The Trouble About Moral Facts

Matthew Scarfone

Department of Philosophy

McGill University, Montreal, Quebec

December 2018

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Doctor of Philosophy

TABLE OF CONTENTS

ABSTRACT

We do not have a good account of moral knowledge. That makes this dissertation a defense of moral scepticism. In arriving at this position I canvass prominent and plausible moral epistemologies and find them all flawed.

In chapter 1 I argue that we should reject an ontological criterion of success in favour of an epistemological one. In chapter 2 I argue that we should accept *cognitivism* about our moral judgments and *nonnaturalism* about moral properties. If we cannot vindicate *nonnaturalist cognitivist realism* then we should accept *epistemic moral scepticism*: the view that, even if there are moral facts, we are not in a position to know them. In chapter 3 I address the evolutionary debunking argument (EDA). I argue that the EDA is not a *justification* debunking argument but instead a *theory* debunking argument. I recharacterize the EDA as an *inference debunking argument* and show that nonnaturalist cognitivist realism need not fall prey to it.

I then move on to canvassing plausible moral epistemologies. In chapter 4 I explain how exploring the history of moral intuitionism shows why it faces an unresolvable dilemma. In chapter 5 I turn to moral intuitionism's natural alternative: reflective equilibrium. I argue that spelling out this view reveals how it is circular at a critical juncture. And in chapter 6 I address a little discussed alternative to traditional moral epistemologies: metaethical Mooreanism. I argue that metaethical Mooreanism is untenable owing to the inference debunking argument I developed earlier.

Because we lack a plausible moral epistemology for nonnaturalist cognitivist realism we are forced to accept epistemic moral scepticism. I wrap up in chapter 7 by showing why, even if we lack a plausible moral epistemology, that scepticism does not generalize to our non-moral epistemic theories. While the search for a plausible moral epistemology is constrained by the need to vindicate nonnaturalist cognitive realism, the epistemic domain deserves an

alternative construal: epistemic constitutivism. For this reason, the sceptical worries in the dissertation do not generalize beyond the moral domain to the epistemic domain.

RÉSUMÉ

Nous n'avons pas de bonne explication de la connaissance morale. Cela fait de cette thèse une défense du scepticisme moral. Pour arriver à cette position, je passe en revue les épistémologies morales les plus éminentes et plausibles et les trouve toutes imparfaites.

Au chapitre 1, je soutiens que nous devrions rejeter un critère de succès ontologique en faveur d'un critère épistémologique. Dans le chapitre 2, je soutiens que nous devrions accepter le cognitivisme à propos de nos jugements moraux et le non-naturalisme à propos des propriétés morales. Si nous ne pouvons pas défendre le réalisme cognitiviste non-naturaliste, alors nous devrions accepter le scepticisme moral épistémique : l'idée que, même s'il existe des faits moraux, nous ne sommes pas en mesure de les connatre. Au chapitre 3, j'aborde l'argument de démystification évolutionniste (evolutionary debunking argument, EDA). Je soutiens que l'EDA n'est pas un argument de démystification de la justification, mais plutôt un argument de démystification de la théorie. Je reconstruis l'EDA comme un argument de démystification de l'inférence et montre que le réalisme cognitiviste non-naturaliste n'a pas en être la proie.

Je passe ensuite à une revue des épistémologies morales plausibles. Dans le chapitre 4, j'explique comment l'exploration de l'histoire de l'intuitionnisme moral montre pourquoi il est confronté un dilemme insoluble. Au chapitre 5, je passe à la solution naturellement opposée à l'intuitionnisme moral : l'équilibre réflexif. Je soutiens que éclaircir cette approche révèle à quel point elle est circulaire à une étape cruciale. Au chapitre 6, je traite d'une solution alternative peu discutée aux épistémologies morales traditionnelles : le mooréanisme métaéthique. Je soutiens que le mooréanisme métaéthique est intenable en raison de l'argument de démystification de l'inférence que j'ai développé plus tôt.

Faute d'une épistémologie morale plausible pour le réalisme cognitiviste non-naturaliste, nous sommes obligés d'accepter le scepticisme moral épistémique. Je terminerai au chapitre

7 en montrant pourquoi, même si nous n'avons pas d'épistémologie morale plausible, ce scepticisme ne se généralise pas à nos théories épistémiques non morales. Alors que la recherche d'une épistémologie morale plausible est limitée par la nécessité de défendre le réalisme cognitiviste non-naturaliste, le domaine épistémique mérite une construction alternative : le constitutivisme épistémique. Pour cette raison, les inquiétudes sceptiques de la thèse ne se généralisent pas au-delà du domaine moral au domaine épistémique.

ACKNOWLEDGEMENTS

I had a lot of help.

The thesis began as a short candidacy paper in 2014. Thanks to Andrew Reisner for his feedback and encouragement during those early stages.

For the past three years I was lucky to be a part of an amazing thesis work group. Alice Everly, Eric Murphy, and Muhammad Velji each read multiple drafts of every chapter here. You folks are heroes. Your consistent insights, written comments, collegiality, and camaraderie has meant a lot to me. Thanks also to Frédérick Armstrong and Raymond Aldred for filling in the wild-card spot and for your feedback on various chapters.

I road-tested ideas at a few McGill and wider-Montreal Philosophy Events. Thanks to my friends and colleagues there as well. Thanks especially to Joey Van Weelden for his feedback on various material, for being a good friend, and for being a metaethical ally in the pluralist sea of McGill. Thanks to both Le Groupe de Recherche Interuniversitaire sur la Normativité (GRIN) and Le Centre de Recherche en Éthique (CRÉ) for financial support and opportunities to present and discuss this work. Thanks to Éliot Litalien for his help translating the abstract.

I received helpful feedback from audiences at McGill University, l'Université de Montreal, the Centre for Research on Religion, the University of Calgary, the University of Sheffield, the 2018 APA Pacific, the Central European University, and the Centre for Moral and Political Philosophy at the Hebrew University of Jerusalem. For commentary, feedback, or other support at these events, I thank Andre Martin, Arnon Levy, Bokai Yao, Christiane Merritt, Dan Baras, Daniel Telech, Dani Attas, David Enoch, Etienne Brown, Hichem Naar, Jens Gillessen, Jessica Isserow, Kathryn Lindeman, Preston Werner, and Said Saillant.

Thanks to my two doctoral supervisors, Iwao Hirose and Sarah Stroud. Both provided consistently quick and devastating feedback. The effect this had on my ability to finish this thing cannot be overstated. I was never waiting more than a couple of weeks: just enough

time to gain some much needed critical distance, but not so much time that I lost interest and motivation.

Iwao's supervision created a much cleaner and sharper thesis. His insistence on clarity and distaste for 'creative writing' meant I couldn't fudge an argument around the borders or bury the lede. He has not only been an excellent thesis advisor but also an advisor about academia more generally. He was imparted to me so much about how to present material, form abstracts for conferences, select journals, and so on that I would be lost without his guidance on all these matters.

Sarah Stroud has not only been involved in every aspect of this thesis, but in all aspects of my doctoral life at McGill. She was a supervisor on the candidacy paper that started this thesis. She knew what my thesis was about before I did. When I thought I was working in (and flailing at) moral metaphysics, she realized my interests were really in moral epistemology. An embarrassing amount of this thesis is the result of me giving tangled yarns to Sarah and her showing me how to weave them into something presentable. During so many meetings over the years I have jotted down her summaries and used them in abstracts and the chapters themselves. Her charitable interpretation of my ideas helped uncover arguments I would never have otherwise seen. Her dedication and patience as an advisor are unparalleled.

Lastly, every idea and argument in the dissertation was at some point bounced off of Martina Orlandi. I'm sorry. You bore the brunt of the half-baked ideas, and I'm sure I ruined many dinner conversations over the past few years. But your love, support, and encouragement have made the past years much easier than they should have been. Thank you. I hope I can do even half as good a job returning the favour.

CHAPTER 1
Introduction

> *North.* — What is the trouble about moral facts? When someone denies that there is an objective moral order, or asserts that ethical propositions are pseudo-propositions, cannot I refute him (rather as Moore refuted those who denied the existence of the external world) by saying "You know very well that Brown did wrong in beating his wife. You know very well that you ought to keep promises. You know very well that human affection is good and cruelty bad, that many actions are wrong and some are right"?
>
> *West.* — Isn't the trouble about moral facts another case of trouble about knowing, about learning?
>
> — P.F. Strawson, "Ethical Intuitionism" (1949), p.23

1.1 Types of Scepticism

This dissertation starts with a simple question: how do we know moral truths? How do we know that lying is wrong, that gender discrimination is bad, that torturing people merely for fun is immoral, and so on? My answer to this simple question is, unfortunately, negative: we do not know moral truths. That makes this dissertation a defense of moral scepticism.

