imperatives, it is natural to draw a similar connection between non-special or familiar normativity and hypothetical imperatives. Hypothetical imperatives are those that are conditional on some particular end that needn't be held by all agents; '*If you want to visit the Parthenon, you should go to Greece.*', and '*If you want a good night's sleep, you ought to drink a glass of warm milk before bed.*' and so forth. Another attempt to excise special normativity (under the guise of categorical imperatives) is Philippa Foot's idea of morality as a system of hypothetical imperatives, where the antecedent is always something like '*If you want to be moral...*' or '*If you care about morality...*'.²⁵⁵ Purely Evaluative Moral Realism is compatible with such a picture as a model of how moral norms work—in fact I will go on to suggest something along these lines in Chapter 4—but the difference again lies in where the view in question puts its focus. Like realism without clout, Foot's model has it that moral claims are fundamentally normative claims, albeit conditional ones.

The other differences between these views and the one I am advancing are subtle.²⁵⁶ In fact, their core claims may be compatible, depending on how they are

than my own—which involves enough additional commitments that Purely Evaluative Moral Realism is, at least, no worse in this way.

²⁵⁵See Foot 1972. Foot (2001, pp. 10, 13-7) did later abandon this model because she came to reject the Humean understanding of reasons for action, but it still provides a valuable point of comparison.

²⁵⁶There is also one fairly clear difference in that these accounts are both naturalistic, where Purely Evaluative Moral Realism is compatible with both moral naturalism and moral non-naturalism. This allows a Purely Evaluative Moral Realist to hew closer to the commitments of Moral Error Theorists, which is good for our purposes, but does not necessarily convey any

understood.²⁵⁷ Purely Evaluative Moral Realism is meant not just as a claim about which moral facts there are, but also as a claim about what the most fundamental moral facts are like. It concerns the basic structure of morality. In other words, the claim that Purely Evaluative Moral Realism makes is that evaluative moral facts are the ‘ground level’ of moral facts. And, while the view allows non-special moral norms about which true claims can be made, these are explained by (or derived from) the evaluative facts—and it is not clear whether there is any single direct route from values to norms, as I will discuss in §4.2. If we understand realism without clout and morality as a system of hypothetical imperatives as claims about the nature of the most basic moral facts, these two views will not turn out to be compatible with Purely Evaluative Moral Realism, but it’s not clear that this is how we are meant to take them. They could also be taken as facts about the structure of moral norms or imperatives without any commitments about which of normative moral facts and evaluative moral facts is prior (or reducible) to the other. That Purely Evaluative Moral Realism is explicit about what priority relationship must hold between norms and values (if any) is a novel aspect that sets it apart from other views that may get around the Argument from Queerness in the same way discussed in §3.2.1.

It may also be a point in its favour over these views. Where special normativity must be understood as ontologically basic, the same is not true of non-special normativity. But the fact that non-special norms can be explained—which is how they avoid being queer—also means that our fundamental moral ontology must be sufficient to explain them, and neither of the other views we have considered offers such an explanation.²⁵⁸

deeper theoretical benefits, so I will not devote any further space to it.

²⁵⁷If we were to hold an entailment theory about value and normativity, then these views would simply be elliptical for one another. Such a theory, however, is one a Purely Evaluative Moral Realist must reject from the beginning, so Purely Evaluative Moral Realism would be incompatible with both given that background assumption.

²⁵⁸Here, I claim only that Purely Evaluative Moral Realism requires that any non-special norms be explained by more fundamental evaluative moral facts. In Chapter 4, I will offer an account of this derivation, but it is strictly additive to the view described in this chapter, which does not

3.2 Purely Evaluative Moral Realism and the Argument from Queerness

In this section I will argue that Purely Evaluative Moral Realism can successfully overcome the Argument from Queerness. Most importantly, it gives us a direct response to the Argument from the Queerness of Special Normativity, so I will first address how it does so. It also allows us to circumnavigate the Argument from the Queerness of Supervenience, which is a less serious threat, but this is still a point in Purely Evaluative Moral Realism's favour, so I will cover this as well.

3.2.1 Purely Evaluative Moral Realism and the Argument from the Queerness of Special Normativity

As stated in Chapter 2, the Argument from the Queerness of Special Normativity can only rule out all moral facts if we are committed to Strong Normativism. Purely Evaluative Moral Realism entails the direct denial of this position, Weak Non-Normativism. It then goes further in saying that not only can there be moral facts that aren't specially normative, but that these facts are the only moral facts that obtain. In other words, it takes the feature of moral facts that the queerness argument from special normativity nominates as the source of their queerness and denies that it's a feature of any actual moral fact.

(At the same time, Purely Evaluative Moral Realism can accept that there are *putative* moral facts—that is, moral claims—that are specially normative, but which aren't actual facts because they aren't true. Weak Non-Normativism and Weak Normativism are compatible, and Purely Evaluative Moral Realism is compatible with the claim that some kinds of moral facts are specially normative. At the same time, the purely evaluative view entails that if there is such a kind, no moral fact of that kind exists.)²⁵⁹

Special normativity, as has been mentioned, is a somewhat nebulous concept

require that there be any moral norms, be they special or non-special.

²⁵⁹As noted previously, to say that some moral facts belong to such-and-such a kind need not imply that any moral facts of that kind exist, per Priest (2008, pp. 42-3).

which encompasses the many attempts to characterize the distinct and unusual kind of normativity so often attributed to the moral domain, whether this be categorical normativity, as suggested by Joyce, irreducible normativity, as suggested by Olson, or inherent and non-relational normativity, as suggested by me. Moral Error Theorists have focused on different details, but they have all been targeting this distinctive kind of normativity. By only embracing evaluative claims (and committing itself to Weak Non-Normativism), Purely Evaluative Moral Realism sidesteps queerness arguments based around any conception of moral normativity as special normativity. At the same time, special normativity shouldn't be taken to include any examples that we can plausibly account for by way of some sort of reduction, such as prudential concerns. Accordingly, Purely Evaluative Moral Realism could remain neutral on the status of non-special norms, both in general and those that might be derived from evaluative moral facts.

To summarize, Purely Evaluative Moral Realism not only admits of moral facts which aren't problematized by the Argument from the Queerness of Special Normativity, but also cannot admit of the kind of moral fact that said argument does threaten. As such, none of the facts it posits can be called queer in this way, and it necessarily avoids this version of the Argument from Queerness.

3.2.2 Purely Evaluative Moral Realism and the Argument from the Queerness of Supervenience

I have already acknowledged that the Argument from the Queerness of Supervenience is not, as things stand, sufficient to motivate a commitment to Moral Error Theory. That said, some philosophers are willing to seriously endorse it, and I remain open to the possibility that an argument for Hume's Dictum or against the possibility of companions in guilt could be given, and so there is some merit in showing how the purely evaluative view avoids it. Purely Evaluative Moral Realism doesn't relate to the Argument from the Queerness of Supervenience in so straightforward a way as it does to the Argument from the Queerness of Special Normativity. This is because Purely Evaluative Moral Realism has nothing to say about the exact nature of the relationship between moral and natural* facts. So we can't take the same approach as we did above. Besides this, I have already

admitted in §1.7 that we can't reasonably deny this supervenience relationship, so if Purely Evaluative Moral Realism can overcome this incarnation of the Argument from Queerness it must be in another way.

To do that, we must ask whether there is some view that undermines the claim that supervenience is queer, whilst also being compatible with Purely Evaluative Moral Realism. I believe that there is one—namely, moral naturalism.²⁶⁰ As I've already noted, Purely Evaluative Moral Realism is neutral on the question of naturalism versus non-naturalism, so there could be both naturalist and non-naturalist versions of it. Under naturalism there is nothing queer about the supervenience of the moral on the natural* because the relationship isn't brute and inexplicable as it is under non-naturalism, but entirely understandable and expected, because supervenience, like entailment, is reflexive. In other words, every fact is supervenient upon (or necessarily covaries with) itself.²⁶¹

This move may raise some eyebrows. After all, moral non-naturalism is one of the assumptions of the Argument from Queerness with which I've been going along up to this point. The only assumption that I have rejected until now is Strong Normativism, because I have been trying to engage with Moral Error Theorists on their own terms as much as possible. As such, suddenly rejecting another assumption may seem illegitimate. Closer examination, however, shows that I have license to do away with the assumption of non-naturalism precisely because I have already done the same to the assumption of Strong Normativism.

²⁶⁰For the purposes of our discussion here, we can treat moral supernaturalism as a kind of naturalism rather than non-naturalism (see §1.7). While this isn't strictly correct, the relevant feature here is that the views in question treat moral facts as identical to some facts about the world, whereas non-naturalism holds that they are distinct from but covariant with such facts. We could perhaps think of such views as moral naturalism*, but I will stick to discussing naturalism as it has more serious adherents and calls for none of the extra ontological commitments that supernaturalism does.

²⁶¹This is because, under naturalism, moral facts are nothing over and above natural* facts. Naturalism doesn't tell us which natural* facts determine moral status or why they do so, and these relationships may be much less clear, but these are questions of substantive ethics that needn't bother us at this juncture.

This is because of the kind of concerns that typically ground a commitment to non-naturalism, and to which Mackie and many other Moral Error Theorists appeal. Non-naturalism, they assert, is the only viable view regarding what kind of fact moral facts are because naturalism can never account for the normative nature of moral facts.²⁶² This is meant to justify non-naturalism as a claim about moral facts in general, so we must take this as meaning that all moral facts have a normative nature that cannot be accounted for naturalistically. But this is just Strong Normativism. So the Moral Error Theoretical commitment to non-naturalism is grounded in the same assumption of Strong Normativism as the Argument from the Queerness of Special Normativity, which is the same commitment that one must already have rejected in order to be a Purely Evaluative Moral Realist.

I am not saying that this is the only concern that ever grounds commitment to moral non-naturalism, but it is one of the most straightforward, and one that some moral antirealists appeal in their own arguments.²⁶³ In the absence of any other arguments, anyone who has rejected (or, at least, reserved judgement on) Strong Normativism, as Purely Evaluative Moral Realists have, needn't be committed to non-naturalism.²⁶⁴

Hence there are at least two versions of Purely Evaluative Moral Realism: a naturalistic one and a non-naturalistic one. As we have already seen, Purely Evaluative Moral Realism itself is neutral between naturalism and non-naturalism, whereas more conventional forms of moral realism (with a distinctly normative element) seem to necessitate non-naturalism. Despite this neutrality however, there is some independent reason to favour the naturalistic account of Purely Evaluative Moral Realism because it is more parsimonious than the alternative. We have already seen why naturalistic views undermine the queerness argument

²⁶²For example, see Mackie 1977, p. 33 and Parfit's (2011a, pp. 324-7) original normativity objection (as distinct from Olson's interpretation, which I discussed in §2.1).

²⁶³For example, see Mackie 1977, pp. 38-40.

²⁶⁴It follows from this that if some sufficiently strong independent argument for moral non-naturalism can be provided, then Purely Evaluative Moral Realism loses its immunity to the queerness argument from supervenience. It would still, however, circumvent the more widely endorsed queerness argument from special normativity.

from supervenience, and we have seen in §3.2.1 why Purely Evaluative Moral Realism must always undermine the queerness argument from special normativity. So what we have here is a version of Purely Evaluative Moral Realism that is immune to every functional version of the Argument from Queerness.

Hence Purely Evaluative Moral Realism can avoid the Argument from Queerness. It is inherently immune to the version of the argument which poses the most serious threat to moral realism—the Argument from the Queerness of Special Normativity—and can be made proof against the other (potentially) threatening version—the Argument from the Queerness of Supervenience—with only one extra commitment, which comes at little theoretical cost. These traits are shared with other forms of naturalistic moral realism but unlike them it does not need to explain how the moral facts it posits came to have a normative structure and it is, in this regard, a superior answer to the problems raised for realism by the Argument from Queerness. Whether the extra commitment to naturalism is necessary for this purpose will depend on how much stock one puts in the queerness argument from supervenience, but in either case we have a way around the issues raised by the Argument from Queerness.

3.3 Objections to Purely Evaluative Moral Realism

Despite my remarks above, it's too early for us to close the book on our central question. What I have described in the preceding sections are *prima facie* reasons to think that Purely Evaluative Moral Realism gives us a way around the Argument from Queerness. The view is such that the problems on which the Argument from Queerness turns simply don't apply to it. But is this loophole plausible?

Purely Evaluative Moral Realism might seem immediately counterintuitive to many readers; it's one thing to suggest that moral facts are not necessarily all specially normative, and quite another to suggest that none actually are. As such, there are some obvious objections I must address. Some of them I will be able to rebut fairly quickly, while others will raise questions that need considerably more involved answers, and this will lead us into the final chapter.

3.3.1 Restriction on Substantive Theories

One line of argument goes that Purely Evaluative Moral Realism would be overly restrictive with regard to which theories within substantive ethics could be true. With some slight modification consequentialist and virtue-based theories can be made amenable to Purely Evaluative Moral Realism but the same cannot be said for deontological systems without significant reinterpretation (see §4.1). Under deontological views, morality is just a system of deontic (and so normative) claims, which is obviously incompatible with Purely Evaluative Moral Realism.²⁶⁵ (Deontologists cannot even compromise with hypothetically normative facts, as Purely Evaluative Moral Realism rejects the idea that normativity of any kind can be found in the base-level normative facts themselves. This is how it differs from the other views discussed in §3.1.2.) Of course, ruling out some particular first-order theory needn't necessarily be a problem. After all, Moral Error Theory rules out all such theories.²⁶⁶ It would, however, be an issue if Purely Evaluative Moral Realism ruled out highly plausible views, which a Normativist may claim is what it does when it eliminates views that posit specially normative moral facts. My

²⁶⁵We will ignore here the possibility that all moral systems can be consequentialized—as suggested by thinkers like Dreier (1993, p. 23), Louise (2004, p. 518), and Portmore (2007), and objected to by Brown (2011)—for a few reasons. Firstly, the claim that they can is controversial enough that it would be unwise to premiss our whole response on it. Secondly, assuming that we can do this with the inherently normative deontological theories seems to presuppose a form of consequentialism that itself outputs normative facts and so may still be incompatible with Purely Evaluative Moral Realism depending on how the norms in question are understood. Finally, as I argue below, we can still answer this criticism even while accepting that Purely Evaluative Moral Realism may exclude deontological theories, and we should address the strongest versions of any objection where we can.

²⁶⁶The crux of the issue may be that Purely Evaluative Moral Realism is insufficiently neutral between substantive theories; Moral Error Theory is, at least, even-handed in its dismissal. However, neutrality is a rather weak virtue. If there are correct and incorrect answers to moral questions, it is good for a metaethical view to rule out incorrect answers and leave room for the correct ones.

response would be that Normativist metaethics is also incompatible with a range of highly plausible substantive moral views. It necessarily rules out any view that is compatible with Purely Evaluative Moral Realism, like the ones I will describe in §4.1. The views in question have enough precedent outside the context of this debate that Normativist metaethics' inability to account for them is no less of a problem than Purely Evaluative Moral Realism's inability to account for strictly deontological systems. So Purely Evaluative Moral Realism cannot be rejected on these grounds.

3.3.2 Queerness Arguments

Another way to object to Purely Evaluative Moral Realism is to submit that evaluative moral facts are themselves queer, even if they don't entail specially normative facts. But are there any successful queerness arguments that nominate value itself as queer? Some incarnations of the Argument from Queerness have been framed in such terms.²⁶⁷ Such arguments, however, usually treat talk of value as encompassing the whole moral domain, and then to turn on some specific feature of value such as are central to the queerness arguments discussed in Chapter 1. To the best of my knowledge, no queerness argument targeting evaluative facts (as distinct from normative ones) in particular exists, but our discussions up to this point give us materials enough to make some suggestions, and I would be remiss in failing to examine these. I am referring to the various accounts of what makes special normativity queer. For each I will have to isolate what it means for value to be queer in the same way as has been suggested for special norms, and then determine both whether it is conceptually true that moral facts have this feature and whether that feature would be sufficient for queerness. As will become apparent, the arguments by and large do not survive this process, in large part due to the disanalogies between normativity and value.

Since I have claimed that moral normativity is queer because it is special, and special normativity is quite loosely characterized, the claim that evaluative moral facts are queer in the same way as normative ones could mean a lot of different

²⁶⁷See Mackie 1977, p. 15 for one significant example; from the outset he speaks of 'objective values'.

things. One would be a *de dicto* reading; special normativity is the distinctive kind of normativity peculiar to moral facts and is queer in how it stands apart from other cases of normativity, so the evaluative nature of moral facts would have to be queer in how it stands apart from other kinds of value. Certainly, there are other kinds of value, such as aesthetic value and pragmatic value, but we would need to determine whether and how moral value stands apart from these. Perhaps the intuitive answer to this question is that there is an onus on all of us to respond to moral value in a way that doesn't hold for other kinds of value. But surely this is just attributing to evaluative moral facts special normativity in a way which will seem illegitimate to the Purely Evaluative Moral Realist, who has already rejected Strong Normativism by denying any necessary entailment from moral value to special normativity.

The other route open to those who want to argue that value is queer is to invoke one of the more particular understandings that special normativity is meant to encompass. As we will see, however, we will have to interpret these arguments creatively if we are to apply them to value, and none of these interpretations will turn out to be particularly compelling. In fact, they all fail in one of the two most obvious ways for a queerness argument to do so; the feature they nominate as queer either is not plausible as a conceptually necessary feature of evaluative moral facts, or it turns out not to be particularly unusual or problematic.

Joyce suggests that the queer feature of moral normativity is its categoricity. When a norm is categorical it applies to all agents regardless of their ends. But, as I noted in §1.8, categoricity alone is insufficient for queerness. The problem comes from treating claims about normativity as though they apply categorically when they are otherwise typically contingent on facts about particular agents in such a way as makes them highly personal, or agent-relative. To see whether an analogous queerness argument can be made, then, we must determine what it would mean for a value to be categorical and whether this would actually be problematic.