This is of course not a new position. Moral scepticism is one of the longest standing positions in philosophy.[1] But I arrive at moral scepticism in a non-standard way. One way

[1] Walter Sinnott-Armstrong notes that one type of scepticism is "defined by the claim that nobody knows or is justified in believing anything" (Sinnott-Armstrong 2006: 10). This

of describing the novelty of my sceptical position is by identifying the ways in which we can fail to secure moral knowledge. There are I think at least two ways that we might come up empty handed: via an *ontological* route or else via an *epistemological* route. Many recent defenses of sceptical positions take the ontological route.[2] But I take the epistemological route to scepticism. To see the difference between these two routes we can turn to some non-moral cases.

Let's first consider Amy, who takes herself to be well-positioned to know about God. There isn't anything about Amy or her epistemic abilities that prevents her from being able to know things about God. But unbeknownst to Amy, God does not exist — so Amy will fail to know anything about God. Amy's failure here is not an internal problem with *her*, but rather an external problem with *reality*. She is certainly capable of knowing things about God, she has her epistemic wits about her, but since God does not exist it is ultimately reality that does not keep up its end of the deal. What we should say about Amy is that she does not in fact know anything about God, and the reason she does not is because God does not exist. This is the ontological route to scepticism.

Now consider Gary, who has an overwhelming but coherent and informed desire "to count blades of grass in various geometrically shaped areas such as park squares and well-trimmed lawns" (Rawls 1971: 432). Suppose that Gary wants to know how many blades of grass there are in Central Park at Noon tomorrow. There seems to be no ontological

---

is a broad type of scepticism that is thought to cover scepticism about moral knowledge or justified moral beliefs. Sinnott-Armstrong says this view "descends from Plato's Academy, so it is called *Academic* scepticism" (Ibid.). See also DePaul 2009 for a discussion of *Pyrrhonian* moral scepticism. However, of the ancient moral sceptics (e.g. Arcesilaus, Aenesidemus, Agrippa, Carneades, Pyrrho, Sextus Empiricus, etc.) it is noteworthy that "there is no skeptical school that would confine itself to 'moral' skepticism, or skepticism about values" (Vogt 2018; see also Vogt 2016).

[2] See for example Burgess 2007; Garner 2007; Joyce 2001, 2005; Joyce & Kirchin 2010; Machuca 2018; Olson 2011, 2014; Streumer 2017, etc.

problem here. Central Park exists, so the park and its many blades of grass are ontologically respectable in that sense. But even if reality keeps up its end of the deal, Gary might nevertheless be prevented from knowing how many blades of grass are in Central Park at Noon tomorrow. Maybe as Gary begins to count he finds that blades of grass are blowing away and throwing off the numbers. Maybe Gary is unable to complete the task as described — if he begins too early he won't be sure how many blades of grass are in the park at Noon, and if he begins too late he will have missed his target time. These are limitations that prevent him from knowing how many blades of grass there are. But, importantly, they are not ontological limitations. They are instead limitations regarding his abilities. This is the epistemological route to scepticism.

Whereas Amy is led to scepticism via the ontological route (because the universe does not contain the sort of thing needed for her to have knowledge), Gary is led to scepticism via the epistemological route (he isn't equipped to have knowledge about a thing that can in principle be known). These two ways of failing to have knowledge are mutually exclusive. On the one hand, if the failure is ontological then that means one could have knowledge if only the universe kept up its end of the deal. So Amy *could* have knowledge of God, if only God existed. On the other hand, if the failure is epistemological then it means that the universe does keep up its end of the deal but the potential knower is deficient or limited in some critical way. Gary *could* know how many blades of grass there are, if only he had different abilities.[3]

[3] There are further distinctions we can draw among purely epistemic failures. Aaron Zimmerman notes that there are also *internal* problems and *external* problems (Zimmerman 2010: 1). One might fail to have knowledge of some thing because of an internal problem, like poor vision. Or one might fail because of some external problem, like bad lighting. Both of these are epistemic problems, and not ontological problems. We are supposing here that the object in question *exists*, but noting that we can fail to know about it in two distinct ways.

In what follows my focus will not be on scepticism in general, but rather on *moral* scepticism in particular. And there is I think a standard way of characterizing the difference between ontological and epistemic failures in the moral domain. On the one hand there are *moral nihilists* or (relatedly) *moral error theorists*.[4] These are people who believe that "there are no moral truths to be known and thus no moral knowledge" to be had (Zimmerman 2010: 43). Moral nihilists are much like atheists: they say that we lack moral knowledge because there are no moral facts, just as atheists say we lack knowledge about God because God does not exist.

On the other hand there are *epistemic moral sceptics*, or what are sometimes just called *moral sceptics*. These are people who say that "whether or not there are moral truths, any evidence, reasons, or grounds we have for our moral beliefs must prove insufficient to provide us with moral knowledge or even justified moral belief" (Zimmerman 2010: 43). Epistemic moral sceptics are much like the aforementioned grass-counter: even if there are facts to be known, one is nevertheless not in a position to know them.

So explained, moral nihilism entails moral scepticism but not vice versa. This is because moral sceptics need not take a stand on whether there are in fact moral truths. Rather, they need only say that *if* there are such truths we either do not know or have justified beliefs about them. Conversely, moral nihilists say that we do not know or have justified beliefs about them *because* they do not exist. If something doesn't exist you cannot know anything about it — so moral nihilism entails moral scepticism. But, of course, your failure to know

[4] I think we should avoid drawing an equivalence between moral nihilists and moral error theorists. The latter are a specific type of moral nihilist. As we will see later, moral error theorists are *cognitivists* about judgments and *nonnaturalists* about purported moral properties. It is this combination that helps generate the error theory. But one need not think that moral judgments are cognitive, nor that moral properties are nonnatural in order to be a moral nihilist. Moral nihilists need only say that moral facts (however they are construed) do not exist.

about something does not entail that it does not exist — so moral scepticism does not entail moral nihilism.

1.2 Realism, Nihilism, and Scepticism

Clarifying the difference between moral nihilism and moral scepticism is important for the rest of the dissertation. I will not be arguing that there are no moral truths. Rather, I will be arguing that *even if there are* moral truths we are nevertheless not in a position to know or have justified beliefs about them. Because of that, much of what I will go on to say will not be of an ontological or metaphysical nature, but instead of an epistemological nature. And that is the 'trouble about moral facts', as I am understanding it: it is a problem about knowing anything about them.

Some of my defense of epistemic moral scepticism will be appealing to moral nihilists, since they too are by definition moral sceptics. But some of what I will go on to say will also be appealing to non-sceptics. In particular, so-called *moral realists* will share many of my background assumptions. As it is usually understood, moral realism is the view that there *are* moral truths.[5] So moral realism is usually taken to be an ontological view. When this ontological view is coupled with the epistemic claim that some moral truths *are known* we get *non-sceptical moral realism.* Because epistemic moral scepticism has some overlap with non-sceptical moral realism, and because epistemic moral scepticism has some overlap with moral nihilism, these three positions can be read as sharing certain assumptions. I describe some of these similarities here.

First, look at how realists, nihilists, and sceptics each relate to the following two claims:

[5] For example, moral realism is usually understood as a view that posits the existence of "moral facts and truths that are objective in some way" (Brink 1989: 14). Similarly, moral realists are usually described as those who think that "moral claims do purport to report facts and are true if they get the facts right. Moreover, they hold, at least some moral claims actually are true" (Sayre-McCord 2001: 1).

Truth-apt Our moral judgments are capable of being true or false

Independence The truthmakers of our moral judgments are mind- and attitude- independent

While it ultimately depends on the details of the particular metaethical position (i.e. the particular version of realism, nihilism, etc.), in principle the three views can all accept both **Truth-Apt** and **Independence**. For example, the non-sceptical moral realist can say that a moral judgment like 'lying is wrong' can be true or false, and that what makes that judgment true or false is independent of our stance (i.e. our attitudes or thoughts) towards it. And both the moral nihilist and the moral sceptic can agree with that. The reason the three views are united here is because **Truth-apt** is a fairly weak claim. All it says is that our moral judgments are *capable* of being true or false, but that alone does not settle the question of whether they are *true*, whether they are *false*, or whether we *do not know*.

Similarly, the three views are also united by a commitment to **Independence**. Realists, nihilists, and sceptics can all agree that whether or not our moral judgments *are* true or *are* false is independent of our individual or collective stance towards the content of those judgments. This makes the three views different than so-called *subjectivist* or *constructivist* accounts of morality. On these latter views, whether or not a moral judgment is true is a function of our attitudes or thoughts.[6] But such views are not best understood as realist, nihilist, or sceptical positions. They are instead better understood as a separate category entirely. For that reason, I will set them aside in what follows.[7]

[6] For constructivist views in particular, see Bagnoli 2013, 2017; Ferrero 2009; Korsgaard 2008, 2009; O'Neill 1989; Street 2008; etc.