This is difficult from the first step, as it's not clear what it would mean for moral value to be categorical. Categorical reasons are reasons for each and every agent, which is to say that they are agent-neutral, so perhaps agent-neutrality is relevant to the supposed queerness of moral values. However, there is an important disanalogy between reasons and values here. There are no normative reasons that

do not ‘belong to’ an agent, so categorical reasons are reasons for each and every agent. We need not understand value in an isomorphic way. Not all goodness is necessarily goodness for someone; some goodness is goodness *simpliciter*.²⁶⁸ According to the Purely Evaluative Moral Realist, evaluative moral facts are just a subset of the descriptive facts about the world. There is nothing queer about descriptive facts being agent-neutral.²⁶⁹ As such, this feature cannot be the focus of a plausible queerness argument. While there are forms of goodness that arguably vary with some particular facts about specific agents (such as their desires), any forms of value that are not subject to such dependencies are ontologically and metaphysically inoffensive because agent-neutrality is a standard feature of familiar descriptive facts that any ontology should account for. So even if moral facts are ‘categorical’ in this way (and each of this reading of ‘categorical’ is something of a stretch), this wouldn’t be grounds to call them queer.

Another suggestion of what exactly is special about special normativity is Olson’s idea that it is irreducible. In the case of normativity, this is sufficient for queerness because, Olson contends, it requires us to make a significant ontological commitment to accommodate it in a way that isn’t true of the reducible normativity he attributes to pragmatic and conventional norms. I’m quite happy to grant that there are kinds of value that are determined relative to particular standards or conventions. Goodness of a kind, for example, is dependent on the kind in question; a good knife is one that cuts well. However, it’s unclear whether moral value would be alone in breaking this paradigm, and the more companions-in-guilt

²⁶⁸Plausibly anything that is good for someone is also to that extent good *simpliciter* but this need not exhaust the ways in which something can be good in itself. Whether it is possible for something to be good without being good for someone (that is, whether all goodness is ultimately personal goodness) is a subject of ongoing philosophical debate.

²⁶⁹There are readings of agent-relative value such that it is something distinct from goodness for, but these have no obvious equivalent in natural language and it is controversial even among realists what these values are and whether there are any—see Schroeder 2007, pp. 270-7. In any case, the presence of such values would do nothing to suggest agent-neutral values were queer unless they could be used to account for all non-moral forms of value.

it has, the less plausible this would be as a queerness argument.²⁷⁰ Things like truth, beauty, and posterity are often supposed to be valuable in themselves without reference to anything further. One could of course offer reductive accounts for any of these, but to back up a queerness argument based on the irreducibility of moral value one would have to offer accounts for each potential companion-in-guilt, or an all-purpose account that managed to exclude moral value and nothing else. So it's not clear that evaluative moral facts would be different enough to evaluative facts from other domains to be called queer.

But even supposing that they were, perhaps just being irreducible is enough to call all these kinds of evaluative facts into question. In the absence of any reduction relationship, we can simply apply the queerness argument from supervenience to such evaluative facts, because if they co-vary with certain natural facts (which moral facts certainly do and their companions-in-guilt plausibly do) they will still require us to posit a brute supervenience relation between the two categories. Of course, this is assuming that the Purely Evaluative Moral Realist grants that moral facts are irreducibly evaluative. As discussed in §3.2.2, they are not forced to do so. If moral facts in general are subject to some kind of naturalistic reduction, there is nothing queer about the relationship between evaluative moral facts and natural facts. A potential objection here is that there could be something queer about the fact that some natural* facts are evaluative where others aren't.²⁷¹ But this kind

²⁷⁰Paul Noordhof has pointed out to me that, for precisely the reason just given, if one accepted both Olson's queerness argument from special normativity, where the specialness is understood as irreducibility, and the queerness argument attempted here for moral value, one would find that each undermined the other by offering a companion-in-guilt for their irreducibility. This initially sounds right but the crux of the issue is not just irreducibility. Irreducibility on its own is not puzzling, as all ontological primitives are by definition irreducible. The issue for both Olson's queerness argument and the analogous one I am describing here is that we have an irreducible example of a kind—norms in the one case, values in the other—whose other members, the arguments posit, are reducible. It is their status as exceptions that makes them queer, and that they are both exceptional among their classes in the same way is, I think, insufficient to make them companions in guilt.

²⁷¹I owe this observation to Stephen Everson.

of relationship is not unusual enough to warrant an attribution of queerness. Some natural facts* are facts about colour where others are not, and some natural* facts are facts about weight where others are not.²⁷² Hence, the naturalistic version of Purely Evaluative Moral Realism (which is open to us whether or not we feel the need to address the queerness argument from supervenience) is immune to a queerness argument from the irreducibility of moral facts.²⁷³

We could even appeal to the suggestion from §1.8 that the queerness of special normativity was due to the way that moral facts are often treated as normative *simpliciter* in and of themselves, whereas we can understand non-special norms as grounded in a relation or set of relations between agents, actions, and outcomes. If we try to make an analogous argument about non-relational value, we end up saying that evaluative moral facts are queer because the value of morally valuable things isn't grounded in this same kind of relationship. For this to be so, non-moral value would have to be necessarily grounded in such relationships, and never intrinsic to its bearers. In other words, the queerness argument we have constructed by analogy turns on a conceptual claim that this is the key difference between moral value and non-moral value. But there is really no good reason to buy into this claim. In fact, we have good reasons to reject it, and so any queerness argument that relies on it.

Firstly, it doesn't seem plausible that it's a conceptual truth that moral value is non-relational; plenty of thinkers take moral value to be directly grounded in personal goodness. John Harsanyi, to give just one example, argued that general value, which is to say value *simpliciter*, is simply what is preferred impersonally (in the sense that it makes no reference to the preferrer's own place in the situation

²⁷²These examples also plausibly reduce to other natural* facts or sets thereof (facts about light and human perceptual systems in the case of colour, and facts about mass and gravity in the case of weight) so their doing so is not a point of disanalogy that could be welded against my counterargument.

²⁷³Some philosophers, such as Schoedinger (2007) and Wedgwood (2007, ch. 9), have offered versions of moral naturalism which are, by their own lights, not reductive, and would thus still be no less vulnerable to this queerness argument than a non-naturalist version.

considered).²⁷⁴ Of course, some thinkers will want to insist that moral value really is inherent to its objects. But, they needn't say that it is always of this kind, nor that moral value is unique in taking this form. For instance, Moore insisted that a beautiful uninhabited world is of greater value than an ugly one, and that its value is intrinsic and held regardless of the fact that there is no one to appreciate its beauty.²⁷⁵ He grants that value can work in the way described, but makes the same claim about a value some would call distinctly non-moral. My point in citing these examples is that there is nothing especially unreasonable or implausible about them, yet they fly in the face of the commitments of this queerness argument. If this argument is to work, then we are owed some persuasive account of how it is that moral value is always and only non-relational, as well as of how it is that other forms of value are always relational, despite the presence of alleged counter-examples which are *prima facie* plausible in both cases. I doubt that such a case is forthcoming, and until it is given I see nothing contradictory in denying any of these claims. Moral Error Theorists cannot, therefore, call the central premise of this would-be queerness argument a conceptual truth about morality and so it fails.

It should be apparent, at this point, that existing queerness arguments give us no grounds to be suspicious of evaluative moral facts. Evaluative moral facts needn't make a view committed to them vulnerable to the queerness argument from supervenience nor the queerness argument from special normativity, and the concerns that motivate the latter fail to ground any further objection to them as well.

3.3.3 Purely Evaluative Moral Realism and Action

Several major objections to Purely Evaluative Moral Realism are concerned with the connection between morality and action. Some will want to insist that moral facts, as a matter of conceptual necessity, must entail specially normative facts,

²⁷⁴Harsanyi 1953, p. 434. Harsanyi also argues that there is only one rationally acceptable set of impersonal preferences, and so there is an objective fact of the matter about what is valuable

²⁷⁵Moore 1903, p. 84.

even at this point. But this is just Olson's reading of the normativity objection, and I've already dealt with it in §2.1.

Others might grant that it's up in the air whether each individual moral fact must be a specially normative reason, but submit that Purely Evaluative Moral Realism is nevertheless incoherent because special normativity is essential to the moral domain as a whole. This is importantly different from both Strong and Weak Normativism (though it entails the latter). It is the claim that if there are any moral facts at all there must be at least one normative one. This holds whether or not there are any moral facts that do not entail specially normative facts. We have already dealt with entailment views, whereby each and every moral fact is a specially normative reason, in §2.2 but here I am addressing a more general philosophical view about the moral domain, which might be held without reference to any more specific justification.

This is the view Joyce seems to be committed to.²⁷⁶ If such a view is correct, then any form of moral realism will necessarily be vulnerable to the Argument from the Queerness of Special Normativity. Yet the claims it makes—which will likely strike many as quite intuitive—are more modest than those of Strong Normativism, and are even compatible with Weak Non-Normativism. As such it would be quite hard to argue against. I certainly don't have a knockdown argument against it. But then, positing this view isn't itself a knockdown argument against Purely Evaluative Moral Realism either, and there are some questions to be asked about the Joycean view.

Firstly, could the Joycean view be justified without appeal to intuition? Joyce justifies his claims by a 'translation test'; whatever must be preserved in a translation from one language to another is an essential part of the meaning of a term or sentence, and he takes it that normativity is such for (at least some) moral language.²⁷⁷ However, this stands more in support of the claim that this is how our moral psychology works (insofar as we can assume that moral language is re-

²⁷⁶Joyce (2006, p. 61) remains neutral on whether there are any moral facts that aren't specially normative, but maintains that the presence of at least some specially normative facts is necessary for the domain in question to be morality.

²⁷⁷Joyce 2001, p. 5.

flective thereof), which I have already granted in §1.3, and which needn't imply that moral facts are structured in the same way. Moral Error Theorists will likely contend that moral facts are specifically those we mean to refer to (whether it's possible to do so or not) when we use moral language, and so their structure will always be what is indicated by moral language and psychology. But if there are moral facts, then I assume that we should all grant that it's possible for people to have false first-order moral beliefs, and more to the point that few if any people have only false moral beliefs, or only true ones. If this kind of mixture of true and false belief is possible with regard to substantive morality, shouldn't we suppose that the same is true about our metaethical assumptions? Even very central assumptions like this one, grounded in our moral psychology and reflected in our moral language, needn't be immune to revision if we grant that there are different kinds of moral fact—evaluative and normative—about which we can correct or incorrect independently.

Second, what kind of realists are Purely Evaluative Moral Realists, if not realists about morality? We can't deny that they *are* realists; one of their central claims is that certain facts do exist, and that those facts form a significant subset of those that Moral Error Theorists would normally be happy to call moral facts. The Joycean would need to explain how decoupling these facts from specially normative moral facts, or just from the claim that the normativity of moral facts is special, makes them into facts of a different domain altogether. I won't assert that these questions are beyond answering, but they do impose a burden that the Joycean will have to discharge.

It's clear that a view whereby the moral domain must contain at least one specially normative fact is incompatible with Purely Evaluative Moral Realism, but it should also be clear from my remarks above that such a view needs some development before it can be levied as an actual objection thereto. But there's another objection grounded in the connection between morality and action, and it warrants a more in-depth response. This objection begins with the observation that ethics is largely concerned with human action, with which I am inclined to agree. From there it goes to say that substantive ethics must be fit to guide our

actions; it must give us some way to decide what to do.²⁷⁸ Again, I have no issue with this premiss, and I would certainly look askance at a purported ethical theory that couldn't—at least in principle—be applied to help us make decisions about how to act. The objection, then, is that Purely Evaluative Moral Realism somehow fails to live up to these requirements.²⁷⁹

There are two ways I can think of for those who would make such a claim to ground it. Both are rooted in Purely Evaluative Moral Realism's rejection of special normativity, and both fail to motivate an outright rejection of the view I have outlined in this chapter for reasons I will explain presently. The first is that Purely Evaluative Moral Realism posits an objective fact of the matter about the goodness and badness of actions, states of affairs, and so on, but it doesn't tell us why we should care about such things, and so it leaves the moral domain disconnected from our actions. The basic observation is correct; there is nothing in Purely Evaluative Moral Realism itself that gives us such reasons. But I wonder if a Normativist realism is any better off in this regard. At least some of the propositions within such a system would be specially normative and thus would apply to all of us, and one cannot remain indifferent in the face of an acknowledged normative reason on pain of irrationality.²⁸⁰ But why should anyone acknowledge the reasons such a system posits? It seems that any system of ethics is unable to give us a reason to care about its content outside of the system itself. (Buying into a Normativist system is, in effect, acknowledging the reasons it posits as normative in the first place, and specially normative to boot, so the failure to respond to such putative reasons may look very out from within such a system.) Some views try to

²⁷⁸Note that this doesn't mean that a given theory's answers to moral questions should always be obvious—though some theories are criticized for being unclear in this way—but that they're clearly meant to offer some kind of recommendation on how to act.

²⁷⁹See, for example Ewing 1947, pp. 57, 134 for the kinds of thoughts that motivate these objections. His remarks indicate that he thought ethics couldn't do without 'ought', and though he took it that the moral 'ought' was a distinct kind we needn't say he thought it was specially normative, as another 'ought' he discusses (that of the fittingness invoked in the buck-passing account of value discussed in §2.2.4 and §2.3) might also fill those shoes.

²⁸⁰Joyce 2001, pp. 80-1.

avoid this problem by subsuming morality under another normative domain like practical rationality or—in the case of views which posit a God who metes out punishment and reward—prudence, but then we have to wrestle with questions of whether the norms invoked are strictly moral, and whether the domains invoked really generate special norms.

Every system of morality will be open to these kinds of accusations of groundlessness, but this needn't be as big a problem as many have taken it to be. Ethics may simply not be equipped to tell us to be good. After all, ethical investigations are premised on our pre-existent investment in questions about the moral status of actions. Instead, it may be that the function of ethics is to tell us *how* to be good, with the assumption that this is something we want—or otherwise have independent reason to pursue—already in place.²⁸¹ This response may fail to satisfy some, but going the other way and trying to claim that we do have reasons to care about morality requires us to justify that claim. This is why so many have posited that moral reasons are categorical, or irreducible, or inescapable, but it is the difficulty of justifying these claims that have led many others to antirealist conclusions about ethics. Special normativity makes a nonsense of the question 'Why be moral?' but fails to do so in a satisfying way, so I believe that my answer—acknowledging that there is no response which will invariably show sceptics and amoralists that they are wrong or irrational in their failure to treat these considerations as reasons—is no worse by that metric.²⁸²

²⁸¹This presupposes the existence of some kind of self-consciously moral motive, a *de dicto* desire to be moral. I think that we do indeed possess such motives, and I will discuss this in more depth in Chapter 4.

²⁸²When the issue of moral normativity is framed in terms of the question 'why be moral?', some philosophers suggest that a satisfactory answer is—and can only be—found from within a moral point of view that can be held in common between different people (see McDowell 1979, pp. 341, 346; Wiggins 1987, p. 123 and 1990, p. 70; and Blackburn 2001, p. 133). Insofar as an agent occupying such a perspective might well be subject to moral reasons, I have no objection to this, but the question for our purposes is whether there are specially normative reasons to occupy this moral point of view. If not, the reasons that come with it need not be specially normative in the way Normativists (including Moral Error Theorists) want to insist moral facts

The second objection from the relationship between morality and action is simply that Purely Evaluative Moral Realism, or any ethical theory that doesn't posit specially normative facts, is unable to guide our actions. To make such a claim, one must first implicitly assume that the only way for a fact to guide action is for it to be a specially normative reason. I have my doubts that this is so. After all, the difference between a fact that is normatively inert and one that is a non-specially normative reason is just context. (By this I mean that any fact can be a non-specially normative reason, but what makes it so is the existence of other facts about, for example, some agent's ends or desires.) But as I mentioned above, a lot of ethical investigation is motivated by questions of how we should act. These are often framed in explicitly deontic (and thus normative) terms: 'How ought I to live?' or 'What should I do in such-and-such a situation?' This objection seems to be grounded in two thoughts; that a question framed in normative terms must be answered in normative terms, and that the only sure-fire way of doing so is with specially normative facts. Both are doubtful. We can answer normative questions with those natural* facts that are meant to serve as normative reasons, so the first is plainly false. The second is not so obviously wrong, but also cannot be assumed to be true.

I think that it is possible to account for the action-guiding nature—even the normativity—of moral facts in terms of non-special normativity, and do so while remaining committed to the core claims of Purely Evaluative Moral Realism. Further, I think that in so doing we may isolate the sources of many of the intuitions which motivate the criticisms of Purely Evaluative Moral Realism that I have anticipated here, and find them to be less well-grounded than we might have initially thought. I don't think, however, that I can argue for either of these claims wholly in the abstract. Instead, I will prove that it's possible in the old-fashioned way; I shall offer an attempt at just such a system. This will take extensive argument, and so I will devote the latter part of the next chapter to outlining this system, and arguing for any additional commitments it requires.

are. The moral sceptic can ask whether we have specially normative reasons to occupy the moral perspective from outside of that perspective, and if no answer is available to them there then Wiggins and company have provided no answer at all.

To be clear, the claims and arguments I make in that chapter will be strictly additive to Purely Evaluative Moral Realism as it is presented in this chapter. Here, we have a few core metaethical claims, and in Chapter 4 we will have a much more comprehensive metaethical system which is compatible with those claims. I don't purport that what I develop therein is the only possible way of extending Purely Evaluative Moral Realism. Instead, it's just one possibility among many, albeit the one that seems most intuitive to me. It follows from this that if my arguments for this extended system fail anywhere, then I may have failed to prove that we can extend Purely Evaluative Moral Realism in such a way that it satisfyingly accounts for moral action-guidingness and normativity, but it will not thereby be shown that all attempts to do so are doomed to failure.

We might also build upon Purely Evaluative Moral Realism by showing how it is—or can be made—compatible with existing first-order ethical theories. If this can be done, then Purely Evaluative Moral Realism can guide action by just the same token that the theories in question can. Again, these are expansions upon the core claims of Purely Evaluative Moral Realism, so they too will be explored in Chapter 4.

In this chapter, I have laid out what Purely Evaluative Moral Realism is: the ontological claim that the only moral facts that exist are evaluative and not specially normative. We have seen to which other positions within metaethics it is committed, and with which ones it is compatible and incompatible.

I have also addressed the question of whether Purely Evaluative Moral Realism constitutes an out or loophole by means of which we can avoid the Error Theoretical conclusion of the Argument from Queerness. We have seen how the specific commitments of Purely Evaluative Moral Realism make it immune to the Argument from the Queerness of Special Normativity. We have also seen that it is compatible with naturalism, which is immune to the Argument from the Queerness of Supervenience. This means that there are versions of Purely Evaluative Moral Realism which resist the two formulations of the Argument from Queerness that pose the most legitimate threat to moral realism.

Finally, I argued that many of the immediate objections to Purely Evaluative Moral Realism fail. But I have been able to offer less argument against one

objection in particular. This is the claim that Purely Evaluative Moral Realism fails to be action-guiding in the way we would typically expect of morality. In the next chapter, I will argue against this contention, first by showing how certain first-order ethical views are compatible with Purely Evaluative Moral Realism, and secondly by showing how this view still allows us to make some true normative claims grounded in moral facts and how the intuitions that seem to undermine it have come about.

Chapter 4

Moral Facts and the Guiding of Action

So far, we have seen how the Argument from Queerness is meant to threaten moral realism, and how Purely Evaluative Moral Realism is immune to the Argument from Queerness. But, as I have already stated, this is not quite enough to answer our central question. It does not matter that the view I have outlined is immune to the Argument from Queerness unless said view also passes the basic tests of coherence and adequacy as an account of the structure of morality. If Purely Evaluative Moral Realism is obviously false, then whatever other features it has are of little interest.

Much of what I have already said in §2.2, §2.3 and §3.3 has been aimed at showing that it passes these tests. However, one issue raised at the end of Chapter 3 was left unresolved. This was the question of whether morality, understood through the lens of Purely Evaluative Moral Realism, is fit to guide action in the way we would typically expect of it. In this chapter, I will argue that it is, and I will do so by examining substantive ethical theories which are compatible with the commitments of the purely evaluative view, and by giving a general account of how morality might guide action by drawing on a common element from said views.

4.1 Purely Evaluative Moral Realism and Substantive Ethics

Must ethical theories be able to help to guide us about which actions to take in certain situations? I am tempted to say that they must. If this is so, then by necessity anything we can rightly call an ethical theory is action-guiding in this way. It follows that if there are ethical theories whose content can be accounted for using a Purely Evaluative Moral Realist framework, then Purely Evaluative Moral Realism itself allows for moral facts to be action-guiding and the objection mentioned at the end of §3.3.3 has no merit. In this section, I shall show how a wide variety of substantive ethical views can be made compatible with Purely Evaluative Moral Realism. The first two I will discuss, Scalar Utilitarianism and Virtue Ethics, require little if any modification to reach this point. The last two, deontology and Divine Command Theory, intuitively seem to be strongly incompatible with the purely evaluative view but can still be understood in such a way that they are compatible with it.

In arguing for this compatibility, I shall make the case that there are substantive views compatible with Purely Evaluative Moral Realism and that they really are action-guiding. I am not claiming that Purely Evaluative Moral Realism is compatible with all first-order ethical theories; if we take deontological views at face value they posit a body of moral facts that consists entirely of the kind Purely Evaluative Moral Realism rejects, and the version I present below will strike at least some deontologists as an unacceptable compromise. I claim only that the examples that I give in the remainder of this section are sufficient to motivate and flesh out the claim that one need not posit special normativity to give moral facts the power to guide action.

In order for them to do this, they must provide non-specially normative reasons. If we follow Olson's view (as described in Chapter 1) this means they must be part of the reduction base for these reasons, alongside other facts about the agent or the situation that would normally be necessary for the existence of a reason. With regard to the relevant non-moral facts, I will focus on the agent's desires (in the broadest sense) because they are a feature of agents that is action-structuring by

nature. Additionally, insofar as they turn on the presence of such a desire, the accounts given below should satisfy a wider array of Moral Error Theorists than just those who buy into Olson's model, since there are many views on which a desire (and a means by which it can be attained) is sufficient for the existence of a reason. As such, it is to this that I will appeal in order to bring motive force into the picture below, but it is worth acknowledging that other reduction bases are possible and could be substituted in with minimal revision. I will make the same appeal in §4.2, where I will build a model to show how we might establish that a moral fact can generate reasons that apply generally (though not universally) in such a way as to explain their apparently categorical nature, but the kind of substitution mentioned above will be less feasible there. My purpose in building this model, however, is simply to show that we can do so and end up at something plausible, and it is strictly additive to the purely evaluative view. I accept from the outset that one could be a Purely Evaluative Moral Realist and still reject any further commitments it makes.

4.1.1 Scalar Utilitarianism

When we speak of utilitarianism, we typically mean maximizing utilitarianism.²⁸³ This is the view that the moral status of an action is determined by the balance of pleasure or happiness against pain or suffering that the action produces as follows: the action that produces the *most* (total or average) pleasure or happiness is right and all the alternatives are wrong.

By contrast, on Alastair Norcross's Scalar Utilitarianism, an action is simply better or worse than alternatives in direct proportion to the amount of pleasure or pain each causes.²⁸⁴ No one action out of a set of alternatives is singled out as what *ought* to be done, because Scalar Utilitarianism doesn't deal in direct recommendations. Instead, what is produced is an account of to what extent actions are good and bad, and how they compare to one another.²⁸⁵ To be sure,

²⁸³Indeed, this seems to be true for every form of consequentialism, but I will limit discussion in the body text to utilitarianism for simplicity.

²⁸⁴Norcross 2020, ch. 3. The account that follows is drawn from Norcross 2006 and 2020.

²⁸⁵Norcross (2020, ch. 3) maintains that any attempt cardinal or absolute account of the

the action that produces the most good is still the best, but this is not taken to imply that it is to be recommended uniquely over other available actions.

It's easy to see how a view that doesn't classify any actions as *obligatory* or *forbidden* might be accused of failing to guide our actions, but Norcross anticipates this and offers a response. He agrees that an ethical theory must guide our actions and submits that the basis of this objection is the idea that a theory can guide our actions if and only if it tells us what we *ought* to do. He contends also that this idea is wrong, and that while Scalar Utilitarianism doesn't tell us what we ought to do, it does guide our actions insofar as the relative goodness (and badness) of certain courses of action gives us *reasons* to perform or forbear from doing those actions.²⁸⁶

These two claims might sound very odd against the background of my own discussion of reasons and 'ought' claims. I have consistently held that 'ought' claims are best understood as a subset of normative claims—claims about our genuine reasons for action—focused on the recommendations of the final balance of all of our reasons. What's more, in later work he characterises these reasons as provided by value in a straightforward way that lines up well with the characterisations of special normativity given in §1.8.²⁸⁷ So the moral facts posited by Scalar Utilitarianism should directly determine what any agent ought morally to do, and have some kind of effect on their overall reasons, even if not an overriding one. Since this implication flies directly in the face of the position's other commitments, we can only conclude that Norcross does not share this understanding of 'ought'. So if someone who is committed to the view that what one ought to do is whatever they have most reason to do (as I am), what are they to make of Scalar Utilitarianism?

Norcross calls the directive elements of morality rejected by his view 'de-

goodness and badness of actions will fail because every candidate for a baseline or neutral action has intuitive counterexamples. There is still, however, a fact of the matter regarding the nature and intensity of the betterness relations at play here, and these are still pointedly evaluative. Also, it doesn't follow from this that other potential bearers of value—such as states of affairs—cannot be assessed in cardinal terms.

²⁸⁶Norcross 2006, pp. 46-8.

²⁸⁷Norcross 2020, p. 4.

mands'.²⁸⁸ Moral facts are not agents and do not make demands, so this is best understood as metaphorical language expressing their deontic force. So what the Scalar Utilitarian denies is that there are facts about what we morally ought to do. Given the sense of 'ought' that I have been using, however, the existence of such facts is entailed by the existence of specially normative reasons.

If we are to square Scalar Utilitarianism with this understanding of 'ought', our only recourse is to change our understanding of moral reasons to non-specially normative ones that have force for particular agents because those agents have some kind of desire, goal, or end such that moral facts are normatively relevant for them. This move is a departure from Norcross, but one that preserves the core commitments of Scalar Utilitarianism; it gives us a theory of value which presents us with considerations relevant to our actions—reasons, or at least potential ones—but makes no deontic claims. This leaves us with a version of the view that is certainly compatible with Purely Evaluative Moral Realism.

In this view, the evaluative moral facts provide non-specially normative reasons for action. Since normative reasons can be motivating reasons just so long as an agent believes them and acts as they do on the basis of that consideration, there is no mystery about how those facts can be a guide to action (though see §4.2 for a more detailed account of how evaluative facts generally might serve as a guide to action). As Norcross points out, if beliefs about rightness and wrongness can serve as motivating reasons, then there is no reason that beliefs about relative goodness and badness cannot do the same, and even motivate the same actions. Hence, an agent with the right motives will change their actions in accordance with what they take the moral facts to be, and this is clearly a case of their actions being guided by these facts.²⁸⁹

The only question that remains is whether the argument for the action-guiding power of Scalar Utilitarianism is any good in itself. The version of Scalar Utilitar-

²⁸⁸See Norcross 2020, p. 42, where the demands of morality and 'ought' facts are treated as equivalent.

²⁸⁹I am assuming motivational externalism here, since I rejected internalism in §1.6. Scalar Utilitarianism itself is not committed to either view, and so Norcross (2006, p. 47) remains neutral when discussing it.

ianism I have just developed entails that this power is contingent on agents having some kind of concern over the goodness (or badness) of their actions, what we might call a ‘self-consciously moral motive’. So it seems evident that under this view moral facts *can* guide our actions, but it isn’t the case that they *must* do so.

Some of the objectors will want to insist that morality must be action-guiding in a *necessary* way, but such a claim will have to be justified carefully, lest the broader argument against Purely Evaluative Moral Realism turn out to be circular. One thing they could argue is that, given the special association between morality and action, it is insufficient for moral facts merely have the potential to be normative reasons in the same way other facts do. Trying to account for this, however, is likely to lead us back to special normativity and so to Moral Error Theory. We might be better served trying to explain the intuition prompting these objections, which I will attempt to do in §4.2.

With those issues more or less handled, we have shown that there is at least one example of a substantive ethical theory which is fully compatible with Purely Evaluative Moral Realism, in the form of Scalar Utilitarianism. But a metaethical system that supports only one substantive moral theory is insufficient for our purposes; I am arguing for the general plausibility of Purely Evaluative Moral Realism, so any one substantive system would be too much of a commitment.²⁹⁰ I must show that other paradigms within substantive ethics can be accommodated by Purely Evaluative Moral Realism.

4.1.2 Virtue Ethics

Like consequentialism, Virtue Ethics is such that it can be fairly easily made amenable to Purely Evaluative Moral Realism. But, also like consequentialism,

²⁹⁰This isn’t to say that supporting only one substantive system means a metaethical system is definitely wrong. It could, after all, support the right substantive system. However, unless that substantive system is thoroughly unlike extant examples while also being almost undeniably correct, showing that the metaethical system supports that substantive system is no kind of argument for the correctness of the metaethical system. Naturally, it is doubtful that any substantive system meets all these criteria.

the most widely known version of the view—in this case the Aristotelian mould—has deontic elements that must be excised or reinterpreted for this to be achieved.

Nevertheless, some of the most fundamental claims of Virtue Ethics, which persist from one version to another, can easily be understood in a purely evaluative way. One of these is that moral facts are not directly concerned with our actions but with the character traits that underlie them, virtues and vices. Facts about virtue are, I submit, evaluative facts because what gives a virtue its status as such is surely that it is a *good* character trait.^{291,292} In fact, some thinkers have suggested that deontological concerns and questions of right action are something of a red herring for moral philosophers, and that one of the primary advantages of Virtue Ethics is its ability to explore ethical questions without invoking them.²⁹³ Naturally, this leaves it with no room for special norms and so prime ground for a purely evaluative metaethics.

As I have mentioned, though, there are aspects and versions of this view that do seem to invoke special normativity, and I must show at the very least that these can be done without before I can declare that Purely Evaluative Moral Realism has room for Virtue Ethics. Firstly, there is the possibility that virtues themselves

²⁹¹What exactly makes for a good character trait will depend on the supervenience base for moral properties, and what these are is a substantive question that falls beyond the scope of this project.

²⁹²Since virtues and vices are examples of thick moral concepts—which build in more specific descriptive or naturalistic content than thin concepts like ‘good’ or ‘wrong’—this serves to show that Purely Evaluative Moral Realism generally is well-positioned to accommodate such concepts. Notably, there are some moral concepts that build in the idea of reasons for action, such as being praiseworthy or abhorrent, but the descriptive content of such terms is only in the warranted response—to claim that something is abhorrent says nothing substantive about that thing except that its natural* traits are such that we have reason to abhor it, so these are not thick in the relevant sense. One other potential counterexample is that of being corrupt, which seems to involve violating norms around the use of power, but there is no reason the norms in question should not be institutional norms of the role that bestows that power, rather than the special normativity that Purely Evaluative Moral Realism rejects.

²⁹³See, for example, Anscombe 1958, pp. 8-9 and Brewer 2009, chs. 3, 5.

imply the existence of specially normative reasons. One way that virtues have been characterized is as a disposition to regard a certain specific range of considerations as reasons for action.²⁹⁴ Of course, while this brings normative reasons back into the picture, it doesn't imply that anyone already has the reasons in question as would be the case if they were specially normative reasons. This reading is easily compatible with the view that to have a virtue is to adopt ends of a certain kind, or to have a certain kind of motive, such that the virtuous agent recognises certain considerations as reasons. Therefore, unless there is some kind of specially normative onus on us to be virtuous, this is not a problem for our attempt to make Virtue Ethics cohere with Purely Evaluative Moral Realism.

The question that naturally follows, then, is whether there is any such onus on us. Are we obligated—do we have specially normative reasons—to be virtuous, or to be as virtuous as we can be? Certainly there are views that seem to imply as much. One common interpretation of Aristotle is that his Virtue Ethics is grounded in a strongly teleological view of the world; it is the *purpose* of a human being to do what is uniquely human, which is to live according to reason, which is to live a virtuous (and, if one is lucky, eudaimon) life.²⁹⁵ This notion of inherent purpose can easily be read as special normativity; if such a purpose gives one reasons it will do so in a way that transcends one's particular personal ends.

Even if we read Aristotle in this way, however, this is not necessarily a problem. For the most part, modern Virtue Ethicists do not endorse Aristotle's teleological view of nature, which doesn't fit well with a modern scientific understanding of the world, where teleological elements do not feature.²⁹⁶ Any appeal to teleology is peculiar to his model and therefore not an essential part of the view. If this

²⁹⁴Hursthouse and Pettigrove 2018, §1.1.

²⁹⁵See, for example, Nagel 1972, p. 253; Irwin 1980, pp. 35-7, 48-9; Hester 1991; Kraut 2018, §2; and Lefka 2021, p. 302. Such interpretations draw on the function argument given in Aristotle's *Nicomachean Ethics*, 1097b20-35, 1098a15-20. I will not consider here the merit of this interpretation, or any rival interpretations, but I will note that any reasons to reject this reading are points in my favour here, as they undermine a major argument for special normativity as a feature of Virtue Ethics.

²⁹⁶Hursthouse and Pettigrove (2018, §3) go as far as to call the Aristotelian view 'discredited'.

can be stripped out, we can have a version of the view that is free from special normativity and so compatible with Purely Evaluative Moral Realism.

The same goes for any other way in which it might be argued that there is a specially normative element to Virtue Ethics. The modern successors to Aristotle are eudaimonist Virtue Ethicists, who may well reject his teleological picture of the world while embracing the claims that virtue is an essential part of eudaimonia, and that all humans have a reason to pursue eudaimonia. Given that eudaimonia roughly means something like flourishing, happiness, or wellbeing (though much is always made of the inadequacy of these translations), it's one of the few things anyone could convincingly claim we all have reason to pursue.²⁹⁷ (Though for a Purely Evaluative Moral Realist to grant this claim, they would need to insist that these reasons are non-special. One way this could be so is if they reduce to some feature that belongs to all humans contingently, and is not a necessary feature of agents as such.) These observations, however, fail to ground the conclusion that all possible versions of Virtue Ethics entail specially normative facts. Firstly, not all Virtue Ethicists relate virtue to eudaimonia.²⁹⁸ Secondly, even if we assume that the reasons to pursue eudaimonia would be specially normative on some versions of the Eudaimonist tradition, this would only imply that these versions of Virtue Ethics would be incompatible with Purely Evaluative Realism.

What is essential to Virtue Ethics is the claim that some traits genuinely are good for their possessor to possess, either because they are good in themselves or good for said possessor. Which traits those are and what grounds their status as such vary across accounts, but these are questions of substantive ethics. What matters is that the base view is only committed to evaluative facts of the kind that Purely Evaluative Moral Realism permits.

So compatibility is established. But there is more we must cover before this becomes an answer to the objection from §3.3.3. The compatible versions of Virtue Ethics must also be fit to guide our actions. This has some potential to be a

²⁹⁷If one took facts about which actions would advance one's wellbeing as a general reduction base for normative reasons, as I suggested in §1.8, such reasons would apply universally without being irreducible and so without being a problem for Purely Evaluative Moral Realists.

²⁹⁸Hursthouse and Pettigrove 2018, §2.2-2.4.

stumbling block, since there are some who would claim that no version of this view is fit to guide action.