[7] As will become clear in the next chapter, the reason for excluding subjectivist or constructivist accounts is because they do not appear to adequately characterize our extant moral discourse and practice. In particular, they reject *nonnaturalism* about moral properties. So they are better understood as revisions on our discourse and practice, rather than descriptions of it.

With **Truth-apt** and **Independence** in mind, realists, nihilists, and sceptics necessarily differ with respect to what follows. Moral realists typically argue that in addition to those two claims we should accept the following:

Realism At least some of the truthmakers of our moral judgments obtain.

This claim is what makes the non-sceptical moral realist a *realist*. So the combination of **Truth-apt**, **Independence**, and **Realism** gets us moral realism. And as we will see in Chapter 2, moral realists will differ with respect the particularities of those truthmakers (are they natural? nonnatural? supernatural? constructed? etc.).

On the other hand, moral nihilists say that we should reject **Realism** and instead accept the following claim:

Nihilism No truthmakers of our moral judgments obtain.

Nihilism is simply the denial of **Realism**. It is primarily an ontological claim. So settling this particular debate between realists and nihilists would require us to wade into ontological and metaphysical debates. But, as I said, I will not do that here.

Instead, I will focus mostly on epistemic matters. And it is here that the moral sceptic disagrees with both the moral realist and the moral nihilist. What makes the moral sceptic's position unique is that it does not turn on an ontological claim — i.e. whether certain truthmakers for our moral judgments obtain or not. Instead of **Realism** or **Nihilism**, the moral sceptic says that we should accept the following claim:

Scepticism We don't know whether any truthmakers of our moral judgments obtain.

The moral sceptic can remain noncommittal regarding the existence of moral facts. Instead, moral sceptics need only say that *even if* there are moral facts we do not have moral knowledge. This is why it is common to distinguish between two questions: "*Assuming* that we have moral knowledge, *how* do we have it?" and "Do we in fact have any moral knowledge?" (McGrath 2004: 209). Both the epistemic moral sceptic and the non-sceptical moral realist

can weigh in on the first question, and they might agree on the answer. But the two will by definition disagree with respect to the second question: the realist affirms it while the sceptic denies it. As said, the problem of knowing anything about moral facts, of learning anything about them, is the trouble about moral facts.

1.3 Layout

In Chapter 2 I will explain two claims about our moral discourse and practice. We should accept *cognitivism* about our moral judgments, and *nonnaturalism* about moral properties. This sets important desiderata for any plausible moral epistemology. If we know moral truths then it is because we have successful judgments about nonnatural moral facts.[8] I will call this view *nonnaturalist cognitivist realism*.[9] What we are after is a plausible moral epistemology for this view. While some think that nonnaturalist cognitivist realism is vindicated if there *are* moral facts (an ontological criterion of success), I argue that we should only take the view to be vindicated if it can explain how we *know* those facts (an epistemological criterion of success).

In Chapter 3 I address the evolutionary debunking argument (EDA). The EDA threatens to undermine the possibility of moral knowledge by appealing to our evolutionary history. If the EDA is successful in its ambitious debunking aim there can be no plausible moral epistemology at all, let alone one for nonnaturalist cognitivist realism. But I argue that the EDA is not a justification debunking argument (undermining any way of justifying our

[8] As we will see, because our search for a plausible moral epistemology is constrained by cognitivism and nonnaturalism, some notable moral epistemologies are off the table. For example, I will not discuss pragmatist or inferentialist accounts, but see Brandom 2000; Chrisman 2010, 2016; Heney 2016; Lekan 2003, 2006; Sepielli 2016; Tiefensee 2016; etc. Neither the pragmatist or inferentialist account is amenable to nonnaturalism, which is why I set them aside.

[9] This is the sort of metaethical view defended by Cuneo 2007, 2008; Enoch 2007, 2011; Fine 2002; FitzPatrick 2008; Hampton 1998; Huemer 2005; Parfit 2011, 2017; Shafer-Landau 2003; Wedgwood 2007; etc.

moral judgments) but instead a theory debunking argument (undermining a particular way of justifying our moral judgments). In particular, the EDA undermines those metaethical theories that are guilty of inferring theoretical beliefs merely from folk beliefs. For this reason I recharacterize the EDA as an *inference debunking* argument. I show that nonnaturalist cognitivist realism need not fall prey to this inference debunking argument, and that leaves intact the hope that we can justify our moral judgments.

I begin canvassing plausible moral epistemologies in Chapter 4, beginning with moral intuitionism. Moral intuitionists promise an epistemology that bridges the gap between our moral beliefs and the nonnatural moral facts. It does this by positing a foundationalist account of moral knowledge. They say that we can intuit certain moral propositions because they are self-evidently justified. This has notoriously led to charges of dogmatism, and with it charges of implausibility as a moral epistemology. I show that 'new intuitionists' fix the dogmatism of 'old intuitionism' by embracing so-called *epistemological over-determination*: a moral proposition that can be justified self-evidently need not *only* be justified self-evidently. But I show that this embrace comes at a significant cost, causing intuitionism to lose it's status as a foundationalist moral epistemology. Intuitionism thus faces a dilemma: without committing to epistemological over-determination the view cannot explain self-evident moral propositions non-dogmatically, but accepting epistemological over-determination causes intuitionism to lose the much needed claim to foundationalism. Without resolving this dilemma, moral intuitionism cannot explain how we know moral truths.

In Chapter 5 I turn to moral intuitionism's most common alternative: reflective equilibrium. This is often said to be the default moral epistemology. Despite it's prominence, I show that reflective equilibrium is poorly understood. It's defenders say that the view has an essentially coherentist structure that explains moral knowledge in a way that doesn't lead to dogmatism. But this opens the view up to charges that it cannot explain how our judgments cotton on to nonnatural moral facts. I show that a recent defense of reflective equilibrium attempts to solve this problem. But unpacking the structure of this defense reveals a serious

flaw. At a critical juncture, reflective equilibrium is question-begging: it relies on reflective equilibrium *itself* to show why reflective equilibrium is a justified moral epistemology. Without resolving this problem, reflective equilibrium (like moral intuitionism) is unable to explain how we know moral truths.

With the two most prominent attempts to explain moral knowledge found wanting, in Chapter 6 I focus on a little discussed alternative to standard moral epistemologies: metaethical Mooreanism. Mooreans say that we can dispense with all sceptical challenges to moral knowledge because we know certain moral claims *better than* we know any premise, inference, or conclusion of any sceptical argument. I argue that even if we should accept Moorean arguments elsewhere, we should not accept them in metaethics. I show that a key premise in Moorean *arguments* (called Moorean *claims*) initially enjoys justification in virtue of *presumptive* reasoning (i.e. taking appearances at face value), but that this justification can be and is defeated by that stemming from *dialogical* reasoning (i.e. responses to arguments). In particular, the inference debunking argument I developed in Chapter 3 provides a potent defeater for the presumptive reasoning in favour of Moorean claims. In light of this, I conclude that we do not know certain moral claims better than we do certain sceptical challenges. Thus, metaethical Mooreanism does not allow us to claim that we have moral knowledge.

Because we lack a plausible moral epistemology for nonnaturalist cognitivist realism, I conclude that we lack moral knowledge. I then end the dissertation by showing why, even if we lack a plausible moral epistemology, and thus should accept epistemic moral scepticism, that scepticism need not generalize to non-moral epistemic theories more broadly. In doing so, I offer a response to so-called 'companion in guilt (or innocence) arguments' that attempt to unite the moral and epistemic domains. I argue that the search for a plausible moral epistemology is constrained by the need to vindicate nonnaturalist cognitivist realism, but the epistemic domain is not best understood as beholden to nonnatural epistemic facts. For this reason, the scepticism defended in the dissertation does not generalize.

CHAPTER 2
Surveying the Metaethical Landscape

I will begin from the assumption that nonnaturalist cognitivist realism enjoys a status as *the view to beat* in metaethics.[1] Since I will use the current chapter to say something in defense of that claim, we can call it a 'motivated assumption'. In the rest of the dissertation I will show why nonnaturalist cognitivist realism fails on epistemic grounds. It is unclear precisely where that leaves us, but I suggest that epistemic moral scepticism looms large. This makes the assumption regarding nonnaturalist cognitivist realism particularly important because it shows why, if the view fails epistemically, there is a significant loss. The failure to find adequate epistemic footing for nonnaturalist cognitivist realism would be of little consequence if the view did not enjoy much plausibility in the first place. But we will see that, owing to its structural features, if it *is* the view to beat *and* it suffers epistemically then we should accept epistemic moral scepticism.

In order to understand why nonnaturalist cognitivist realism is the view to beat here we will need to survey the metaethical landscape. This chapter does just that. This will consist in a guided tour through the metaethical terrain, and along the way I show what I take to be convincing extant arguments against certain positions and in favour of others. The culmination will be an admittedly partial defense of the assumption that nonnaturalist cognitivist realism enjoys a default status in metaethics. It is the metaethical view that other views necessarily depart from.