This criticism is grounded in the practical application of the view; Virtue Ethics, the argument claims, cannot guide our actions because it can never produce a clear answer on how we should act in a given situation. In short, it cannot be used as a decision procedure. This isn't strictly correct—what it fails to provide is a straightforward criterion of rightness, but at this juncture that should not come as a surprise. It can still, however, be used to determine what course of action is prescribed by the balance of moral reasons.

According to Virtue Ethics, there is a fact of the matter about which elements of a situation constitute reasons for a perfectly virtuous agent. Anyone who is concerned about being virtuous themselves then only has to apply their phronesis or practical wisdom to determine what these reasons are, their weights, and how they balance against one another, and they will have their answer as to the virtuous course of action. This process is liable to go wrong due to human imperfection, but it still implies that a sufficiently well-informed judge could decide what to do in this way.

This decision procedure may be hard to use, but problems with the availability of information are equally applicable to utilitarianism. There are so many factors at play in some situations that it is really anyone's guess which action will produce the most happiness.²⁹⁹ Of course, one could be a subjective utilitarian and claim that people ought to perform whichever action has the highest *expected* utility, but even on subjective grounds it can be very difficult to make the requisite calculations.³⁰⁰ So utilitarianism, too, can face an incomplete information problem when trying to determine right action.

We can keep going with this. There are questions about how to formulate the Kantian categorical imperative, how broad or granular we can afford (or are required) to be when framing our maxims, and how to judge when edge cases really

²⁹⁹See Smart 1973, p. 34, Norcross 1990, p. 254, and Greaves 2016, §I. Note that these thinkers do not think this is a death knell for utilitarianism any more than I think the obscurity of virtue is one for virtue ethics.

³⁰⁰Feldman 2006, p. 56.

do amount to lying or promise-breaking or anything else that might be forbidden by a deontological system. Even Divine Command Theory is subject to these ambiguities—if what is right and wrong turns on the will of God, then we will not know how to act until we know which god(s), if any, exist and what precisely they would have us do. So problems of this type seem to arise across substantive ethics—with the exception of those who reject theory entirely, and they run into a different slew of problems—and thus pose no special threat to a purely evaluative form of Virtue Ethics.

Further, given that both Virtue Ethics and Purely Evaluative Moral Realism take the idea of right action to be ill-founded in the first place, the difficulties they face in pinning down a single action to recommend should perhaps come as no surprise. Yet there is, as we have seen, a path to such recommendations. It is obscure, but not more so than some rival views.

In any case, the possibility of making moral decisions on the basis of Virtue Ethics is enough for it to clear the hurdle of being theoretically able to guide action, which in turn means it fits the minimum standard that we have adopted. As is the case for Scalar Utilitarianism, this is because moral facts generate non-special norms with reference to some self-consciously moral motive such as a drive to be virtuous or, more broadly, to be and do good. Proponents of eudaimonist Virtue Ethics could even cite the desire to achieve eudaimonia as such a motive. (This move would be questionable, since it treats morality as a means to the end that is a life well lived and the motive itself thus needn't make direct reference to any moral concepts, but we could also claim that a desire for eudaimonia properly understood is self-consciously moral since virtue is a necessary part thereof.) Again, I will have more to say about such motives and their role in human psychology in §4.2.

4.1.3 Deontology by Way of Honouring Values

Of the three main traditions within substantive ethics, deontology is clearly the one that plays least well with Purely Evaluative Moral Realism. A system of ethics concerned first and foremost with duties could be understood as nothing more than a series of specially normative claims, which would render it directly incompatible with the model I have suggested. But other ways of understanding deontology may

be more amenable to being understood through a purely evaluative framework.

The one of greatest interest to us is Philip Pettit's suggestion (already discussed in §2.2.2) that the fundamental difference between consequentialist and non-consequentialist systems, including deontology, is the way they tell us to respond to value.³⁰¹ Under consequentialism we ought to promote that which is valuable (or rather, the moral status of actions is determined by the degree to which they promote that which is valuable) whereas non-consequentialist systems demand that at least some values be honoured instead. Honouring a value, in this context, means acting out of respect for the value itself and (insofar as this is possible) instantiating it in one's own actions.³⁰²

This paradigm faces some challenges. It struggles to deal with some values that are widely accepted, most conspicuously happiness.³⁰³ Deontologists generally grant that happiness is valuable, but it's unclear what it means to honour happiness and how this would be distinct from promoting it.³⁰⁴ Of course, Pettit's characterization says only that some values are to be honoured, not that none are such that they are to be promoted, so we could plausibly claim that happiness is to be promoted while other values are to be honoured. But we would then need to explain the difference between the values that are to be honoured and the ones that are promoted, and if we are to square this view of value with Purely Evaluative Moral Realism we must do so without appealing to anything like the Fitting-Attitude Analysis where we say they warrant (in some specially normative

³⁰¹See Pettit 1989, 1991.

³⁰²McNaughton and Rawling 1992, p. 835.

³⁰³McNaughton and Rawling 1992, p. 836.

³⁰⁴There is a deeper problem here insofar as the core difference between promoting and honouring values is unclear. Why does committing an injustice to prevent some more serious injustice violate the requirements of honouring justice? It seems like the answer would have to appeal to deontological ideas, which introduces an unpalatable level of circularity into the view, at least insofar as the idea of honouring values is meant to explain deontology. Our interest, however, is not in the overall merits of this or any other version of deontology, but in whether it can be made coherent with Purely Evaluative Moral Realism, so I will not pursue this objection any further.

sense) such a response. So this may be a dead end for our purposes.

Alternately, suppose we had an axiology or theory of value consisting entirely of the kind of values that we can honour, such as honesty and loyalty, or the values of life and human dignity.³⁰⁵ Suppose even that this axiology allows us to account for those difficult values like happiness in a framework based entirely around honouring values. (This might strike some as counterintuitive, but there is no substantive system of ethics that will not provoke that response from someone somewhere.)

In such a case, an agent with a self-consciously moral motive that made reference to honouring values, rather than or alongside references to doing or being good, would then be subject to a set of non-special norms that mirrored the content of a system of deontological ethics based on the same axiology. In other words, Purely Evaluative Moral Realism is capable of accommodating an ethical system with almost entirely the same substantive content as at least some forms of traditional deontology.

This will be insufficient to satisfy some deontologists. Being non-special, the norms in question are conditional in a way that doesn't gel with conventional understandings of duty or obligation. Further, no compelling reason has been given for agents to have a motive that entails honouring values rather than any other self-consciously moral motive or none at all. Purely Evaluative Moral Realism doesn't admit of specially normative reasons, so deontologists cannot say that we are obligated to have any such motive, and so morality loses the privileged position as something to which we cannot help but be subject. This, however, is just another instance of a broader objection to the purely evaluative view, which I already addressed in §3.3.3.

While the picture given above is full of provisos and is especially open to some of the objections to Purely Evaluative Moral Realism, it nevertheless presents us

³⁰⁵One could even go so far as to use some existing deontological system as the basis for one's axiology. One might say, for example, that actions are morally valuable if and only if they conform to the Kantian categorical imperative. Such a system will strike many as needlessly roundabout—and I am inclined to agree with them—but it is not off the cards entirely, and the remarks that follow all apply equally well to such a system.

with a case where all of the prescriptions of some deontological ethical system are true for some agents (and potentially all agents with a self-consciously moral motive, depending on how one's own values provide reasons; see §4.2.) So we can give an account of deontology through the lens of Purely Evaluative Moral Realism, albeit one that involves more (and perhaps more egregious) structural changes than either of our previous attempts.

There is no question as to whether such a system could guide action; once formulated it is a series of (qualified) prescriptions and so directly suited to the task.

4.1.4 Divine Command Theory

Another form of substantive ethics that turns out to be surprisingly amenable to Purely Evaluative Moral Realism is Divine Command Theory. Like the version of deontology discussed in §4.1.3, it requires some additional commitments, but these ones should not seem especially egregious.

The possibility that Divine Command Theory might be made to cohere with Purely Evaluative Moral Realism is surprising, because what Divine Command Theory centrally claims is that the demands of morality are coextensive with that which is commanded or willed by God. So moral facts are facts about imperatives, in that they are facts about what has been demanded of us. As discussed in §1.8, this kind of structure is strongly connected to the special normativity of moral facts, so accounting for such a system in purely evaluative terms does present a challenge.

Whether we can do so will turn on what it means for God to will or command something. The account of this which is of interest for our purposes began life as an attempt to dissolve the Euthyphro dilemma. This is the question of whether God wills that we do what is good because it is good, or whether what God wills is good because God wills it (both of which have rather unpalatable implications for Divine Command Theorists). The response in question begins with a claim about the nature of ethical goodness, namely that it consists of resemblance to God.³⁰⁶ Whatever features are good-making in any entity are so because they are

³⁰⁶Alston 1990, pp. 319-20.

features of God. This includes actions (although these might be more appropriately be called compatible with God's nature rather than resembling God) and so the claims Divine Command Theory makes about what we ought to do are constrained, perhaps even determined, by claims about goodness insofar as the kind of things God can will or command are likewise constrained.³⁰⁷

Such an account makes Divine Command Theory a view explicitly about what we ought to do, which is grounded in further commitments to a specific theistic theory of value. This still looks incompatible with the purely evaluative view.

However, the presumption that a command is equivalent to a special norm is just that. We can easily make sense of the reasons yielded by commands, even commands from God, in a non-special way. A command, by its nature, comes from a person in a position of authority, like one's superior within an organisation or an agent of the government. The commander has some kind of power over the recipient of the command, and there is some kind of incentive to comply in place. This may only be implicit—a soldier does not need to be told with every order that they will be court-martialled or demoted if they disobey—and it needn't be the only motivation behind their compliance—a person may comply out of genuine respect for the authority of the one issuing the command—but this structure of extrinsic punishment and reward is always present.³⁰⁸

This being the case, it seems that once someone has been issued a command, they have normative reasons to comply with it. These reasons are non-specially normative, since they are so dependent on the personal ends of those involved

³⁰⁷Adams 1999. Quinn 2007, p. 78.

³⁰⁸This does not extend to all imperatives in the sense of any speech act that uses a verb in its imperative form. This can also include sentences like 'Please pass the butter,' or 'Duck!' shouted upon seeing a softball sailing toward the listener's head. The former is a request (asking, rather than telling) and the latter is more like advice delivered quite urgently than a command (though there is a consequence to failing to comply, it is not extrinsically imposed by the speaker, but an intrinsic consequence of failing to dodge the softball). Likewise there may be social reasons to acquiesce to a request, albeit without the asymmetrical power relationship implicit in a command. Such cases are of little concern to us, though; as far as I am aware, there is no religion which claims that God politely requests that we be good.

(or another reduction base). Insofar as they are commands, we should expect commands from God to be isomorphic with ordinary commands, with the proviso that God, as typically understood by theists, is the ultimate authority. God's commands can thus be understood as categorical in two ways, though only one of these matters for our purposes.

Firstly, they are categorical in the sense that they are directed at all agents. This is the sense that is not relevant for our purposes—I can give an order for every agent in the universe to jump and it would be a 'categorical imperative' in this way, but none of those agent would thereby come to have normative reason to jump because there is no authority behind my orders. Secondly, they are categorical in the sense that they actually provide normative reasons to all agents. As an omnipotent being, God is always in a position of power over any and every agent such as to give God authority (and so the power to provide reasons) in the way outlined above. Hence his authority is unconditional, which we might take as a sign that the normativity of God's will is special normativity. However, this conclusion is not inevitable.

Remember that commands typically occur against a background of extrinsic incentives which are the source of their normative power. Typically theists posit just such a system of reward and punishment through ideas about the afterlife. If the consequence of acting immorally is eternal suffering, then it seems like any agent will have non-specially normative reasons to perform good actions rather than bad ones. This does not even require us to posit a self-consciously moral motive; prudential concerns alone are sufficient for this kind of normativity.³⁰⁹

³⁰⁹We might be inclined to question whether a form of normativity grounded in prudence is really morality, though this problem is one shared with lay versions of Divine Command Theory that hinge their authority on God's power and inclination to reward and punish. But, the fact that we can ground these norms in prudential concerns doesn't rule out a more self-consciously moral version so much as imply that both the prudentially and morally motivated should come to the same practical conclusions. Consider also that although Divine Command Theorists often believe in such an afterlife, it is not essential to their view; there could be a Divine Command Theorist who believes there is no afterlife but that we should act in accordance with God's will simply because that is the moral thing to do.

Some theists will want to insist that the commands of God carry more weight than this, but such commitments are an extravagance when it is possible to account for the reasons He gives us in terms of non-specially normative reasons.

So with just a couple of additional commitments—one of which solves a long-standing problem for Divine Command Theorists, and the other of which is already widespread among them—we have an account of how the facts about what is good (here meaning which things resemble God or are possible for a being insofar as that being resembles God) are a guide to action, because they determine what commandments God can and will make. This account differs from those given for Scalar Utilitarianism and deontology in that it does not require that we have a self-consciously moral motive (though it does not rule it out either) but it still comfortably answers the objection from §3.3.3.

In this section, I have argued that consequentialism, deontology, Virtue Ethics, and Divine Command Theory each have at least one incarnation that is compatible with Purely Evaluative Moral Realism and lives up to the demand that morality give us a guide to action. Each such theory does indeed give us a way to respond to the objection in §3.3.3, but they are fairly limited ones calling for highly specific sets of commitments. In §4.2, I will lay out a more general account of how the moral facts permitted by the purely evaluative view can guide action and how, even without special norms, there can be normative moral facts which are true for the vast majority of agents. What follows can be considered an extension of Purely Evaluative Moral Realism beyond the basic claims laid out in Chapter 3, such that one could reject these further claims while still being a Purely Evaluative Moral Realist of some stripe. I won't claim that it is the only such extension possible, only that it is the most obvious and intuitive to me. It allows us to account for the action-guiding nature of morality and to say that certain moral normative claims are true in a way that may placate some critics, without committing ourselves to anything so ontologically extravagant as traditional realism requires of us, which I believe are significant virtues. This picture will invoke (and describe in greater depth) a common element among many of the views I have just discussed: self-consciously moral motives.

4.2 Realism about Non-Special Moral Norms

In this section, I will develop a general account of how even purely evaluative morality can still have normative implications. I will do this by way of an analogy with epistemology.

In Chapter 2, we introduced Normativism, the view that moral facts (or some kinds of moral facts) entail specially normative facts. The various forms of Normativism that were discussed are all, strictly speaking, Normativism about morality. We could generate Normativism about other putatively reason-giving domains just by swapping the appropriate descriptors into the definitions given there. For example, some people might subscribe to a form of epistemic Normativism, where some or all epistemic facts are taken to imply the existence of specially normative facts. Moral Error Theorists who base their position on the queerness argument from special normativity are best able to do so when they are Normativists about morality and nothing else. If one has already accepted a Normativist view of some other domain of fact, then admitting specially normative moral facts makes no fundamental additions to one's ontology and so the debunking account loses its advantage over moral realism. This is why companions-in-guilt are such a threat to the Argument from Queerness.

As it happens, there are a few domains beyond the moral about which many thinkers would be willing to endorse some form of Normativism. Examples include logic, epistemology, and semantics. Anyone who means to motivate Moral Error Theory by way of the Argument from the Queerness of Special Normativity must respond to the presence of any putative companions in guilt. Olson is sensitive to this, and suggests that we could simply accept that such norms are equally problematic to moral ones (though this leaves them as fodder for a companions-in-guilt argument as discussed in §1.4 and §1.8).³¹⁰ However, he also offers models that can explain away these norms' apparent special status in a way much more conducive to the overall Moral Error Theoretical argument.

Olson points out that wherever we try to determine whether some proposition p is true, it is because we have some interest in whether p is true, and we therefore

³¹⁰Olson 2014, pp. 138, 155.

implicitly have non-special reasons to believe whatever results our investigation turns up.³¹¹ Elsewhere, he suggests that the norms of deductive logic are reducible to a standard of correct reasoning, which likewise makes their normativity non-special (assuming correctness itself is not a specially normative concept).³¹² These are good enough as explanations of why these domains might seem specially normative without actually being so, but Olson's approach has a number of superfluous elements because he addresses each of these domains separately. I think we can take things further, and explain a wide class of these norms in a single move.

I submit that human beings have a general interest in truth. By this, I don't mean that we seek to maximize the number of truths known, or that we seek out truth as an end in itself; it is not as though we see people collecting any and every trivial fact within reach purely for the sake of knowing true things (at the very least, such behaviour is far from common).³¹³ Rather, I mean that we are concerned with whether the beliefs we form are true. We generally want to form true beliefs rather than false ones. The propositional content of our beliefs may not always be such that we want it to obtain, but insofar as we form beliefs about any state of affairs—even bad states—we want to form accurate beliefs. I believe this trait to be universal, or close to it, across the human race.³¹⁴

Such a trait as I am suggesting could be easily explained with reference to evolution by natural selection. This is the same kind of explanation invoked by debunking accounts of moral phenomena, so Moral Error Theorists who have come this far should have no objection to it. Generally speaking, an interest in or desire in believing p when p is the case and believing not- p when p is not the case, or a disposition to form such an interest or desire, would be adaptive—and thus favoured by natural selection—simply because of the widespread usefulness of true

³¹¹Olson 2014, p. 161.

³¹²Olson 2014, p. 137.

³¹³I owe this observation (and so my awareness of the need to draw the distinction I outline below) to Johan Gustafsson.

³¹⁴This is not to say that it must also be overriding. It's perfectly consistent for a person to desire that the beliefs they form be accurate, and also to wish they had remained deluded about some particular truth for non-epistemic reasons, such as that truth being particularly upsetting.

beliefs (and the corresponding inutility of false ones). The instrumental value of an accurate picture of the world is hard to understate because for practically any given end, true beliefs about the relevant matters of fact will be conducive to its attainment. This includes the ends of survival and reproduction, which are the only ones that need concern us in the process of natural selection, though an interest in truth will usually also serve the overall good of an individual possessor insofar as true beliefs will serve whatever other ends they might have.

Having a motive like this would make an agent subject to a wide array of non-special norms. Importantly, there are objective facts of the matter about which methods of belief formation are most truth-conducive. These are epistemic facts, and they are directly relevant to those who wish for the beliefs they form to be true ones. They can lead us to a set of non-special norms that are relativized to what we might call the characteristic concern of epistemology, which is knowledge (as opposed to false belief rather than as opposed to ignorance of any stripe).³¹⁵ Likewise the rules of deductive logic, which are explicitly meant to preserve truth, must apply to anyone who hopes to acquire knowledge and avoid delusion.

If, as I suggest, all or very nearly all humans have this concern for truth, then some associated set of epistemic norms is true for all or very nearly all humans. It is therefore unsurprising, and perhaps even very reasonable, that when we state epistemic norms, we tend to do so without qualifying them. We say things like ‘You shouldn’t believe everything you hear,’ and ‘One ought not believe both p and not- p ,’ without appending something like ‘... if you want to form true beliefs and avoid forming false beliefs,’ because of an implicit (and mostly safe) assumption that all other agents do, as a matter of fact, want this.

Brief reflection will reveal many domains where we speak in these terms. We say things like ‘You ought to eat more fruits and vegetables,’ and ‘You shouldn’t buy lottery tickets,’ and ‘Don’t touch the electric fence,’ without further comment. Of course, we can justify these prescriptions if it is demanded of us; we might for

³¹⁵There is, of course, more to knowledge than true belief, which is what we have been discussing. But, having knowledge implies that one has a true belief, at the least, so whatever epistemic facts are relevant to true belief formation are relevant to knowledge acquisition and vice versa.

each of these respectively say ‘They’re good for you,’ and ‘They’re a waste of money,’ and ‘It’ll hurt like the dickens.’ So each of these is a non-special norm whose reduction base includes some end or desire. But we don’t phrase them as hypothetical imperatives. Rather, since they are relativized to very general goods such as health, wealth, and the avoidance of pain, we simply assume that the addressee has the relevant motives.³¹⁶ Usually, we are correct to do so, and thus we very frequently make normative claims that are, at least *prima facie*, correct.³¹⁷

What is important is that, in explaining these norms, we invoke something further that we assume to be an end for those to whom they are meant to apply, but we don’t need to appeal to anything like special normativity. I see no reason that we shouldn’t think that a similar model applies in the case of epistemic norms. Certainly, this would be more parsimonious than ascribing them a special kind of normativity for the reasons discussed in §1.8. Some might resist the model I propose as it applies to epistemology in particular by claiming that such a motive itself is an extravagant commitment that violates the same principle of parsimony that motivates many of the arguments I have endorsed over the course of this project. Believing some proposition entails—at minimum—holding it to be true, so the idea of an agent believing in anything but a consciously truth-tracking way, which aims to hold only true things to be true, is incoherent. There are, though, some beliefs where (for at least some believers) the literal truth of their content is not paramount, such as religious belief, but those who wish to reject my model may dispute that this is really the same phenomenon, and the discussion would end at an impasse of opposing intuitions as to what counts as a belief, so I will not

³¹⁶Finlay (2004, pp. 217-8) suggests a similar approach to judgements of goodness *simpliciter*. While the metaphysics of value differs from what I have advanced—at the very least, I don’t want to assume that moral goodness is related to a single characteristic end—the conversational pragmatics are much the same. He also grants that unqualified value prescriptions could instead be grounded in the needs and expectations of society generally (p. 222) and the same is plausible as a reduction base for normative claims of the kind I have been concerned with here.

³¹⁷I specify that these claims are *prima facie* correct here only because I take it that what a person actually ought to do is what they have most reason to do and so a very small subset of those things that they have some (or sufficient) reason to do.

hinge my response on these. Instead I will point out that even if the idea of forming beliefs without the intention that they track truth is somehow contradictory, this need not imply anything specially normative, or rule out the existence of a motive to form true beliefs rather than false ones. Claims about the nature of belief, like the nature of anything else, are descriptions rather than prescriptions, and I think that desires for one's physical and mental capacities to operate normally are not particularly surprising or extravagant, since pursuit of all of one's other ends is contingent on these.

To summarize, I have suggested a model whereby the normativity of epistemic facts is non-special, but contingent on an end that is universal (or close to it) across the human race.³¹⁸ We attribute them special normativity because this explains why epistemic norms seem to apply to everyone, but we can explain the same observations more parsimoniously by realising that this is just an extreme case of our natural tendency to assume that others have certain ends. There are, of course, other views of how we are to account for the normativity of epistemic facts, but what matters for our purposes is that we are not forced into accepting epistemic Normativism.³¹⁹

This has been a lengthy preamble, but we have finally made our way to the point. While we are firmly in the realm of the possible rather than the certain, we

³¹⁸In fact, the interest in truth may be even more universal than that, since it is hard to imagine a propensity to form beliefs without concern for their truth would survive a process of natural selection, so it might be that any agent with an evolutionary history also has such an interest. Nevertheless, said interest would still only be one among many, and a substantive fact about the agent, so the normativity it underlaid would still be non-special.

³¹⁹For example constitutivist views (see Wedgwood 2002 and Shah and Velleman 2005) start from the aforementioned notion that beliefs are meant to track truth by nature, and submit that in believing without aiming at truth an agent would be believing in a fundamentally incorrect way. In other words, it would be irrational to believe anything without at the same time aiming for truth in said belief. Whether correctness or rationality are specially normative notions is a separate question I won't engage with here. Others like Kelly (2003) will directly deny the instrumentalist picture I have sketched above and insist that epistemic norms (perhaps among others) are precisely the kind the queerness argument from special normativity aims to problematize.

have a way to account for epistemic norms without invoking special normativity. The facts in question are still concerned with reasons, but the reasons in question are very much explicable and so raise none of the explanatory burdens that are so central to the Argument from Queerness. Given that this possibility is open to us in one case where we might otherwise have relied on a Normativist view of the domain, the natural next step is to attempt to generalize this to other other domains where we have been tempted to invoke special normativity. So I suggest we extend this model to the moral domain.

For this analogy to hold, it must be the case that humans generally have a self-consciously moral motive. When I say that a motive is self-consciously moral, I mean that, if expressed in natural language, it would make direct reference to moral concepts and need to be expressed using moral language. I have invoked such a motive, a desire to do or be good *de dicto*, in most of my accounts of how traditional substantive views can be brought in line with a purely evaluative metaethics. To me, it seems very plausible that we have such a motive. Just as we ask whether *p* because we care whether *p*, the fact that we engage in ethical debates and moral philosophy, and that people face moral dilemmas suggests that there is a widespread and abiding interest in the moral domain. As discussed when we outlined our debunking account of morality, this mostly manifests as thinking in terms of categorical imperatives (be that specific prescriptions or a more general desire to do ‘the right thing’) whether or not this is reflective of the actual nature of moral facts.

If we do indeed have such a motive, and whether or not it is dressed up in the language of special normativity, then it follows that we should be able to construct non-special norms that apply to all or most agents based in this motive. This picture is complicated slightly by a degree of flexibility in the terms I have invoked.

When I talk about a motive, I mean anything internal to a person’s psychology that can give them ends and move them to action. In other words, I mean a desire in the broadest philosophical sense, and whether this has a single unified meaning or is a polysemous term that encompasses many conative mental states is

a matter of some controversy.³²⁰ I won't try to solve this puzzle here, but I should acknowledge that remaining neutral means I must be prepared for self-consciously moral motives to take many forms.

I have already specified that a self-consciously moral motive is one that would need to be expressed in moral language. Moral language includes both evaluative and normative terms. As I argued in Chapter 2, there is no necessary entailment between these categories, and as I discussed in Chapter 3 the Purely Evaluative Moral Realist must regard the latter as empty insofar as it refers to the special normativity often attributed to morality.

We can profitably understand self-consciously moral motives as drives to respond to perceived moral value in particular ways. But, as discussed in §4.1.3, there are multiple ways to respond to value, so it's plausible that there are multiple such motives at play within and across agents. We needn't be closed to the possibility that valuing some particular thing entails, in one way or another, desiring certain states of affairs or treating them as ends (and actually doing so, rather than it being the case than one ought to do so in any specially normative sense). In such a case, that desire or treatment is the self-consciously moral motive when that which the agent values is moral goodness. It would follow from this that there is some single action that a morally motivated agent has a genuine reducible normative reason to perform in any situation, though not necessarily one that they ought to perform since they may have other motives that overrule it.

It may be objected that we are in no position to assume this, or even that we have some good reason to think it is false given that, as we have already acknowledged, people do respond to the things they value in more than one way in the actual world. But, I have only said that we needn't rule out the possibility that valuing something entails certain desires or ends. We can be equally open to the alternative, which is that there is a plurality of self-consciously moral motives,

³²⁰A good example of the kind of polysemy that is of interest here is 'jade', which historically referred to a range of minerals with superficially similar properties, but failed to pick out a natural kind in terms of chemical structure or composition. Motives likewise have enough in common (neurophysiologically or phenomenologically, for example) that we can usefully refer to them together whether or not there is a single natural kind to which the term refers.

none of which holds a unique position above the others (whether this be within agents, across them, or both).³²¹ This includes the perverse agent discussed briefly in §1.6 who seeks to do what is bad or wrong, though I think as a contingent fact of human psychology (born out of the same evolutionary history already discussed) that this is not the kind held by most agents. In this pluralist situation, for the same reason we can suppose people have self-consciously moral motives in the first place, it seems reasonable to conclude that such motives still fit within a fairly narrow spectrum. Hence, there would be a lot of overlap in the norms generated from these motives in much the same way that different substantive theories of ethics often overlap in their recommendations and assessments while each having their own rationales that lead to divergent logical extremes and different approaches to edge cases. This overlap allows for many unqualified moral prescriptions to still be true for most listeners and so fit into the same paradigm as epistemic norms.

So we have a general explanation of how non-special moral norms apply to most human beings.³²² Like any domain of fact, moral facts serve as normative reasons for any agent with certain motives to which they are relevant, and we have good reason to believe that all or most humans have some motive (be it ‘the’ self-consciously moral motive or one or more out of a possible range) that turns moral facts into reasons for them. As such, it is clear that moral facts can guide our actions (but not that they must do so, as that would be typical of special normativity). Whether there is one ‘right’ action from the moral perspective, as

³²¹This is the picture suggested by Svavarsdóttir (1999, p. 197), who takes it that an individual’s self-consciously moral motive has its content settled by the moral concepts that the individual employs. I certainly agree with her that the characterization given by Michael Smith, a will to ‘do the right thing’ is unlikely to be universal, which is just as well for Purely Evaluative Moral Realism, which must take ‘the right thing’ to be either an empty descriptor or (if that means whatsoever one *ought* to do) entirely contingent on agents’ other motives, including any moral ones.

³²²The term ‘non-special moral norms’ is a gloss I will use for readability. Strictly these are non-special norms grounded, at least in part, in evaluative moral facts. Per the core claims of Purely Evaluative Moral Realism, they are not moral facts in themselves but facts about practical reason which can be explained only with reference to moral facts.

there appears to be in the case of epistemology, depends on whether there is a single desire or end that one must have as a consequence of one's own acts of valuing (rather than value in the objective sense, which the Purely Evaluative Moral Realist takes to be normatively inert in and of itself).³²³ Nevertheless, it remains the case that there is a role for moral facts to play in structuring at least some of our actions.

This is meant to answer to the objection from §3.3.3, that purely evaluative moral facts cannot guide our actions. It only works, however, insofar as the view I have outlined above is acceptable. Earlier, I referred to this view as an extension of Purely Evaluative Moral Realism and this is accurate inasmuch as it is really the addition of a few further commitments. The first addition is that non-special norms are unproblematic. I have been treating this as an assumption until now, and I take it for granted that most readers will be happy to go along with this, even if they need to give a highly deflationary account of what non-specially normative claims amount to. The second new commitment is that moral facts can be the grounds of non-special norms, and this too is fairly innocent. Consider a journalist who wants to report on a story of modern heroism; they have a reason—rooted in their own desires—to respond to moral facts in a particular way.³²⁴ Even normativist realists, who hold that moral facts always have specially normative implications should be willing to grant this much—with the proviso that the journalist has these reasons in addition to those that normativists take moral facts to intrinsically entail.³²⁵

³²³Where I speak of acts of valuing, I could just as well speak about value judgements. Neither is strictly accurate, since I do not see them as discrete or conscious actions so much as background facts about an agent's psychology, but what is important is that the fact that some agent X values a thing can ground their reasons in relation to that thing, and is distinct from any facts about the objective value of that thing.

³²⁴I take it that heroism is morally relevant since it is usually understood as a form of supererogation (see Urmson 1958). This can be conceptually true even if it turns out that there are no normative facts, and Purely Evaluative Moral Realists might account for heroism with goodness that exceeds reasonable expectations.

³²⁵Sturgeon (1985) might also suggest that moral facts can be relevant to scientific or historical purposes. Under the view he advances, we cannot really explain the end of slavery without

The last and most potentially contentious addition is the idea that most people have a self-consciously moral motive (or a number of such motives) that underlies a set of non-special moral norms that apply to those people. I will discuss the plausibility of this final claim more in §4.3.2. Taken together these amount to a realism about a certain class of non-special norms that are dependent (though not solely dependent) on moral facts. This view can be held in conjunction with Purely Evaluative Moral Realism, and doing so shows how moral facts can be action-guiding on the purely evaluative view.

I have argued that this model is both plausible and parsimonious, but if it could be undermined, so would its power to answer the objection in question. As such, the final section of this chapter will be devoted to anticipating and responding to some possible objections to this extended picture.

4.3 Objections to Realism about Non-Special Moral Norms

I have argued that evaluative moral facts can play a role in guiding our actions and that the truth of Purely Evaluative Moral Realism (which only posits moral facts of this kind), therefore, would not prevent morality from being a guide to action. To do so, I posited a model whereby morality generates non-special norms that apply to all or most agents because of a commonly held self-consciously moral motive. This model was based on an analogy with epistemic norms.³²⁶ In this section, I will consider some objections to this model. These broadly fall into two

appealing to the moral fact that slavery is bad. Like my journalist, a historian might have strictly non-moral reasons to acknowledge and appeal (or respond in some other way) to moral facts. The point I am making above, however, is one that is meant to be agreeable even to someone who does not share the Purely Evaluative Realist's commitments, much less Sturgeon's, hence my use of the journalist example.

³²⁶Here as elsewhere a norm is a *pro-tanto* reason for action. What an agent actually ought to do is a function of all the norms that apply to them and we need not assume that moral motives override other reason-giving motives.

camp; problems with the analogy between the epistemic and moral domains (in the form of disanalogies or counterintuitive implications about either domain), and objections to the idea of a self-consciously moral motive.

4.3.1 Problems with the Analogy between the Moral and Epistemic Domains

Above, I described models of how both moral facts and epistemic facts have the power to guide action. According to the view I described, insofar as the latter is plausible, the former is as well, because they have analogous structures. But there are places where the parallels between moral facts and epistemic facts break down. The question we need to consider is this: do any of these differences serve to undermine the conclusions to which the analogy is meant to lead us? That is, do any of them imply that purely evaluative moral facts are not, as I have claimed, fit to guide action?

One disanalogy that I have already devoted some brief discussion to is between the motives invoked by the two accounts. I made a fairly concrete suggestion for what we might call the self-consciously epistemic motive; we have a general interest in truth insofar as that when we form beliefs, we want to form true ones.³²⁷ Not so when it comes to the self-consciously moral motive. Rather, I have granted that there could be many such motives. For reasons already discussed, I don't think that this kind of pluralism strictly rules out moral facts' ability to guide action, but it would complicate the picture considerably.

For one thing, it seems like it may eliminate moral disagreement in the same way that non-cognitivism does. Non-cognitivism entails that, in making a moral claim, we aren't trying to state a fact or express a belief. Rather, we are elliptically communicating something else like an emotional response or a prescription. Because these statements do not express propositions, they are not truth-apt and

³²⁷Here, 'interest' is ambiguous between true-rather-than-false beliefs being beneficial to us and our having some actual motive to pursue such beliefs. In fact, I think it is fair to attribute both of these to humans in general, though the latter sense is more directly relevant to the model of epistemic normativity I have described.

cannot contradict one another. Opponents of Purely Evaluative Moral Realism might try to make a similar argument; when two agents A and B make prescriptive moral claims, these are relativized to their self-consciously moral motives and so the two will talk past one another. This wouldn't undermine moral facts' power to guide actions under Purely Evaluative Moral Realism, but it would rob this particular version of the view of one of the features that is both typical and advantageous of moral realism in general. It would also run counter to our shared commitment with Moral Error Theorists that the meaning of moral terms is determined by the intentions of ordinary speakers when they use them; there seems to be real and vehement disagreement in discussion of moral reasons for action.

This argument, however, fails due to imprecision; if both speakers are talking about what A ought or ought not to do (in the non-specially normative sense), then it is not the case that A's claims are relativized to A's self-consciously moral motive and B's claims are relativized to B's self-consciously moral motive. A's self-consciously moral motive makes it the case that A ought to act in such-and-such a way—this is a matter of fact by the same token that A's other motives make it the case that they ought to act in certain ways—and since both A's and B's claims are about what A ought to do, the presence and nature of this motive is what determines the truth or falsity of said claims.³²⁸ B claims that A should act in certain ways not because B has a self-consciously moral motive but because B assumes A has a similar self-consciously moral motive (and is *ceteris paribus* justified in doing so).³²⁹

Thus, if A says 'I should ϕ ,' and B says to them in reply 'No, you shouldn't ϕ ,' then they genuinely disagree. Likewise, if A makes a much more general claim like 'People should ϕ ' or ' ϕ -ing is wrong' then they are making a claim about the moral reasons of people generally based on assumptions about the self-consciously moral

³²⁸What these motives actually provide are normative reasons, but all other things being equal one ought to do anything one has such a reason to do, and I have stuck to deontic language here because these are the terms in which moral debates are usually couched.