[1] Nonnaturalist cognitivist realism is the sort of metaethical view defended by Cuneo 2007; Enoch 2007, 2011; Fine 2002; FitzPatrick 2008; Hampton 1998; Huemer 2005; Parfit 2011, 2017; Shafer-Landau 2003; Wedgwood 2007; etc. And I am taking the locution 'the view to beat' in metaethics from Enoch 2017.

In section 2.1 I explain what I take to be the point of offering a metaethical theory in the first place: we are attempting to explain our moral judgments along semantic, psychological, ontological/metaphysical, and epistemic lines. I then divide the metaethical landscape using three cuts: along semantic and psychological lines with a distinction between cognitive and noncognitive judgments, along ontological/metaphysical lines with a distinction between nonnatural and natural properties, and along epistemological lines with a distinction between an epistemic criterion of successful judgments and an ontological criterion of successful judgments. While this does not exhaust all extant metaethical options, it will provide considerable background for the reader to understand the location of this thesis within recent metaethical debates. In section 2.2 I explain the difference between cognitivism and noncognitivism regarding moral judgments, and I show why we should prefer cognitivism. In section 2.3 I explain the difference between nonnaturalism and naturalism regarding moral properties, and I show why we should prefer nonnaturalism. And in section 2.4 I introduce the notion that nonnaturalist cognitivist theories require an epistemological criterion of success, rather than the standard ontological criterion. The remainder of the dissertation can then be understood as an extended argument for why we cannot accept that our moral judgments succeed in that way.

2.1 What is Metaethics?

Metaethics is the area of philosophy that attempts to characterize our moral discourse and practice. As a distinguishable area of moral philosophy, metaethics typically contrasts against both normative ethics and applied ethics, where normative ethics deals with the structure of first-order moral theories (think debates between consequentialists and deontologists), and applied ethics with specific moral issues (think the rightness or wrongness of euthanasia). Metaethicists attempt to explain normative ethics itself, and they do this by uncovering what, if anything, we are committed to in our first-order moral thought, discourse, and practice.

We can uncover commitments within four key areas: within our moral language, within our moral psychology, within our moral metaphysics or ontology, and within our moral epistemology. If we are interested in metaethics then we are interested in analysing first-order morality along any or all of these lines.

One question we might have at the outset is this: who is the 'our' in 'our moral discourse and practice', or the 'our' in 'our moral language', and so on? Philosophers?[2] WEIRD people?[3] For the most part, metaethicists do not make this explicit. To remedy this, we can note that there are two senses here. At one extreme, we can say that there is an *absolute* sense of 'our' which refers to all humankind. This sense is intended to capture the moral discourse and practice, or the moral language, psychology, etc., of "everyone, always and everywhere" (Lewis 1989: 126). And at the other extreme, we can say that there is an *indexical* or *relative* sense of 'our' which refers to "a population consisting of the speaker and those somehow like [him or her]" (Ibid. 127).

The sense of 'our' that I have in mind is *absolute*. I intend for my arguments here to cover the moral discourse and practice of everyone, always, and everywhere. This is because I take it that there is a general yet essential feature of our moral discourse and practice, and that characterizing this general feature in a particular way shows that everyone is already committed to understanding morality along similar lines.

So what is this essential feature of our (in the *absolute* sense) moral discourse and practice? I take it that the key feature is this: we make and have moral judgments. This is something that I think we all do. The way I am using it here, moral judgments are evaluations

[2] It may be interesting to examine whether and how philosophers and non-philosophers differ in their understanding of first-order moral theorizing. See Bourget & Chalmers (2014) for a meta-survey of philosophers metaethical views, and see Goodwin & Darley (2008) for the metaethical views of non-philosophers.

[3] People from *western*, *educated*, *industrialized*, *rich*, and *democratic* societies. See Henrich et al (2010).

of things as morally good or bad, right or wrong, better or worse, and so on.[4] I use 'things' as a shorthand for actions or inactions, persons or groups, motives or consequences, and any and all such things that we deem it appropriate to evaluate in specifically moral terms. I think that we all make judgments of this sort. And if you agree then we are on our merry way towards the absolute sense of 'our'. The goal of metaethics, as I understand it, is to explain this aspect of our everyday lives.

These judgments can be understood in two distinct ways: as *linguistic* utterances and as *psychological* states. When someone *makes* or *has* a moral judgment, they are either *making* an utterance, or they are *having* a propositional attitude. In what follows, I will understand moral judgments as covering both utterances and propositional attitudes. Examples of moral judgments are linguistic claims or psychological attitudes like "torturing people merely for fun is wrong", or "gender discrimination is bad", and so on. These are examples of the sorts of moral judgments that metaethicists attempt to understand along linguistic, psychological, metaphysical or ontological, and epistemological dimensions.

But even if they are an essential feature of our moral discourse and practice, there is a lot of disagreement about how best to understand moral judgments, and moreover which sort of metaethical theory best accounts for those judgments. I cannot canvass all such options here.[5] Rather, I will partition the landscape using three cuts: cognitivism versus noncognitivism, nonnaturalism versus naturalism, and epistemic success versus ontological success.

[4] I also take it that our moral judgments often take the form of so-called *thick evaluations*, like 'so-and-so or such-and-such is courageous', or '. . . just', or '. . . honest', or '. . . despicable', or '. . . cruel', and so on. But I will not refer to these sorts of evaluations in the remainder of the thesis. See Kirchin 2017.

[5] See instead McPherson & Plunkett (2017) for a far more comprehensive overview.

2.2 First Cut: Cognitivism versus Noncognitivism

To begin carving up the metaethical landscape, the first cut we can make will divide the plane in two. On one side, there will be *moral cognitivists*, and on the other side, there will be *moral noncognitivists*. Since we are understanding 'moral judgments' as covering both utterances and propositional attitudes, we can understand these views as characterizing judgments in both a *linguistic sense* and a *psychological sense*. Moral cognitivism (hereafter just *cognitivism*) is the view that moral judgments, in the *psychological sense*, are particular cognitive states, like beliefs, and it is the view that moral judgments, in the *linguistic sense*, are truth-apt statements. And moral noncognitivism (hereafter just *noncognitivism*) is the view that moral judgments, in the *psychological sense*, are noncognitive attitudes, like approval or disapproval, and it is the view that moral judgments, in the *linguistic sense*, are non-truth-apt statements. I will get clearer on the details of these two positions and then show why we should accept cognitivism rather than noncognitivism.

But first a quick word on terminology. Some might object to the term *noncognitivism*, preferring instead a term like *expressivism*, *quasi-realism*, *prescriptivism*, or some other *-ism*, since these latter terms more appropriately describe a given metaethical theory.[6] Nevertheless, I eschew the latter terms in favour of the former. This is because I agree that noncognitivism is better

> understood as a purely negative doctrine or set of doctrines, for this way of understanding noncognitivism allows us to understand how there can be expressivist and prescriptivist species of the genus noncognitivism. The expressivist

[6] Others might object to the term noncognitivism because they think that purportedly noncognitive judgments are in an important sense *cognitive*. If you have this worry then take *cognitivism* and *noncognitivism* as imperfect labels that describe two (apparently) different types of judgments.

> and the prescriptivist should be seen as adding their own distinctive positive accounts of the meanings of moral language to the purely negative doctrine(s) of noncognitivism (Ridge 2013: 3620).

It is in this sense that I use the term noncognitivism, in order to identify at the genus level those metaethical theories that position themselves against cognitivism (where cognitivism is understood as *psychologically representationalist* and *semantically factualist* — to be explained below). While noncognitivism is itself an inadequate term for distinguishing between expressivists (e.g. Chrisman 2008; Gibbard 1990, 2003; Price 2011), quasi-realists (e.g. Blackburn 1984, 1998), or prescriptivists (e.g. Carnap 1935; Hare 1952), my use of the term is intended to capture the similarities between these theories and the collective sense in which they reject the characteristic features of cognitivist metaethical theories.

It is commonplace to note that cognitivism and noncognitivism can be understood as either semantic views concerning moral language or as psychological views about moral evaluation.[7] When they are construed as semantic theses, the two views diverge over how we should understand the moral terms operative within our moral discourse.

According to cognitivism, the moral terms that we use in utterances like "torturing people merely for fun is wrong" or "gender discrimination is bad" are *predicates*. If the moral terms here ('wrong' and 'bad') are predicates then using these terms is our way of describing a particular action, person, or target of moral evaluation. When we use a moral predicate like "...is wrong" or "...is bad", a natural way to read and understand the utterances containing them is that those utterances are capable of being true or false. In a somewhat stilted way, we can say that the action 'torturing people merely for fun' satisfies the predicate "...is wrong" just in case that action is *in fact* wrong, or that 'gender discrimination' satisfies the predicate "...is bad" just in case it is *in fact* bad. More colloquially, however, we would

[7] See for example Svavarsdóttir 2001 and Van Roojen 2014.

just say that that 'torturing people merely for fun is wrong' is a true statement when *in fact* torturing people merely for fun is wrong (*mutatis mutandis*, 'gender discrimination is bad').