³²⁹Finlay (2014, pp. 181-5) observes that there may also be pragmatic reasons to proceed under this assumption even when one has some reason not to make it, such as the rhetorical force of unqualified normative claims.

motives of people generally. If there are a plurality of such motives, fewer such claims will turn out to be true, but it remains the fact that anyone who contradicts A meaningfully disagrees with A.³³⁰

These remarks may seem odd; Purely Evaluative Moral Realism is effectively an error theory about normative moral claims and beliefs, and yet I am saying that even in light of it some claims about our moral reasons, or what we morally ought to do, can and will turn out to be true. But Purely Evaluative Moral Realism is only an error theory about special normativity, and entails nothing with regard to non-special norms. Claims like ‘People should ϕ ’ concern our overall normative reasons. And, if we take it that their lack of qualification means they are unlimited in scope and so categorical, then those reasons will seem to be specially normative. Those claims would then turn out to be strictly false. However, per my remarks in §4.2, I think this is the wrong way to take them. I submit that ‘People should ϕ ’ is isomorphic with ‘People should brush their teeth’, in that both are in fact spoken against a background of assumptions about certain common ends of all or most people. (It is not simply the case that people should brush their teeth; they have most reason to do so because they are invested in their general health, their freedom from pain, and their ability to easily eat hot and cold foods, and the same relationship holds between moral reasons that apply to most people and the self-consciously moral motive most people possess.) So ‘People should ϕ , is as true as ‘People should brush their teeth’. This is to say that it is true enough for most contexts, because they are non-special norms backed up by implicit assumptions about the ends of people in general. And the fact that they are carelessly given without qualification should not be a mark against them.³³¹

³³⁰Truth conditions for most statements vary with context, especially for statements without qualification, so claims about which moral norms apply to people in general needn’t be false if it turns out that there is some small number of actual amoralists in the world. Just so long as there is enough of a majority whose self-consciously moral motives lead to more or less the same recommendations, we can make statements of this kind and speak truly.

³³¹This would not prevent the same remark made under circumstances that called for more stringency, or a claim that was explicitly meant to apply to all people, from being false at the same time.

This model handily answers the objection above, but also highlights another disanalogy between the motives invoked by my models of epistemic and moral normativity. This is that there is a lot less controversy over what it takes for a thing to be true—at least in lay circles—than over what it takes to be good.³³²

This too contributes to making the structure of the self-consciously moral motive less clear. We can figure out the self-consciously epistemic motive from fairly consistent patterns of human behaviour with regard to truth. When it comes to ethics, however, things aren't so clear. If we try to determine the self-consciously moral motive by examining patterns of human behaviour with regard to moral goodness or value, it quickly becomes clear that people not only have very divergent ideas about which things are good, but also about how to respond to value. At the very least, we have the paradigms of respecting and promoting values (see §2.2.2 and §4.1.3), but these are far from the only possible answers. Whether or not valuing (as discussed in §4.2) generates a reason to respond in some single particular way, it's clear that in actuality people respond in a wide variety of ways. I won't deny that all of this complicates the picture, but I'm not certain it undermines my analogy. What matters is that just as people typically value truth, they also typically value goodness or moral values in the most general sense. This act of valuing entails one motive or another, which in turn entails some non-special norm or norms that apply to humans in general. (As in §2.2.3, I hold acts of valuing to be an emotional phenomenon distinct from a thing's actually having value *simpliciter* or being believed to do so.) I have been deliberately noncommittal on what the self-consciously moral motive actually looks like because humans' concern with goodness seems to move humans to a much wider variety of actions than their concern with truth. This might seem ad hoc to some, but the connection between the fact of the matter over what is good and the norms that apply to agents is what matters most, and it remains in place despite this disanalogy.

The other way in which this analogy threatens to fall apart is also based in a

³³²It would be naïve to suppose that the question of what truth is or when a proposition is true is settled, but I take it that philosophers and non-philosophers alike have a level of intuitive understanding of truth that meets far fewer challenges from others than any equivalent intuitions about moral qualities.

difference between truth and moral goodness. A self-conscious motive in relation each is attributed to the whole or the vast majority of humanity on the grounds that having such a motive is adaptive. In the epistemic case, our capacity and inclination to track true facts is adaptive because the facts in question are true and we tend to act in more adaptive ways when we believe them. If the analogy holds, our ability to ascertain moral facts and our self-consciously moral motives are adaptive because the facts in question are moral, which is to say that they are facts about what is morally good or bad in actions, and believing them can likewise lead us to act more adaptively.³³³ The problem, according to this line of objection, is that only the former case requires its object to be actual facts; if it is the morality of moral facts that is doing the work then all we really need to posit are moral beliefs, just as we did in our debunking account of moral phenomena.

The crux of this objection is that the explanation of why the self-consciously moral motive is adaptive doesn't necessitate the existence of moral facts as actual facts, whereas in the epistemic case the explanation works only if the facts in question really are facts, which is to say true. This would be a problem, if the point of the analogy were to demonstrate the existence of moral facts, but this is not what the analogy is meant to achieve. The account I have given is not meant as an argument for the existence of moral facts any more than Olson's account of epistemic normativity is meant as an argument for the existence of epistemic facts. Rather, each assumes that a given domain of fact exists, and posits a certain attitude toward that which the facts within that domain are facts about (truth in the epistemic case and goodness in the moral case). By positing such an attitude, each explains how those facts can be normatively significant. In other words, for the purposes of this analogy we are assuming that Purely Evaluative Moral

³³³How we come to have the particular moral beliefs that we do is a question which (we have assumed) can be answered by a debunking account. Wiggins 1990, pp. 68-72 suggests that moral beliefs can be explained by the fact that 'there is nothing else to think' without abandoning 'the point of view that shall be common between one person and another'. I grant that this point of view is strongly relevant to morality, but a Purely Evaluative Moral Realist can simply insist that we are not bound through special normativity to take account of this point of view, and that an inclination to do so is the kind of evolved trait that I have suggested above.

Realism is correct so that we can show how, if it were true, moral facts could still be action-guiding.

The above observation—that, if the analogy holds, it is the morality of moral facts that makes our sensitivity to them adaptive—has one other implication which doesn't amount to a disanalogy per se, but which is counterintuitive enough to fuel another potential objection to the conjunction of Purely Evaluative Moral Realism and realism about non-special moral norms. Such an objection would be another version of the problem discussed in §3.3.1 that the view may be too restrictive over the content of substantive moral theories. In this case, however, the issue is not that it might rule out deontological theories. Instead, according to this objection, commitment to the usefulness of moral facts from an evolutionary standpoint entails that the actions most favoured by the moral perspective must always be adaptive ones. This should trouble us because it seems like there are plenty of laudable actions that would actively undermine the agent's reproductive fitness; sacrificing one's own life to protect others, for example, is often seen as very noble.³³⁴ This objection has it that the view I have suggested implies that these actions could not really be good, which is a problem in itself. It is a deeply counterintuitive commitment, and not one that I intend to make.

But just as this problem arises from the analogy with epistemology, it shall be put to bed by it too. This objection rests on a mistaken assumption that if the self-consciously moral motive is adaptive, it must be because it always and only recommends courses of action which increase the agent's reproductive fitness. This is not at all how natural selection works. The actual process is less discriminating than this; all it demands is that a trait turn out to be useful (or fail to be a hindrance) most of the time, in the specific circumstances in which creatures possessed of that trait actually live. That is to say, while the model I am proposing would not be plausible if the self-consciously moral motive, and the actual moral facts, did not by and large lead to us acting in adaptive ways, this doesn't imply that the

³³⁴Of course, there are cases of this which can increase reproductive fitness, such as a person sacrificing themselves to protect their children or other blood relatives, but it would (plausibly) be no less good for a person to sacrifice themselves to save strangers or adopted children, which would have the opposite effect on one's reproductive fitness.

course of action that is morally best in any given situation must necessarily always be the most adaptive one. The same holds true for the epistemic case, preserving the analogy; true beliefs are generally good for us but this needn't always be the case. Some beliefs may be so distressing as to fundamentally undermine an agent's ability to survive and reproduce.³³⁵ Nevertheless, the practical value of true beliefs remains in place for a sufficient majority of cases that our concern for truth is adaptive. So rather than implying that what is moral is always what increases reproductive fitness, which is implausible, the model I have suggested implies only that sensitivity to moral considerations (when reflected in an agent's actions) will generally or mostly increase the reproductive fitness of the agent, which is already implied by debunking accounts of moral phenomena and so should be unsurprising to anyone who has followed my arguments this far.

While there are problems arising from the analogy between the epistemic and moral cases, none of these seem to undermine the part of the analogy that matters, which is the way in which normative facts—facts about what given agents have reason to do—can depend on the facts within one of these domains so long as its characteristic concern is one which is reason-giving for the agent in question. As such we have seen that of the arguments based on these disanalogies which are meant to undermine my model for non-special moral norms (in conjunction with Purely Evaluative Moral Realism), none succeed. If another disanalogy could be found, and found in the way moral facts relate to reasons and action (without appealing to Normativism, of course), this might be sufficient to ground a more successful argument in this vein. At present, however, no such disanalogy presents itself to me, so this class of argument is insufficient to motivate a rejection of the extended purely evaluative view.

4.3.2 Denial of the Self-Consciously Moral Motive

Disanalogies between the epistemic and moral cases are not the only possible grounds on which one might object to the view that agents are subject to non-special moral norms because of their self-consciously moral motives. Another class

³³⁵For a fanciful example, consider the writings of H. P. Lovecraft, where characters are driven to madness or suicide by the truths they learn. This very much reduces their reproductive fitness.

of objection is those rooted in the rejection of my claim that humans generally have some self-consciously moral motive. This could be taken as another example of an objection rooted in disanalogy, if it is made in concert with a commitment to the claim that we do have the kind of general concern for truth that I have posited, though the arguments themselves entail no such commitment. My realism about non-special moral norms relies on the claim that we have a self-consciously moral motive to explain the action-guiding power of moral facts, so denying this leaves us back where we started: trying to deal with the objection that the purely evaluative view has no room to grant moral facts this power. The self-consciously moral motive could also explain the existence and (potential) truth of many of our moral intuitions, which is another of my expanded picture's major virtues. So, this view would be in serious trouble if it could be shown that we do not, in fact, have any such motive as part of our typical psychological makeup.

I am not the first to posit self-consciously moral motives of this kind.³³⁶ This means that I do not need to preempt those who would take issue with the view I have suggested; many arguments against the existence of these motives have already been made. None of them, however, strike me as particularly persuasive. Below, I will outline the ones I have found and state why I believe them to be unsuccessful.

One such argument comes from Michael Smith, who originally raised it against Brink's motivational externalism. Smith claims that if a model of non-special moral norms like mine were accurate, this would make our concern for moral (or morally relevant) features of the world such as honesty and loyalty derivative.³³⁷ The issue, Smith claims, is that we ought to see these features as reason-giving in a more direct way and that there is something perverse about responding to them only because one is committed to good in the *de dicto* sense and happens, at the

³³⁶They are also invoked by Foot's (1972) morality as a series of hypothetical imperatives and the realism without clout advanced by Tresan (2010) and Baeten (2012). Svavarsdóttir (1999) also posited a very close parallel under the name 'the desire to be moral', and Morgan (2006, p. 325) suggests that certain forms of naturalism need to posit motives of the kind I have described to do precisely the work I have invoked them to do.

³³⁷Smith 1994, p. 75.

time in question, to take it that these features of the world are good.

The problem with this response, though, becomes apparent when we notice terms like ‘ought’ and ‘perverse’ making their way into the argument. Smith infers motivational externalism’s commitment to a self-consciously moral motive by noting that a change in moral opinion on the part of a good and strong-willed person reliably leads to a change in behaviour, and that said motive is necessary to explain this connection within an externalist framework. So Smith’s argument centres on the motives and interests of good people (and so implicitly on the motives and interests that we ought to have, though of course Purely Evaluative Moral Realists deny the validity of this move). This is a problem, but not because it runs counter to Purely Evaluative Moral Realism’s denial of special normativity. Rather, the issue is that this argument doesn’t target the claims I make in my analogy. My claims about self-consciously moral motives are concerned with the actual psychology of human beings in general. Perhaps Smith is right about what it takes to be a good person, but then what of it? It is still plausibly the case that humans have a self-consciously moral motive (or even some set thereof) all the same.

In fact, the kind of direct concern for morally valuable things alluded to by Smith isn’t mutually exclusive with a self-consciously moral motive. It has been suggested that such a motive helps us to balance our priorities when our direct (moral) concerns pull us in multiple different directions.³³⁸ Even the very existence of these kinds of conflicts, in the form of moral dilemmas, points to how a good person can be concerned with what is good or right or best in the *de dicto* sense even when they are not sure what it is.³³⁹ Again, it doesn’t matter much to my view what good people are like so much as what people are like on the whole, but surely finding oneself in a moral dilemma is something that can happen to anyone and certainly not a sign that there is anything unusual or perverse about their moral psychology. I do not claim that self-consciously moral motives are the only moral motives. All I claim is that some kind of self-consciously moral motive will typically figure in human moral psychology, so it might well be held alongside

³³⁸Railton 1984, pp. 149-50 and Beaulieu 2007, pp. 755-8.

³³⁹Sadler 2003, p. 69.

more direct concern for certain moral values (although I suspect such concerns involve—at least—an implicit recognition that their objects are the same kind of thing targeted by a more general concern for good, which is to say a self-consciously moral motive).³⁴⁰

One further response open to us is that, insofar as Smith's claims can be taken as claims about what is morally required of us, they beg the question against Purely Evaluative Moral Realism. Notably, his remarks are largely framed in terms of what makes for a good person, which are well within the bounds of what the purely evaluative view can accept, but it is quite plausible that, as a motivational internalist, Smith intends us to take his remarks about what is good as not just normative but specially normative. If this is so, then this argument is unfit to undermine the conjunction of Purely Evaluative Realism and realism about non-special moral norms in two ways.

Another case against self-consciously moral motives is rooted in the suggestion that the kind of views that require us to posit them render self-consciously moral motives (as I have characterized them) impossible. So, like Smith's argument, this argument isn't focused on what motives agents actually have. However, since it concerns the possibility that we have self-consciously moral motives in light of precisely the kind of claims Purely Evaluative Realism makes, it poses a more significant threat to the extended view.

Seiriol Morgan has pointed out that one issue for systems that must posit a self-consciously moral motive on the part of good agents is that they do not allow for one set of values to have 'normative superiority' over alternatives.³⁴¹ This seems to be because such systems need to posit these motives in the first place because they cannot appeal directly to special normativity. If this is so, then the same concern can be raised against any form of Purely Evaluative Moral Realism. But so far these claims amount to nothing that has not already been acknowledged. The point at which this becomes a legitimate objection to the view I advance is

³⁴⁰We see this kind of pluralism about moral motives in Railton 1984 and Norcross 1997, both in response to Williams's (1981, p. 18) claim that certain kinds of moral reasoning involve 'one thought too many', which is a concern specifically invoked by Smith (1994, p. 75).

³⁴¹Morgan 2006, p. 333.

when we make the further leap of supposing that a lack of special normativity implies an inability to privilege one value set, and that as a consequence a motive that makes explicit reference to moral concepts has nothing to ‘latch onto’ and has no practical upshot.

The crux of this objection is that, lacking special normativity, Morgan takes it that there is nothing to determine which set of values is the right one, morally speaking. Naturally, Purely Evaluative Moral Realists will want to deny this. I think that the exact answers to these kinds of questions are for substantive ethicists to determine, but this may not satisfy some. Fortunately, we already have everything we need to put together the beginnings of an answer to this question.

Those who hold the purely evaluative view need not reject the picture of moral psychology presented by our debunking explanation of moral phenomena, and my realism about non-special moral norms is explicitly premised on that picture. Part of this picture is that moral thought, or at least the capacity therefor, is a basic feature of our psychology. As Joyce observes, morality seems to arise in every society, and there are certain commonalities with regard to what it takes as its central concerns.³⁴² Specifically, it seems to be primarily concerned with the way we relate to other people, and certain other things in the world.³⁴³ This frequently relates to the distribution of non-moral goods (and evils), and I suspect—but cannot demonstrate definitively—that in more internally diverse societies the goods in question become increasingly abstract. That these patterns should be so prevalent indicates that there are limits on what could conceivably count as a moral fact, which further suggests that concerns that self-consciously moral motives would

³⁴²Joyce 2006, p. 9.

³⁴³In some regards this seems very intuitive, though certain elements of it may raise eyebrows among contemporary ethicists. That morality governs how we relate to other humans and certain non-human animals sounds reasonable enough, but that it should govern what it is appropriate to eat—which is one of three major patterns Joyce (2006, ch. 2) highlights—will strike many, myself included, as at least somewhat odd. However, since Purely Evaluative Moral Realists are committed to the idea that our moral thought is at least somewhat error prone—it does, after all, tend to operate in terms of categorical imperatives, which they reject—perhaps this should not be a cause for surprise or concern.

lack content to be unfounded.³⁴⁴

One final way to challenge the idea that we have self-consciously moral motives is perhaps the most obvious and direct, but at the same time the least promising. As I have already noted, my concern here is more a matter of empirical reality than structure. The question is: do we, in fact, have self-consciously moral motives? If it could be shown that the empirical evidence suggests we do not, then that would be a significant threat to the expanded view I have presented in this chapter. However, I do not think that straightforwardly denying that we have self-consciously moral motives is a tactic that is likely to succeed. One could make an argument from redundancy, appealing to our direct moral concerns in the same way Smith does. But this would not strictly rule out the presence of a self-consciously moral motive, nor could it account for our experience of moral dilemmas. It seems as though when we find ourselves pulled in different directions by two moral values (even if we value each of them directly), we feel that one option is genuinely better and feel driven to determine which it is. This indicates both that moral agents posit more abstract moral qualities and that they place non-derivative value on the pursuit thereof.