So cognitivism is the view that moral judgments rely on moral terms that operate as predicates, and thus our moral judgments are to be understood as purportedly factual statements. When we say that "gender discrimination is bad" we are making a statement of fact claiming that (or to the effect that) 'gender discrimination' satisfies the predicate "... is bad". Because this is a semantic thesis that says moral judgments are statements of fact, we can call this commitment of cognitivism *semantic factualism* (Van Roojen 2014).

In addition to this semantic thesis cognitivists accept a particular psychological thesis. Because cognitivism is semantically factualist, we can see that a sincere moral judgment like "gender discrimination is bad" appears to indicate one's *belief* that gender discrimination is *in fact* bad. This is because, as we saw, one is making clear that he or she *believes* that gender discrimination satisfies the predicate "... is bad". In this way, there is agreement between cognitivism's semantic and psychological theses: understanding moral language according to semantic factualism reveals that a sincere moral judgment concerns one's *beliefs*.

Another way to characterize the psychological claim here between is in terms of so-called *direction of fit*.[8] The direction of fit for beliefs is said to be *mind-to-world*. Our beliefs aim at accurately capturing or matching or representing (any of these locutions will do) the way things actually are independent of us.[9] A belief that 'gender discrimination is bad' aims to capture or match or represent the world such as it actually is (i.e. as the world being such that gender discrimination is in fact bad). In this way, we can say that our beliefs are *representational*. Cognitivism, then, is standardly characterized as a *representationalist*

[8] See for example Anscombe 1957 and Searle 1984.

[9] Note how this allows us to readily distinguish a *belief* from a *hope* or *wish*, since hopes and wishes are not representations of the world so much as they are desires that the world be a certain way.

view, because it says that moral judgments are attempts at representing the world. This follows if our moral judgments are understood as beliefs, because beliefs have a mind-to-world direction of fit.[10]

So cognitivism is the semantic view that moral judgments are statements of fact (i.e. semantic factualism), and it is the psychological view that moral judgments are beliefs (i.e. representationalism). But even if this does seem to be a natural gloss on moral judgments, both of these theses are controversial. Complaints mostly issue from noncognitivists, who deny both representationalism and semantic factualism. Moreover, it is because noncogntivists reject representationalism that they are led to reject semantic factualism.

Noncognitivists reject that moral judgments are beliefs, and it is important to see why. As we saw, cognitivists say that moral judgments are beliefs and so they are committed to representationalism. Importantly, they think that representationalism naturally follows from semantic factualism: i.e. because we have *this claim* about moral language, it implies *that claim* about moral beliefs. But noncognitivists reject that moral judgments are beliefs, and that means they can reject representationalism. We saw that beliefs have a mind-to-world direction of fit, which shows why cognitivists are committed to representationalism. But if the relevant attitude for moral judgments is not *belief* but instead *something else* then we might not accept representationalism.

Though cognitivists move from a view about moral language to a view about moral psychology, noncognitivists appear to move in the opposite direction: from a view about moral psychology to a view about moral language. Noncognitivists say that moral judgments

[10] Andrew Sepielli says that metaethical views can be characterized according to their stance on *the representation claim*, which is the claim that "there are metaethically possible conditions under which something thinking a thought like [A is morally wrong] accurately represents the world in so doing" (Sepielli 2015: 266).

are non-doxastic attitudes, like approval or disapproval (and we will see why below).[11] If they are attitudes of this sort, then representationalism can easily be rejected. Beliefs, on the one hand, and attitudes approval or disapproval, on the other, have different directions of fit.

While the direction of fit for beliefs is said to be mind-to-world, the direction of fit for noncognitive attitudes like approval or disapproval is reversed. Approval or disapproval have a direction of fit that is said to be *world-to-mind*. These attitudes are for the world to be a certain way. We can see this direction of fit more clearly by first looking at *desire*. Desires are a prime example of an attitude that is world-to-mind because they are attitudes towards the world being a certain way. When one *desires* or *wants* one is not attempting to *represent* the world being a certain way, but rather one has a non-representational attitude that the world be a certain way (usually, as as allowing that the desire or want is satisfied). If we understand how desires are world-to-mind then we can generalize to other noncognitive attitudes as well. Disapproval of 'gender discrimination', for example, does not aim to capture or match or represent the world as being a certain way, but rather that attitude indicates that we wish or desire that the world be a certain way (usually, as such that 'gender discrimination' doesn't take place). In this way, noncognitive attitudes like approval or disapproval are often described as *non-representational*. Noncognitivism, then, is *non-representationalist*, because it says that moral judgments are not representations but rather attitudes towards the world.

From this view about moral psychology, noncognitivists derive a particular view about moral language. According to noncognitivists, the moral terms in the above utterances are not actually predicates at all, but rather are indicators of attitude. The moral utterance

[11] That said, I think it is important to remain neutral about the exact nature of these attitudes. Because I am understanding noncognitivism as a negative doctrine, the exact nature of these attitudes will allow for distinct positive noncognitive positions (e.g. expressivist, prescriptivist, etc.).

"torturing people merely for fun is wrong" might look and sound like it involves predication, but according to noncognitivism the moral term 'wrong' is actually an indication of a noncognitive attitude like approval or disapproval. When we utter that sentence, we should be understood as indicating that the action 'torturing people merely for fun' is in some relevant sense *disapproved* of by the speaker, or that the speaker *disapproves and urges that you do so as well*, or some similar characterization that turns on a propositional attitude other than belief.

Noncognitivism therefore denies that moral judgments are statements of fact. When we say that "gender discrimination is bad" we are not saying that the action 'gender discrimination' satisfies the moral predicate "...is bad", or that the sentence is true if *in fact* gender discrimination is bad (and 'not bad' otherwise). Rather, these sentences are not capable or being true or false at all, because the moral terms within them are not actually predicates. Instead, the moral term "...is wrong" indicates a particular attitude that we have towards the action in question. We are indicating our attitude of disapproval (rather than our attitude of belief) regarding gender discrimination. And because this is a semantic thesis that says moral judgments are not statements of fact, we can call this commitment of noncognitivism *semantic non-factualism* (Van Roojen 2014).[12]

We can now put these pieces together. Cognitivism and noncognitivism are competing and incompatible combinations of moral semantic and moral psychological theses. On the one hand, cognitivism is the view that combines *semantic factualism* with *representationalism*, and on the other hand noncognitivism is the view that combines *non-representationalism* with *semantic non-factualism*. Moreover, semantic factualism implies representationalism,

[12] I am aware that there are sophisticated versions of noncognitivism that make use of *minimialist* or *deflationist* accounts of truth. Because of this, some versions of noncognitivism might be able to incorporate more appearances or features of semantic factualism than others. But since noncognitivism is not my primary focus in what follows, I will not deal with these interesting accounts here. See Camp 2017 for an overview of these versions of noncognitivism that incorporate a deflationist or minimialist sense of truth.

because beliefs are representational or have a mind-to-world direction of fit. And *non-representationalism* implies semantic non-factualism because moods or approval or disapproval are non-representational or have a world-to-mind direction of fit.

2.2.1 Argument for Cognitivism

With the two views on the table, let us look at whether we should accept one over the other. There are readily available arguments in favour of each position, which I will quickly present before showing why we should accept cognitivism over noncognitivism.

The primary argument for cognitivism is what we can call *the easy argument.* The easy argument says that our moral discourse and practice certainly appear representationalist, in the same way as other areas of discourse that are paradigmatically representationalist (think about an utterance like "snow is white"). Because of this, we are licensed to think that our moral discourse and practice are also representationalist. Again, representationalism is the claim that an area of discourse is attempting to capture a reality that is in some important sense independent of our thought and speech about it. If an area of discourse is representationalist, then thought and speech can more or less accurately capture, match, or represent the reality that the discourse is about. As we saw above, typical moral judgments like "torturing people merely for fun is wrong" and "gender discrimination is bad" look and sound like they involve our predicating wrongness and badness of those actions and events. Cognitivists think that we should take the representationalist appearance of our moral language at face value. This is what makes the easy argument *easy.* The view reads off of our moral language and concludes that when we say that "gender discrimination is bad" what we really mean is that we believe (stiltedly) that 'gender discrimination satisfies' the predicate "...is bad" or (colloquially) that 'gender discrimination is bad' is true just in case gender discrimination is *in fact* bad. This is similar to other paradigmatically representationalist utterances, like "snow is white". When we say "snow is white" the most natural way to understand it is as our indicating a sincere belief that 'snow' satisfies the predicate "...is

white", or that our sentence 'snow is white' is true just in case snow is *in fact* white. Cognitivists urge us to accept this representationalist appearance of our moral language, and the easy argument works by pointing to the similarities that moral language has with other representationalist language.

Of course, as we saw above, noncognitivists deny that our moral discourse and practice is actually representationalist. So the easy argument has no truck with them. Noncognitivists suggest that we have been misled by the syntactic structure of our moral language, that we have fallen "into thinking that a [moral] statement is really an assertive proposition ... but actually a [moral] statement is nothing else than a command in a misleading grammatical form" (Carnap 1935: 24). Though it can look and sound like moral terms are in fact predicates, and that we are thereby suggesting that certain things satisfy those predicates, this appearance is misleading. Our moral terms are something else entirely: they are indications of approval or disapproval, or indications of approval or disapproval with a recommendation or command that you do so as well (the latter of course being different species of noncognitivism).

But why think that we are so misled about our own moral utterances? It can be difficult to understand the noncognitivist position here. This is because noncognitivists generally do not offer an explicit explanation for how it is that we have fallen into confusion regarding our moral language. That is, they do not offer any sort of 'debunking' account for the *prima facie* representationalism of moral language.[13] But this should be surprising. After all, noncognitivists do not deny that moral language *appears* explicable in terms of semantic factualism and psychological representationalism. Rather, they will insist that this appearance, however intuitive or initially compelling, is ultimately illusory.

[13] A debunking theory would explain how it is that we have come to have beliefs that are mistaken. This is importantly different than merely offering a competing account. This is roughly the difference between an *undermining* defeater and a *rebutting* defeater. For more on this latter difference, see Pollock (1986).

Because they do not provide a debunking theory, noncognitivists must pursue an alternate route. Rather than explain our moral language, they offer a competing account of our moral judgments themselves. While the primary argument for cognitivism is *the easy argument*, the primary argument for noncognitivism is what we can call *the motivation argument*. The motivation argument implies that the easy argument for cognitivism is easy because it focuses on the wrong feature of our moral judgments. Understanding this move will allow us to see why noncognitivists think that we are deeply mistaken about moral language. Look again at a typical moral judgment like "torturing people merely for fun is wrong". What is involved in our issuing this judgment? Cognitivists tell us that what is involved is that we are predicating *wrongness* of 'torturing people for fun'. But noncognitivists direct our focus somewhere else. They ask us to think about a person who might *make* that judgment. More to the point, think of the way in which a person who makes that judgment will *act*. Noncognitivists suggest that if we want to understand the particularities of *making* a moral judgment, then we need to appreciate just how tightly our moral judgments are connected to our *actions*.

For example, if Caligula readily makes the moral judgment that "torturing people merely for fun is wrong", and yet he also readily tortures people merely for fun, then that seems to be grounds for thinking that Caligula's moral judgment is insincere. Maybe we could explain away this discrepancy if there were some sort of extenuating circumstance compelling Caligula to act in such a way. But if he were not at least *pro tanto* motivated to avoid torturing people for fun, then this seems to be sufficient grounds for thinking that Caligula is not really judging torture as wrong. Noncognitivists urge us to focus on *this* aspect of our moral judgments, on the tight connection between making a moral judgment and being motivated to act in a correspondingly appropriate way, rather than on certain syntactic features of moral language.

And once we appreciate the way in which moral judgments are bound up with motivation, noncognitivists say that we should give up on the representationalist appearance of

moral language. Recall that cognitivists say that our moral judgments are beliefs, and beliefs aim at representing the world. But if our moral judgments are beliefs then how can we make sense of the tight connection between moral judgment and motivation? It is widely held that beliefs alone are unable to account for motivation.[14] For this reason, the representationalist appearance of our moral language must be simply that: a 'mere' appearance. Thus, noncognitivists think that their view (which accounts for judgments in terms of semantic non-factualism and non-representationalism) better accounts for our moral judgments.

With those arguments on the table, what I want to do now is show why we should accept cognitivism. Here is a sketch of the basic argument.[15] Noncognitivism is a negative thesis that says moral judgments are not cognitive judgments, so when we utter moral sentences we are not actually asserting anything. Because noncognitivism is supposed to be a *description* of our moral discourse and practice (rather than a *revision* of it), this entails the following: when we utter moral sentences we are not *intending* to assert anything.[16] Rather, noncognitivists say that our moral discourse and practice consists in something else: expressions (per expressivism), or prescriptions (per prescriptivism), or some other noncognitive attitudes. But here is the objection: it is not true that we do not intend to assert anything with our moral judgments. So noncognitivism is false, and we should accept cognitivism.

---

[14] This is the 'Humean theory of motivation'. Michael Smith says that "the Humean theory is a dogma in philosophical psychology", and moreover that it "is the expression of a simple but important truth about the nature of motivating reasons" (Smith 1987: 37). More recently, Terence Cuneo notes that "it is widely accepted by philosophers that a broadly Humean account of moral motivation is correct" (Cuneo 2002: 465).

[15] Because this first chapter is an overview, the view to follow comes largely from Cuneo (2006).

[16] If we *were* intending to assert something, then our moral discourse and practice would be in error - but noncognitivists do not think that it is. That is what distinguishes noncognitivists from moral error theorists.

Let us unpack this argument. Set aside moral language for a moment. In understanding language in general, it is standard to make a distinction between *locutionary acts* and *illocutionary acts* (Austin 1962). On the one hand, *locutionary* acts consist in "'saying something' in the full normal sense" (Ibid. 94). For example, we can understand the utterance "snow is white" in the 'full normal sense' by looking at the construction of the utterance itself, understanding its typical meaning, and so on. It is in English, it expresses a proposition, it uses a predicate to describe an object, and so on. So all utterances are locutionary acts and can be understood in the full normal sense.

On the other hand, *illocutionary* acts are locutionary acts with a particular force or guise. A locutionary act (i.e. a normal sentence construction) can have different forces: some have an *interrogative* force, others have *assertoric* force, still others have *expressive* force. Some utterances express desires, some express promises, some are commands.[17] How do we know which illocutionary force a particular locutionary act has? The illocutionary force varies depending in part on what we are *trying to do* with our locutionary act. When we say that "snow is white", the illocutionary act here depends on the context, the speaker's intention, what the speaker thinks the listener will understand, and an innumerable list of other factors. It may be the this locutionary act has assertoric force, if we are simply asserting that snow *is in fact* white. It may also be that this locutionary act has some other illocutionary force. Perhaps we utter that sentence in response to some claim that the speaker takes to be insightful but which we find completely obvious or vacuous (more like "duh, and snow is white"). In this latter case, the illocutionary force of "snow is white" is not merely assertoric, but rather it is expressive.[18]

[17] Illocutionary acts also include "making statements, asking questions, giving orders, making promises, apologizing, thanking, and so on" (Searle 1975: 319).

[18] This is similar to the following insight from Wittgenstein (1958, sec. 23): "But how many kinds of sentences are there? Say assertion, question, and command? There are *countless* kinds: countless different kinds of use of what we call 'symbols', 'words', 'sentences'. And

If this is the way in which we typically understand language, then we can turn back to moral language. We can now describe our moral utterances using the locutionary/illocutionary distinction. We should now be able to appreciate that all moral utterances are, at base, *locutionary* acts, simply because they are sentences in the full normal sense. They have a readily identifiable construction, obey the typical rules of language, and so on. But these locutionary acts can also take on a particular *illocutionary* force. Some of our moral utterances are assertions, some are expressions, some are promises, etc. Like all language, the illocutionary force will vary based on context, intention, and so on. But *which* illocutionary force a particular locution has depends in part upon what we intending to do with those words.

It is typical to note that "performing an illocutionary act is something that [one] does deliberately or intentionally" (Cuneo 2006: 40). This is important to highlight for present purposes. While we can accidentally *utter* a sentence (e.g. in our sleep), it does not seem as if we can accidentally *assert* something, for example. So while locutionary acts can be accidental or non-deliberate, illocutionary acts cannot. Whether or not an utterance of ours qualifies as an assertion or expression or whatever depends on what we are *trying to do* with our utterance.

But stressing the intentional element of illocutionary acts reveals a deep problem for noncognitivism. According to noncognitivists, when we make moral utterances we are not *intending* to assert anything. Rather than looking at what we intend to do, noncognitivists focus instead on what we *in fact* do. And what they say is that we are in fact indicating some noncognitive attitude, like approval or disapproval. On that picture, what we are *intending* to do with our moral language more or less drops out of the picture, and with it goes the

---

this multiplicity is not something fixed, given once for all; but new types of language, new language-games, as we may say, come into existence, and others become obsolete and get forgotten."

claim that we are intending to *assert* something. But here's the problem: it is not true that we do not intend to assert anything with our moral utterances. Here is why.