None of the attempts to deny the presence or possibility of the self-consciously moral motive have succeeded. As such, a view that combines Purely Evaluative Moral Realism with realism about non-special norms grounded in moral facts cannot be dismissed on the basis that it invokes such motives; available evidence points to them being not only innocent but highly plausible.

In this chapter, I have argued against the claim that Purely Evaluative Moral Realism leaves no possibility for moral facts to be action-guiding. I have done this in two ways. Firstly, I have done it by laying out how existing substantive ethical systems can fit into a purely evaluative framework and still offer us guidance for our actions. Secondly, I have done it by detailing one of many potential ways of

³⁴⁴This may prompt the immediate objection that this view takes away the specialness of morality, the reason to be concerned with it over any other form of value. This isn't much of an objection because, as I have made clear previously, I take it that the moral domain is of special concern to humans for contingent reasons but does not inherently hold the kind of privileged position Normativists want to assign to it.

expanding on the basic commitments of the purely evaluative view. This view is a sort of abstraction from the substantive views I have laid out, making explicit the way that a self-consciously moral motive (a feature that can be invoked by all of those views) can turn moral facts into genuine, though non-special, normative facts.

We have seen that this extended view answers the objections raised in §3.3.3 and survives a number of further objections. Hence there is at least one way to account for the action-guiding nature of morality under Purely Evaluative Moral Realism. With that, the last—and perhaps the greatest—of the objections to Purely Evaluative Moral Realism that I have foreseen has been dealt with.

Conclusions and Summary

At the start of this project, we saw how the Argument from Queerness presents a threat to moral realism by challenging the special normativity of moral facts. Next, we saw how this issue only presents itself on the assumption of Strong Normativism. I called this assumption into question by denying a necessary entailment from evaluative moral facts to specially normative facts. Then I defended that denial by rejecting a number of entailment views about these categories of moral facts.

This set the stage for Purely Evaluative Moral Realism. Purely Evaluative Moral Realism avoids the Argument from the Queerness of Special Normativity since it denies that a normative structure is intrinsic to any actual moral fact. What's more, I argued for the possibility of a naturalistic version of this view, which would allow us to find a way around the Argument from the Queerness of Supervenience in the event that an argument could be given for one of the claims that might motivate it. So it seems that Purely Evaluative Moral Realism can avoid the Argument from Queerness. This is assumes that it can pass the test of basic defensibility and plausibility, which was the next question to which we attended.

In the course of this defence, I had to respond to several objections. One particular objection required an especially involved response. This was the question of whether moral facts as, Purely Evaluative Moral Realism understands them, are fit to guide action in the way typically expected. However, a response to these concerns is readily available, for we have seen that many defensible views within substantive ethics can operate with a Purely Evaluative Moral Realist system of metaethics at their back. From there, we can even abstract to a more general picture of how evaluative moral facts can guide action by being realists about non-

special norms that are reducible to those moral facts and a self-consciously moral motive.

My defence of Purely Evaluative Moral Realism has been by and large just that; a defence. While I have made mention of the view's virtues (and will continue to do so in Appendix A), my primary objective has been to show that a committed Purely Evaluative Moral Realist would be at least as reasonable as those who plump for more established alternatives.

It does give moral realists a way around the Argument from Queerness, but that way is itself a limited form of antirealism and it requires the realist to give up much that they may be unwilling to surrender. While it preserves what are perhaps the most intuitive moral facts—that some things are good or bad—it also forces us to either severely restrict our understanding of normative moral facts, or to do away with them altogether.

Appendix A

Virtues of Purely Evaluative Moral Realism

When I outlined Purely Evaluative Moral Realism, I devoted considerable space to defending the view from anticipated objections, but little to extolling any virtues of the view beyond avoiding the Argument from Queerness. Yet more can be said in support of the purely evaluative position.

If the best thing we could say about Purely Evaluative Moral Realism were that objections to it failed, it would remain open to one further objection that I didn't examine in Chapter 3, namely that it is an ad hoc response to the Argument from Queerness, and thus undermotivated.

The question this thesis aims to answer is strictly concerned with whether Purely Evaluative Realism is coherent in itself and how it relates to the Argument from Queerness, which is why I have reserved this discussion for an appendix. But clearly there is still value in discussing the virtues of the view I have developed. It will show that Purely Evaluative Realism can be motivated by concerns beyond the Argument from Queerness.

A.1 Dissolving the Is-Ought Problem

David Hume claimed that works of moral philosophy always make an unacknowledged shift somewhere from claims about what is or is not to claims about what

ought or ought not to be.³⁴⁵ He clearly sees this move as objectionable, though it is unclear whether this is because he believes that it cannot be justified, or simply that it has not been.³⁴⁶ Either way, this observation highlights a challenge for moral philosophers.

This is the challenge of explaining the relationship between what *is* good (whatever that amounts to in naturalistic terms) and what we *ought* to do. Some systems of moral philosophy, like utilitarianism, posit a very explicit connection of this kind, but even where the precise nature of the connection is not spelled out there is a strong intuition that there must be some kind of link between what is good and what is right or obligatory. The point Hume highlights is that the two categories do seem to be radically different. This is something I highlighted in §1.7 and §2.1. Given that, any connection between them cannot simply be posited or assumed; it must be explained. Hence moral realists in particular, who posit these qualities as real constituents of the world, are saddled with an explanatory burden.

One of the advantages of Purely Evaluative Moral Realism, relative to more conventional forms of realism, is that it dissolves this problem entirely. Its central claims are that there are evaluative moral facts, and that there are no normative moral facts. Evaluative moral facts are specifically what we might call is-type facts; they do not entail anything normative (and so they cannot entail anything deontic, which are what we might call ought-type facts). If is-type moral facts exist and ought-type moral facts do not exist, as Purely Evaluative Moral Realism maintains, then there simply cannot be any such connection between them as we are wont to assume, and if no such connection exists there can be no call to explain it.

It might be objected that all Purely Evaluative Moral Realism does is push the problem back. The issue that it addresses is the apparent gap between the evaluative and the normative, but even having sidestepped that problem, it still has to deal with the same gap between the descriptive and the evaluative. A Purely Evaluative Moral Realist can respond that this problem is one it shares with rival forms of moral realism, in which case it is none the worse for being unable to

³⁴⁵Hume, T 3.1.2.27; 1978, p. 469; 2007, p. 302.

³⁴⁶Cohon 2018, §5.

solve it (and still superior insofar as it does not shoulder the explanatory burden arising from the gap between the evaluative and the normative). But perhaps both these gaps represent a more general problem around supervenience, in which case there is only one solution, and a view that has only one of these problems is in no less need of it than a view that has both. I see no reason to assume that this is the case, but there is still a way for Purely Evaluative Moral Realists to preserve their view's advantage over the alternatives. By breaking the connection between the evaluation and normativity, Purely Evaluative Moral Realism makes evaluation nothing more than a special case of description. This being the case, the purely evaluative view is still free of any explanatory burden because there is no explanatory gap to fill between the descriptive and the evaluative.

Purely Evaluative Moral Realism also gives us the means to explain the intuition that there should be some connection between the evaluative and the normative. Remember that, as discussed in §1.3, the ideal debunking account can explain a disposition to think of the moral domain as categorically imperative (this would, after all, reinforce the adaptive behaviour inculcated by moral beliefs) regardless of its actual structure or content. Purely Evaluative Moral Realists need not be committed to this picture, of course, but they are not barred from it either, and those who reject more conventional forms of realism on the basis of the Argument from Queerness will already be amenable to it. As such, there is no ground to object that Purely Evaluative Moral Realism achieves this dissolution only by running unacceptably counter to our intuitions about the connection between goodness and what we ought to do.

A.2 Free Will and Moral Responsibility

Another advantage of Purely Evaluative Moral Realism is that it offers us a way around at least one other sceptical argument besides the Argument from Queerness. Specifically, it defuses concerns raised for substantive ethics if one denies the existence of free will.

The denial of free will is well motivated. There are a number of good arguments backing up this position. Peter van Inwagen points out that if some form of determinism is true and so everything that happens is a consequence of everything

that came before and the laws of nature (as seems plausible), then, because these things are not in our power to control, our actions are not in our power to control either.³⁴⁷ We cannot do otherwise than we actually do.

If, on the other hand, some form of indeterminism is true and there is a range of possible futures open to us, we seem to reach the same conclusion by a different path. The range of futures available to us must be constrained in the sense that they must all be compatible with the nature of each agent who plays a part in making them.³⁴⁸ The actions available to any particular agent must all be equally compatible with their nature. The choice between them, then, comes down to chance rather than to the agent.³⁴⁹ Random chance is not a decision, and there is nothing we can do to influence which one of the range of possible actions we will take. So we cannot do otherwise than we actually do.

Crucially, the claim that we have free will seems to be a claim about what we *can* do. If one is free to choose between soup and salad, then it is true both that one can choose the soup and that one can choose the salad. The standard way to understand free will is as the ability to do (or have done) differently than one actually does (or did).³⁵⁰ If we cannot change what we will do, we cannot do (or have done) differently than we actually do (or did), and it seems that regardless of whether determinism or indeterminism is the case, we cannot act otherwise than we actually do.

Meanwhile, there is a widely accepted principle that ‘ought’ implies ‘can’. In other words, it follows from the fact that you ought to perform some action that said action is one that you can perform. By modus tollens it follows that if you cannot do something, whatever it might be, it is not the case that you ought to

³⁴⁷van Inwagen 1983, p. 16. Arguments to this effect are also to be found in Kant 1996, *KpV* 5:94–95, pp. 215–216 and 1998, *KrV* A536/B564, p. 535.

³⁴⁸van Inwagen 1998, pp. 462–3.

³⁴⁹van Inwagen 1983, p. 16.

³⁵⁰This account dates back at least as far as the mid 1600s (see Hobbes 1999, p. 16). What the relevant sense of ability is has been a topic of much debate (see O’Connor and Franklin 2020, §2.2 for a summary) but all that matters for our purposes is that it be the same sense of ‘can’ that is invoked elsewhere in this section.

do that thing. (Likewise, if you cannot act otherwise than you actually do, then—assuming that there is something that you ought to do—you cannot fail to do it.)

Traditional forms of realism are replete with ‘ought’ facts, or deontic facts. In §1.8 and §2.1, I defined these facts as a subset of normative facts. They are facts about what we have most reason to do. By contrast Purely Evaluative Moral Realism posits no such facts.

The problem for traditional forms of moral realism, which posit ‘oughts’ in no short supply, is this: if it turns out that we don’t have free will, there will be little sense in making prescriptions of any kind, moral ones included. If ‘ought’ implies ‘can’ and no agent could ever have done differently than they actually did or exert any kind of real control over their actions, we cannot reasonably claim that they ought to have acted in any other way. Moral claims of the form ‘Some agent A ought to ϕ ’ could then only be true if A did, will, or would ϕ under the appropriate circumstances. This runs strongly counter to our intuitions about moral facts and normativity. We normally accept that what a person should do can differ from what they actually do.

This seems to make a nonsense out of moral ‘ought’ claims, and thus to invalidate a large proportion of the claims made by traditional forms of moral realism. It has the same effect whether we see these ‘oughts’ as grounded in special normativity, or in the hypothetical normativity invoked by some of the views discussed in §3.1. Purely Evaluative Moral Realism doesn’t face this problem. Because this view only posits evaluative moral facts, none of the claims it permits its adherents to make are threatened by this consequence of denying free will.³⁵¹ The denial of free will undermines claims to the effect that someone ought or ought not to act or have acted in a certain way, but we can still meaningfully say of any given action that it is or was good or bad, or at least better or worse than alternatives (and while there are strictly no other possible actions if determinism is true, we can still easily imagine reasonable counterfactuals to use as a basis for comparison).³⁵² This

³⁵¹Claims about non-special moral norms of the kind discussed in §4.2 may be threatened, but they are not part of the core of Purely Evaluative Moral Realism.

³⁵²A similar claim can be found in Haji 1998, chs. 3, 9, 11 and Pereboom 2001, pp. 143, 149;

latter claim is true *ex hypothesi* if we already accept Purely Evaluative Moral Realism, because is it committed to the claim that values don't entail norms. (That said, even Strong Normativism only commits us to the possibility that evaluative facts will entail deontic facts, because it remains an open question whether you can have reasons to do something that, on balance, you ought not to do, as discussed in §2.3.6.)³⁵³

Thus it is that the denial of free will undermines most, if not all, of the substantive claims made by those presupposing traditional realism, but not the ones made by those presupposing Purely Evaluative Moral Realism. In other words, the view that I have laid out allows its adherents to retain all of their moral judgements where other views will be forced to abandon at least some of them. There are ways by which traditional moral realists might regain the ability to make at least some of these claims—by modifying the use or meaning of certain moral terms, for example—but further debate will always be needed to show these moves to be legitimate, and Purely Evaluative Moral Realism has the advantage of not requiring them in the first place.

It should be clear at this juncture that Purely Evaluative Moral Realism has merits beyond its immunity to the Argument from Queerness.

2014, p. 131 about the compatibility of evaluative claims with the denial of moral responsibility, which is itself rooted in the denial of free will. I will not try to determine here whether deontic facts are necessary or sufficient for moral responsibility, but if they turned out to be so, then the above claim might entail or be equivalent to the one made by Haji and Pereboom. Either way, I take it that the elimination of moral responsibility and related categories like praiseworthiness and blameworthiness (which, unlike the thinner evaluative terms discussed in Chapter 2, do seem to entail norms where they apply) is perfectly compatible with Purely Evaluative Moral Realism.

³⁵³There is a related but separate debate over whether one can have reason to do something you cannot do (Streumer, 2017, §65), which is particularly relevant here. To say one cannot is a neat (though not the only possible) explanation of the implication from 'ought' to 'can' but both sides of this question are compatible with all variants of Normativism and Non-Normativism.

Bibliography

- Adams, R. M. (1999). *Finite and Infinite Goods: A Framework for Ethics*. New York: Oxford University Press.
- Alston, W. P. (1990). Some suggestions for divine command theorists. In M. D. Beatty (Ed.), *Christian Theism and the Problems of Philosophy*. Notre Dame, IN: University of Notre Dame Press.
- Alvarez, M. (2017). Reasons for action: Justification, motivation, explanation. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Metaphysics Research Lab, Stanford University.
- Anscombe, G. E. M. (1958). Modern moral philosophy. *Philosophy* 33(124), 1–19.
- Anscombe, G. E. M. (1963). *Intention*. Cambridge, MA: Harvard University Press.
- Armstrong, D. M. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Audi, P. (2013). How to rule out disjunctive properties. *Noûs* 47(4), 748–766.
- Audi, R. (1997). *Moral Knowledge and Ethical Character*. Oxford: Oxford University Press.
- Baeten, E. (2012). Another defense of naturalized ethics. *Metaphilosophy* 43(5), 533–550.
- Beaulieu, G. (2007). Meta-ethical rationalism and the amoralist challenge: An externalist response to Michael Smith’s argument. *Dialogue* 46(4), 751–760.

- Bedke, M. S. (2009). Might all normativity be queer? *Australasian Journal of Philosophy* 88(1), 41–58.
- Blackburn, S. (1984). *Spreading the Word*. Oxford: Clarendon Press.
- Blackburn, S. (1993). *How to be an Ethical Anti-Realist*, pp. 166–181. Oxford: Oxford University Press.
- Blackburn, S. (2001). *Being Good: A Short Introduction to Ethics*. Oxford: Oxford University Press.
- Boyd, R. (1988). How to be a moral realist. In G. Sayre-McCord (Ed.), *Essays on Moral Realism*, pp. 181–228. Ithaca, NY: Cornell University Press.
- Brandt, R. B. (1959). *Ethical Theory*. Englewood Cliffs, NJ: Prentice-Hall.
- Brewer, T. (2009). *The Retrieval of Ethics*. Oxford: Oxford University Press.
- Brink, D. O. (1986). Externalist moral realism. *Southern Journal of Philosophy* 24(Supplement), 23–41.
- Brink, D. O. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Brown, C. (2011). Consequentialize this. *Ethics* 121(4), 749–771.
- Brown, P. (2013). The possibility of morality. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 163(3), 627–636.
- Bykvist, K. (2009). No good fit: Why the fitting attitude analysis of value fails. *Mind* 118(469), 1–30.
- Carlson, E. (1995). *Consequentialism Reconsidered*. Dordrecht: Kluwer Academic Publishers.
- Chalmers, D. J. (1996). Naturalistic dualism. In M. Velmans and S. Schneider (Eds.), *The Conscious Mind*, pp. 359–368. Oxford: Oxford University Press.
- Chilovi, S. (2021). Grounding entails supervenience. *Synthese* 198(Suppl 6), 1317–1334.

- Cline, B. (2016). Against deliberative indispensability as an independent guide to what there is. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 173(12), 3235–3254.
- Cohon, R. (2018). Hume's moral philosophy. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 ed.). Metaphysics Research Lab, Stanford University.
- Crisp, R. (2005). Value, reasons, and the structure of justification: How to avoid passing the buck. *Analysis* 65(1), 80–85.
- Cuneo, T. (2007). *The Normative Web*. Oxford: Oxford University Press.
- Dancy, J. (2000). *Practical Reality*. Oxford: Oxford University Press.
- Dancy, J. (2004). *Ethics without Principles*. Oxford: Clarendon Press.
- Dancy, J. (2006). Nonnaturalism. In D. Copp (Ed.), *The Oxford Handbook of Ethical Theory*, pp. 122–145. Oxford: Oxford University Press.
- Danielsson, S. and J. Olson (2007). Brentano and the buck-passers. *Mind* 116(463), 511–522.
- Davidson, D. (1963). Actions, reasons, and causes. *The Journal of Philosophy* 60(23), 685–700.
- Davidson, D. (1978). Intending. In Y. Yovel (Ed.), *Philosophy of History and Action*, pp. 41–60. Dordrecht: D. Reidel Publishing Company.
- Dreier, J. (1990). Internalism and speaker relativism. *Ethics* 101(1), 1–26.
- Dreier, J. (1993). Structures of normative theories. *The Monist* 76(1), 22–40.
- Durrant, R. G. (1970). Identity of properties and the definition of good. *Australasian Journal of Philosophy* 48(3), 360–361.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.

- Ewing, A. C. (1939). A suggested non-naturalistic analysis of good. *Mind* 48(189), 1–22.
- Ewing, A. C. (1947). *The Definition of Good*. London: Routledge and Kegan Paul.
- Ewing, A. C. (1959). *Second Thoughts in Moral Philosophy*. London: Routledge and Kegan Paul.
- Feldman, F. (2006). Actual utility, the objection from impracticality, and the move to expected utility. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 129(1), 49–79.
- Finlay, S. (2004). The conversational practicality of value judgement. *The Journal of Ethics* 8(3), 205–223.
- Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- FitzPatrick, W. (2008). Robust ethical realism, non-naturalism, and normativity. In R. Schafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 3, pp. 159–205. Oxford University Press.
- Fletcher, G. (2009). Rejecting well-being invariabilism. *Philosophical Papers* 38(1), 21–34.
- Foldy, M. S. (1997). *The Trials of Oscar Wilde: Deviance, Morality, and Late-Victorian Society*. London: Yale University Press.
- Foot, P. (1972). Morality as a system of hypothetical imperatives. *The Philosophical Review* 81(3), 305–316.
- Foot, P. (2001). *Natural Goodness*. Oxford: Clarendon Press.
- Frankena, W. K. (1958). Obligation and motivation in recent moral philosophy. In A. I. Melden (Ed.), *Essays in Moral Philosophy*. Seattle: University of Washington Press.
- Garner, R. T. (1990). On the genuine queerness of moral properties and facts. *Australasian Journal of Philosophy* 68(2), 137–146.

- Gauthier, D. (1986). *Morals by Agreement*. Oxford: Oxford University Press.
- Geach, P. T. (1958). Imperatives and deontic logic. *Analysis* 18(3), 49–56.
- Geach, P. T. (1965). Assertion. *The Philosophical Review* 74(4), 449–465.
- Gert, J. (2016). A fitting end to the wrong kind of reason problem. *Ethics* 126(4), 1015–1042.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis* 23(6), 121–123.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Goldman, A. I. (1967). A causal theory of knowing. *The Journal of Philosophy* 64(12), 357–372.
- Goldman, A. I. (1970). *A Theory of Human Action*. Princeton, NJ: Princeton University Press.
- Greaves, H. (2016). Cluelessness. *Proceedings of the Aristotelian Society* 116(3), 311–339.
- Greene, J. (2013). *Moral Tribes: Emotion, Reason, and the Gap between Us and Them*. New York: Penguin Books.
- Haji, I. (1998). *Moral Appraisability: Puzzles, Proposals, and Perplexity*. New York: Oxford University Press.
- Hampton, J. E. (1998). *The Authority of Reason*. Cambridge: Cambridge University Press.
- Hare, R. M. (1952). *The Language of Morals*. Oxford: Clarendon Press.
- Harman, G. (1977). *The Nature of Morality: An Introduction to Ethics*. Oxford: Oxford University Press.
- Harsanyi, J. C. (1953). Cardinal utility in welfare economics and in the theory of risk-taking. *Journal of Political Economy* 61(5), 434–435.

- Hester, M. (1991). Aristotle on the function of man in relation to eudaimonia. *History of Philosophy Quarterly* 8(1), 3–14.
- Hobbes, T. (1654/1999). Of liberty and necessity. In V. Chappelle (Ed.), *Hobbes and Bramhall on Liberty and Necessity*, pp. 15–42. Cambridge: Cambridge University Press.
- Huemer, M. (2009). Values and morals: Outline of a skeptical realism. *Philosophical Issues* 19, *Metaethics*, 113–130.
- Hume, D. (1739/1978). *A Treatise of Human Nature* (2nd ed.). Oxford: Clarendon Press.
- Hume, D. (1739/2007). *A Treatise of Human Nature: A Critical Edition*. Oxford: Clarendon Press.
- Hurka, T. (1998). Two kinds of organic unity. *Journal of Ethics* 2(3), 299–320.
- Hursthouse, R. and G. Pettigrove (2018). Virtue ethics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2018 ed.). Metaphysics Research Lab, Stanford University.
- Ichikawa, J. J. and M. Steup (2018). The analysis of knowledge. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2018 ed.). Metaphysics Research Lab, Stanford University.
- Irwin, T. H. (1980). The metaphysical and psychological basis of Aristotle's ethics. In A. O. Rorty (Ed.), *Essays on Aristotle's Ethics*, pp. 35–53. Berkeley and Los Angeles, CA: University of California Press.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.
- Joyce, R. (2001). *The Myth of Morality*. Cambridge: Cambridge University Press.
- Joyce, R. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Kahane, G. (2011). Evolutionary debunking arguments. *Noûs* 45(1), 103–125.

- Kant, I. (1996). *Practical Philosophy*. Cambridge: Cambridge University Press.
- Kant, I. (1998). *Critique of Pure Reason*. Cambridge: Cambridge University Press.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research* 66(3), 612–640.
- Kim, J. (1984). Concepts of supervenience. *Philosophy and Phenomenological Research* 45(2), 153–176.
- Kim, J. (1993). *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.
- Kirchin, S. (2013). Evaluation, normativity, and grounding. *Proceedings of the Aristotelian Society, Supplementary Volumes* 87(1), 179–198.
- Kitcher, P. (2011). *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Korsgaard, C. M. (1986). Skepticism about practical reason. *The Journal of Philosophy* 83(1), 5–25.
- Korsgaard, C. M. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Kraut, R. (2018). Aristotle's ethics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2018 ed.). Metaphysics Research Lab, Stanford University.
- Kripke, S. (1980). *Naming and Necessity*. Oxford: Blackwell.
- Lefka, A. (2021). Ancient greek philosophy and current definitions of eudaimonia. In P. Mitsis and H. L. Reid (Eds.), *The Poetry in Philosophy: Essays in Honor of Christos C. Evangeliou*. Sioux City, IA: Parnassos Press.
- Lewis, D. K. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. K. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

- Louise, J. (2004). Relativity of value and the consequentialist umbrella. *Philosophical Quarterly* 54(217), 518–536.
- Mackie, J. L. (1946). A refutation of morals. *Australasian Journal of Psychology and Philosophy* 24(1-2), 77–90.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin Books.
- Mackie, J. L. (1980). *Hume's Moral Theory*. London: Routledge and Kegan Paul.
- Mackie, J. L. (1982). *The Miracle of Theism: Arguments for and Against the Existence of God*. Clarendon Press.
- McDowell, J. (1979). Virtue and reason. *The Monist* 62(3), 331–350.
- McNaughton, D. and P. Rawling (1992). Honoring and promoting values. *Ethics* 102(4), 835–843.
- McPherson, T. (2012). Ethical non-naturalism and the metaphysics of supervenience. In R. Schafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 7, pp. 205–234. Oxford: Oxford University Press.
- Mehta, N. (2017). Phenomenal, normative, and other explanatory gaps: [a] general diagnosis. *Philosophy and Phenomenological Research* 98(3), 567–591.
- Meinertsen, B. R. (2021). Against disjunctive properties: Four armstrongian arguments. *Philosophia* 49(1), 95–106.
- Moberger, V. (2020). Hume's dictum and metaethics. *Philosophical Quarterly* 70(279), 328–349.
- Moore, G. E. (1903). *Principia Ethica*. Cambridge: Cambridge University Press.
- Moore, G. E. (1912). *Ethics*. London: Williams and Norgate.
- Moore, G. E. (1942). A reply to my critics. In P. A. Schlipp (Ed.), *The Philosophy of G. E. Moore*, pp. 535–687. New York: Tudor Publishing Company.

- Moore, G. E. (1959). *Philosophical Papers*. London: George Allen and Unwin.
- Morgan, S. (2006). Naturalism and normativity. *Philosophy and Phenomenological Research* 72(2), 319–344.
- Nagel, T. (1972). Aristotle on eudaimonia. *Phronesis* 17(3), 252–259.
- Nolan, D. (1997). Quantitative parsimony. *British Journal for the Philosophy of Science*, 329–343.
- Norcross, A. (1990). Consequentialism and the unforeseeable future. *Analysis* 50(4), 253–256.
- Norcross, A. (1997). Consequentialism and commitment. *Pacific Philosophical Quarterly* 78(4), 380–403.
- Norcross, A. (2006). Reasons without demands: Rethinking rightness. In J. Dreier (Ed.), *Contemporary Debates in Moral Theory*, pp. 38–53. Oxford: Blackwell.
- Norcross, A. (2020). *Morality by Degrees: Reasons without Demands*. Oxford: Oxford University Press.
- Nowell-Smith, P. H. (1954). *Ethics*. Harmondsworth: Penguin Books.
- O'Connor, T. and C. Franklin (2020). Free will. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2020 ed.). Metaphysics Research Lab, Stanford University.
- Olson, J. (2006). G. E. Moore on goodness and reasons. *Australasian Journal of Philosophy* 84(4), 525–534.
- Olson, J. (2011). In defense of moral error theory. In M. Brady (Ed.), *New Waves in Metaethics*, pp. 62–84. Basingstoke: Palgrave Macmillan.
- Olson, J. (2013). Buck-passing accounts. In H. LaFollette (Ed.), *The International Encyclopedia of Ethics*, Volume II, pp. 625–636. Chichester: Blackwell.
- Olson, J. (2014). *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press.

- Olson, J. (2017). Error theory in metaethics. In T. McPherson and D. Plunkett (Eds.), *The Routledge Handbook of Metaethics*, pp. 58–71. London: Routledge.
- Parfit, D. (1997). Reasons and motivation. *Aristotelian Society Supplemental Volume 71*(1), 99–130.
- Parfit, D. (2011a). *On What Matters*, Volume 2. Oxford: Oxford University Press.
- Parfit, D. (2011b). *On What Matters*, Volume 1. Oxford: Oxford University Press.
- Pereboom, D. (2001). *Living Without Free Will*. Cambridge: Cambridge University Press.
- Pereboom, D. (2014). *Free Will, Agency, and Meaning in Life*. Oxford: Oxford University Press.
- Pettit, P. (1989). Consequentialism and respect for persons. *Ethics 100*(1), 116–126.
- Pettit, P. (1991). Consequentialism. In P. Singer (Ed.), *A Companion to Ethics*, pp. 230–240. Oxford: Blackwell.
- Pigden, C. R. (2007). Nihilism, Nietzsche, and the doppelganger problem. *Ethical Theory and Moral Practice 19*(5, Moral Skepticism: 30 Years of Inventing Right and Wrong), 141–156.
- Platts, M. (1980). Moral reality and the end of desire. In M. Platts (Ed.), *Reference, Truth, and Reality: Essays on the Philosophy of Language*, pp. 69–82. London: Routledge and Kegan Paul.
- Portmore, D. W. (2007). Consequentializing moral theories. *Pacific Philosophical Quarterly 88*(1), 39–73.
- Priest, G. (2008). The closing of the mind: [h]ow the particular quantifier became existentially loaded behind our backs. *The Review of Symbolic Logic 1*(1), 42–55.
- Putnam, H. (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.

- Quinn, P. L. (2007). Theological voluntarism. In D. Copp (Ed.), *The Oxford Handbook of Ethical Theory*. New York: Oxford University Press.
- Rabinowicz, W. (2008). Value relations. *Theoria* 74(1), 18–49.
- Rabinowicz, W. and T. Rønnow-Rasmussen (2004). The strike of the demon: On fitting pro-attitudes and value. *Ethics* 114(3), 391–423.
- Railton, P. (1984). Alienation, consequentialism, and the demands of morality. *Philosophy and Public Affairs* 13(2), 134–171.
- Railton, P. (1986/2003a). Facts and values. In *Facts, Values, and Norms*, pp. 43–68. Cambridge: Cambridge University Press.
- Railton, P. (1986/2003b). Moral realism. In *Facts, Values, and Norms*, pp. 3–42. Cambridge: Cambridge University Press.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Raz, J. (1999a). *Engaging Reason: On the Theory of Value and Action*, Chapter The Amoralist, pp. 273–302. Oxford University Press.
- Raz, J. (1999b). *Engaging Reason: On the Theory of Value and Action*, Chapter Agency, Reason, and the Good, pp. 22–45. Oxford University Press.
- Raz, J. (2001). *Value, Respect, Attachment*. Cambridge: Cambridge University Press.
- Robb, G. (2005). *Strangers: Homosexual Love in the Nineteenth Century*. London: W. W. Norton and Company.
- Rowland, R. (2019). *The Normative and the Evaluative: The Buck-Passing Account of Value*. Oxford: Oxford University Press.
- Russell, B. (1948). *Human Knowledge: Its Scope and Limits*. London: George Allen and Unwin.
- Sadler, B. J. (2003). The possibility of amoralism: A defence against internalism. *Philosophy* 78(303), 63–78.

- Scanlon, T. M. (1998). *What We Owe to Each Other*. London: Bellknap Press of Harvard University Press.
- Scanlon, T. M. (2014). *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Schoedinger, A. B. (2007). Nonreductive ethical naturalism. *The Proceedings of the Twenty-First World Congress of Philosophy 1*, 3–6.
- Schroeder, M. (2005). Cudworth and normative explanation. *Journal of Ethics and Social Philosophy 1*(3), 1–28.
- Schroeder, M. (2007). Teleology, agent-relative value, and ‘good’. *Ethics 117*(2), 265–295.
- Shah, N. and J. D. Velleman (2005). Doxastic deliberation. *The Philosophical Review 114*(4), 497–534.
- Shepski, L. (2008). The vanishing argument from queerness. *Australasian Journal of Philosophy 86*(3), 371–387.
- Sinnott-Armstrong, W. (2006). *Moral Skepticisms*. New York: Oxford University Press.
- Slote, M. (1989). *Beyond Optimizing*. Cambridge, MA: Harvard University Press.
- Slote, M. and P. Pettit (1984). Satisficing consequentialism. *Proceedings of the Aristotelian Society, Supplementary Volumes 58*, 139–176.
- Smart, J. J. C. (1973). An outline of a system of utilitarian ethics. In J. J. C. Smart and B. Williams (Eds.), *Utilitarianism: For and Against*, pp. 1–74. Cambridge: Cambridge University Press.
- Smith, M. (1994). *The Moral Problem*. Oxford: Blackwell.
- Sober, E. (1982). Why logically equivalent predicates may pick out different properties. *American Philosophical Quarterly 19*(2), 183–189.

- Stampe, D. W. (1987). The authority of desire. *The Philosophical Review* 96(3), 335–381.
- Stratton-Lake, P. (2002). Introduction. In P. Stratton-Lake (Ed.), *Ethical Intuitionism: Re-evaluations*, pp. 1–28. Oxford: Clarendon Press.
- Street, S. (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 127(1), 106–166.
- Streumer, B. (2017). *Unbelievable Errors: An Error Theory about All Normative Judgements*. Oxford: Oxford University Press.
- Sturgeon, N. L. (1985). Moral explanations. In J. Rachels (Ed.), *Ethical Theory 1: The Question of Objectivity*, pp. 49–78. Oxford University Press.
- Sturgeon, N. L. (2006). Ethical naturalism. In D. Copp (Ed.), *The Oxford Handbook of Ethical Theory*, pp. 91–121. New York: Oxford University Press.
- Svavarsdóttir, S. (1999). Moral cognitivism and motivation. *The Philosophical Review* 108(2), 161–219.
- Thompson, J. J. (2008). *Normativity*. Chicago, IL: Open Court.
- Tresan, J. (2010). Question authority: In defense of moral realism without clout. *Philosophical Studies* 150(2), 221–238.
- Urmson, J. O. (1958). Saints and heroes. In A. I. Melden (Ed.), *Essays in Moral Philosophy*, pp. 198–216. Seattle: University of Washington Press.
- van Inwagen, P. (1983). *An Essay on Free Will*. Oxford: Clarendon Press Press.
- van Inwagen, P. (1998). The mystery of metaphysical freedom. In P. van Inwagen and D. W. Zimmerman (Eds.), *Metaphysics: The Big Questions* (2nd ed.), pp. 456–65. Malden, MA: Blackwell.
- van Roojen, M. (2010). Moral rationalism and rational amoralism. *Ethics* 120(3), 495–525.

- Way, J. (2013). Value and reasons to favour. In R. Schafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 8, pp. 27–49. Oxford University Press.
- Wedgwood, R. (2002). The aim of belief. *Noûs* 36(s16), 267–297.
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- Wedgwood, R. (2009a). The “good” and the “right” revisited. *Philosophical Perspectives* 23(1), 499–519.
- Wedgwood, R. (2009b). Intrinsic values and reasons for action. *Philosophical Issues* 19, *Metaethics*, 321–342.
- Wiggins, D. (1987). *Truth, Invention, and the Meaning of Life*, pp. 87–138. Oxford: Clarendon Press.
- Wiggins, D. (1990). Moral cognitivism, moral relativism and motivating moral beliefs. *Proceedings of the Aristotelian Society* 91(1), 61–85.
- Williams, B. (1979). Internal and external reasons. In R. Harrison (Ed.), *Rational Action*, pp. 17–28. Cambridge: Cambridge University Press.
- Williams, B. (1981). *Moral Luck: Philosophical Papers 1973-1980*, Chapter Persons, Character, and Morality, pp. 1–19. Cambridge University Press.
- Williams, B. (1985). *Ethics and the Limits of Philosophy*. London: Fontana Press.