Terence Cuneo draws our attention to a particular type of person: the traditional religious believer. Unlike moral sceptics, amoralists, or other such characters, "the traditional religious believer is not a philosopher's fiction" (Cuneo 2006: 62). Traditional religious believers are moral realists. That is, they believe that there are moral facts. Some of these moral facts are divine and about God — for example, that *God is just*. And some of these moral facts are ordinary and concern human beings — for example, that *torturing people merely for fun is wrong*. Because there is no doubt about whether or not traditional religious believers exist (as there might be for genuine moral sceptics or amoralists), Cuneo suggests that we focus on them to draw an important conclusion about moral utterances.

Recall that we are trying to figure out what we are intending to do with our moral language. The noncognitivist says that we are not intending to assert anything, rather we are intending to indicate an attitude. But Cuneo points out that "considerations of charity dictate that we should not interpret the traditional believer's moral discourse as the ... expressivist suggests" (Cuneo 2006: 62). When a traditional religious believer says "God is just", what are they intending to do? If we accept the noncognitivist account, this utterance does not attempt to assert anything about God, and must instead be something like approval about God. But by the principle of charity this is false. Traditional religious believers are in fact intending to assert something: they are asserting that God has the property of being just. They are not merely expressing or indicating an attitude of approval of God. It is of course besides the point whether God in fact exists, or whether God is in fact just. It matters only what traditional religious believers are intending to do with their utterances. And what it means to be a traditional religious believer is that you think God exists and has certain properties, among which is 'being just'.

Noncognitivists are forced to disagree with this account. They must say that traditional religious believers are not actually asserting anything with their moral utterance "God is

just". Rather, they are merely expressing something like approval about God. But this is not a charitable interpretation of what traditional religious believers are doing. It is to flatly ignore their background religious commitments. Moreover, the noncognitivist gloss here does not concern "the actual illocutionary act intentions of traditional religious believers who moralize" (Ibid. 65). Rather, it seems to be a revisionary take on religious belief.

We can then generalize from this take on traditional religious believers to our moral discourse and practice more generally. And if we seeking to understand our moral discourse and practice *as we find it*, then cognitivism is clearly the way to go. Only this view allows us to capture the fact that moral utterances are assertions. This of course does not rule out that sometimes moral utterances are expressions. They might be. After all, we cannot understand the illocutionary force of an utterance merely by looking at the sentence construction. So we should accept that sometimes our moral utterances are assertions, other times they are expressions, and so on. But the noncognitivist cannot say this, because their view is defined negatively: our moral judgments are not cognitive and thus not assertions.[19]

This isn't to say that noncognitivists cannot offer a *revisionary* account of our moral discourse and practice. It is of course open to them to say that, while our moral judgments are clearly sometimes intentions to assert something we should instead come to use moral judgments in a different way. I have no objections to noncognitivists pursuing this revisionary line of argument. However, what it demonstrates is that cognitivism is a better account of our moral discourse and practice *as we find it*. As such, cognitivism is the view to beat.

[19] One might ask whether noncognitivists can accept that moral judgments are sometimes assertions. But it seems obvious to me that they cannot. This is because noncognitivists deny that there is a moral reality that is being represented. Indeed, this anti-representationalism is an important motivation of the view. As a result, noncognitivists cannot accept that some moral judgments are sometimes assertions, since they deny that there is any moral reality about which one can assert something. However, as I noted above, there are some sophisticated versions of noncognitivism that rely on a minimalist or deflationary sense of truth, and so may in principle be able to capture some of these representationalist insights.

So to quickly summarize this section: I surveyed an argument in favour of cognitivism in order to motivate my assumption that nonnaturalist cognitivist realism is the view to beat in metaethics. Nonnaturalist cognitivist realism of course requires cognitivism, and I have shown that cognitivism is independently plausible. It better accounts for our moral discourse and practice than noncognitivism.

2.3 Second Cut: Nonnaturalism versus Naturalism

Recall that in section 2.1 I characterized metaethics as the examination of our moral discourse and practice along four lines: linguistic, psychological, metaphysical/ontological, and epistemological. In section 2.2 I looked at our moral judgments along linguistic and psychological lines, endorsing cognitivism. I will now turn to metaphysical and ontological matters. This will set up the remainder of the thesis, which is epistemological.

If the arguments in section 2.2 are roughly correct, then when we make a moral judgment we are indicating or expressing our belief that an action or inaction, person or group, motive or consequence, or any such thing that we deem it appropriate to evaluate in specifically moral terms satisfies a moral predicate. The linguistic and psychological characterization of moral judgments as beliefs sets up the following question: beliefs about *what*? The cognitivist answer is that our moral judgments are beliefs that whatever we are evaluating in specifically moral terms has a certain moral property. When we say and think that an action, like 'torturing people merely for fun', satisfies the predicate "...is wrong" what we are saying is that we believe that this action has the property of wrongness. And in this particular utterance or belief, 'wrongness' is a specifically moral property.

Another way to characterize the matter is in terms of the *truthmakers* of our moral judgments. If when we judge that "torturing people merely for fun is wrong" we are expressing a belief, and beliefs aim at truth, then what would make that belief *true*? Whatever makes this belief true is its truthmaker. Plausibly, this belief is true if and only if torturing people merely for fun is *in fact* wrong, just as our belief that "snow is white" is true if and

only if snow is *in fact* white. In this latter case, *snow is white* is true just in case snow has the property of 'whiteness' — in which case, our belief that "snow is white" would be true because the property of whiteness serves as the truthmaker for our belief. Likewise, in order for the action 'torturing people merely for fun' to in fact be wrong, it would have to have the property of wrongness — in which case, our belief that "torturing people merely fun for is wrong" would be true.

So if our moral judgments are beliefs about moral properties, then we need to ask about the *nature* of these properties. What metaphysical and ontological characteristics do moral properties have? Again, I cannot canvass all such views.[20] But there are typically two competing positions: either moral properties are *natural* properties or else they are *nonnatural* properties. Following the literature, call these two positions *moral naturalism* and *moral nonnaturalism*, respectively (hereafter just *naturalism* and *nonnaturalism*). In this section I will focus on these two views of moral properties, and defend nonnaturalism.

Probably no philosophical label is more fraught with controversy and disagreement than the label *naturalism*. Many people want to identify as a naturalist, or endorse some form of naturalism, but do not say much about what that means. Still, it has become a platitude to endorse naturalism to some degree or another, with some saying that "with each passing decade we get further indications that naturalistically is the way we must learn to understand ourselves" (Gibbard 2012).

One demarcation that we might make is between *things in the natural world* as opposed to *things in the supernatural world*, where the latter is meant to capture entities like gods and devils, heavens and hells, and so on. But supernaturalism of this sort is clearly not what is being positioned against specifically metaethical naturalism. Probably there are few metaethicists who would describe their own positions as *supernaturalist* (the label would

[20] But see McPherson & Plunkett (2017) for a comprehensive overview.

probably sound pejorative) so we should not rely on the distinction between naturalism and supernaturalism.

And yet, if we have an understanding of the supernatural wherein it describes things that some people think genuinely exist, but which are not captured by or reducible to the natural world studied by science, then we might be in a position to see how there may be *the natural world* and then *something else*.[21] So let us begin by saying that naturalists are those who think that the only things that exist are reducible to the natural world, and that nonnaturalists (of which supernaturalists would be a species) are those who think that in addition to those things there are also some other things.[22]

But which sort of things? Because our focus is on *properties*, specifically on *moral properties* (rather than on things or entities), a good way to zero in on the distinction between naturalism and nonnaturalism is as follows: naturalists are those who think that moral properties are part of the natural world, and nonnaturalists are those who think that moral properties are not part of the natural world. For most purposes this distinction would be fine, even if it leaves things vague regarding the exact characterization of the natural world. Indeed, it is up to naturalists and nonnaturalists to fill in the details about what counts as part of 'the natural world' and what does not. But for present purposes we will need more precision. In particular, the phrase 'not part of the natural world' is insufficient for distinguishing between two views: nonnaturalist cognitivist realism and moral error theory. And if our characterization of naturalism is too vague to account for that distinction, the

[21] Talk of 'worlds' here is, while metaphorical, also a bit misleading. Referring to the supernatural world makes it sound as if that world exists in a particular location. But, strictly speaking, we should instead talk of things that have a spatiotemporal existence and things that do not have spatiotemporal existence.

[22] See also Leary (2017) who suggests that a core commitment of nonnaturalism is the claim that "[science] provides an incomplete account of reality" (Leary 2017: 79).

rest of the thesis will not be able to make headway arbitrating between those two views similar but incompatible views.

To see why the phrase 'not part of the natural world' is too vague, consider a fanciful parallel. Suppose that we have the following two claims about unicorns:

(1) Unicorns are part of the natural world

(2) Unicorns are not part of the natural world

These claims are of course mutually exclusive. And suppose that there are three sorts of people:

(a) people who think unicorns exist on Earth

(b) people who think unicorns exist but only in an unearthly mystical dimension

(c) people who think unicorns do not exist

These three sorts of people are mutually exclusive as well. When we cross-reference these two sets it is clear that (a) will endorse (1), since the Earth is part of the natural world. It is also clear that (b) will endorse (2), since unearthly mystical dimensions are not part of the natural world. Likewise, (c) will also endorse (2), since things that do not exist are not part of the natural world.

So far so good. But suppose that we want to distinguish between (b) and (c). Neither will endorse (1), and both can endorse (2), and yet the positions are importantly different. (2) seems to obscure an important difference between *existence in a nonnatural world* and *non-existence*. In this case, if we are attempting to distinguish (b) from (c), the distinction between (1) and (2) is not helpful.[23] So the locution 'not part of the natural world' is insufficiently precise for some purposes.

[23] Note that adding in the qualifier 'if they exist' is similarly unhelpful, since it rules out (c) entirely.

Ditto for characterizing naturalism and nonnaturalism in terms of whether moral properties are part of the natural world or else not. Someone who thinks they are is of course a naturalist. But someone who thinks they are not could be either a nonnaturalist or else a moral error theorist. So, strictly speaking, the characterization of naturalism and nonnaturalism in terms of whether moral properties are 'part of the natural world' is too vague.

For this reason, we will have to distinguish naturalism and nonnaturalism at a more general level. Not unlike noncognitivism, nonnaturalism is first and foremost a rejection of a particular view (as we saw, *noncognitivism* captures any metaethical characterization of our moral judgments that rejects the cognitivist claim that moral judgments are best described as cognitive states). Hence, we will want to understand naturalism *simpliciter* in order to understand nonnaturalism.

There are roughly two strategies that we can pursue in order to characterize naturalism. The first strategy is to identify some unique characteristic that is possessed by all and only natural properties (Shafer-Landau 2001: 58). Call this the *feature strategy*. The second strategy is to identify some unique characteristic of the subject matter that is possessed by all and only naturalist disciplines (Ibid.). Call this the *subject matter strategy*.

The *feature strategy* starts by looking at plausibly naturalistic *things*, like moderate-sized dry goods, and then attempts to uncover what, if anything, they all share in common.[24] On the other hand, the *subject matter strategy* starts by looking at plausibly naturalistic *subjects*, like biology, physics, economics, etc., and then attempts to uncover what, if anything, they all share in common. One might suppose that the *subject matter strategy* entails that there is some particular feature that the focus of that subject share. But strictly speaking the *subject matter strategy* can be silent on whether there is such a feature.

[24] "But then we have to ask, of course, what this class comprises. We are given, as examples, 'familiar objects' — chairs, tables, pictures, books, flowers, pens, cigarettes . . . But does the ordinary man believe that what he perceives is (always) something like furniture, or like these other 'familiar objects' — moderate-sized specimens of dry goods?" (Austin 1958: 8).

For the *feature strategy*, we might start by trying to fill in the consequent for the following sentence: something is natural if and only if ____. From there we might canvass options that seem to pick out all and only natural things. For example, something is natural if and only if it is TOUCHABLE, or MATERIAL, or CAUSALLY EFFICACIOUS, or whatever. And in doing so we will be able to say that something is natural if and only if it is TOUCHABLE, and thus nonnatural if it is NOT TOUCHABLE. In this way, the *feature strategy* promises to furnish us with a distinction between naturalism and nonnaturalism.

However, there is an important and well-known problem with the *feature strategy*. There does not seem to be any accepted (or acceptable) feature that picks out all and only natural things. It seems that however we fill in the consequent above, we will always be able to identify some seemingly natural thing that is not captured. To see why, let us look at some problems with each of these candidates already mentioned, and from there draw a conclusion about the failure of the *feature strategy* in general.

Take the feature of being TOUCHABLE. This seems to capture many natural things. Moderate-sized dry goods are TOUCHABLE. But what about quarks? Surely quarks are natural but it does not seem that they are themselves TOUCHABLE. Or what about being a brother? Certainly that too is a natural thing (it's not supernatural, or nonnatural), and yet being a brother is not itself TOUCHABLE (Shafer-Landau 2001: 59). So being TOUCHABLE seems like an inadequate feature for distinguishing natural things from nonnatural things.

Similar problems arise when we slot in other candidates. MATERIALNESS? Physical fields are natural yet not themselves MATERIAL. CAUSAL EFFICACIOUSNESS? Self-identify is natural yet not itself CAUSALLY EFFICACIOUS. What if we instead said that something is natural if and only if it is something that WOULD HAVE EXISTED PRIOR TO OR APART FROM HUMAN BEINGS? Shafer-Landau notes that on this definition of naturalism, human beings themselves and all objects that human beings create would have to be considered nonnatural (Shafer-Landau 2001: 59). And while that might rule in some natural things, like rocks and trees, it rules out other natural things, like tables and chairs.

In the absence of a plausible feature that all and only natural things share, it seems that the *feature strategy* is going to fail. So if we want to preserve the distinction, our best way of distinguishing naturalism and nonnaturalism will not be in virtue of the *feature strategy*. Instead, if we are going to maintain a distinction between naturalism and nonnaturalism it will be through the *subject matter strategy*.

According to the *subject matter strategy,* naturalism is the view that the only properties that exist are "those that would figure ineliminably in perfected versions of the natural and social sciences" (Shafer-Landau 2001: 59). Notice that here the emphasis is not on the *feature* of natural things themselves, but instead on the *subject's* methodology for coming to know about natural things. Certain subjects have a unique methodology, and cluing into this methodology will allow us to capture the natural/nonnatural distinction. Some have suggested that we can say that a property is natural if and only if we can discover empirically that an object has some property, and a property is nonnatural if and only if we cannot discover empirically that an object has that property (Streumer 2017: 5). Note how this relies on empirical observation for coming to know about a particular natural property (and non-empiricism for coming to know about a particular nonnatural property) but remains silent about the property itself. So there is a methodological difference here for discovering different sorts of properties: empirical observation for the natural, and something other than empirical observation for the nonnatural. Following up on this suggestion will allow us to fill in the details of the *subject matter strategy*.

Empirical observation is the standard methodology for the natural and social sciences, which includes subjects like physics, chemistry, biology, and so on, but also psychology, economics, sociology, anthropology, and the like. If we adopt the *subject matter strategy* then we may think that whatever physicists et al. in their perfected disciplines can study is *ipso facto* natural. Ditto for psychologists, economists, and so on. Empiricism is typically thought to be divisible into three types of knowledge: what we know by *direct observation*, what we know by *deduction* from direct observation, and what we know by *induction* from

direct observation (Huemer 2017: 592). I will take it for granted that physics and psychology, chemistry and economics, and so on, generate knowledge via direct observation, and then valid deductions and inferences from observations. Empiricists, then, are those who think that everything we know is owing to these sorts of processes.

Empirical disciplines are probably our most familiar way of generating knowledge. However, other disciplines purport to be non-empirical. Moreover, they purport to be non-empirical while also claiming to generate knowledge. For example, astrologists do not purport to discover things via direct observation (nor from deductions or inferences thereof). So astrology is non-empirical, and thus nonnatural. Yet you might just as well think that if astrology is non-empirical it is because it is hokum. If you think that art criticism is more respectable than astrology, then note that it too does not seem to rely on direct observation: while art objects are (usually) construed as natural things, the properties that art critics say they can possess need not be (e.g. beauty, sublimity, etc.).[25] Nor is theology or (some parts of) religious studies empirical. Whatever methodology these subjects and disciplines use seems non-empirical. Thus, what the *subject matter strategy* suggests is that whatever properties astrologists, art critics, theologists, and so on, study that is disciplinarily unique is *ipso facto* nonnatural. But maybe you're happy to lump religious studies, theology, art criticism, and astrology together and assign them to the intellectual garbage can. In that case, you should be prepared to do the same to logic, mathematics, and metaphysics. While these latter disciplines of course rely on deductions and inferences, they do not seem to stem from empirical observation. But I take it that critiques of astrology on the grounds of its non-empiricism will not so quickly carry over to mathematics.

All this to say that we should be cautious of writing off all non-empirical subjects. And if we are going to allow that such subjects are *prima facie* legitimate in their generation of

[25] What we might want to say here is that a nonnatural property like beauty *supervenes* on the natural art object. However, this would not make the property of beauty itself natural.

knowledge, then we may have a suitable way of distinguishing naturalism from nonnaturalism. Without identifying a particular feature, we can say that a property is natural if and only if it can be discovered via empirical methods. And a property is nonnatural if and only it cannot be discovered in this way.

If that is roughly correct, then where should we slot morality: with the natural or with the nonnatural? While some stress that 'naturalistically' is the way in which we must learn to understand ourselves, our moral discourse and practice makes trouble for this suggestion. Morality, in particular its requisite properties, does not seem to find a comfortable home alongside physics and chemistry, nor does it find much kinship with psychology and economics. This seems to indicate that morality is a nonnatural subject.

Supporting this claim, Shafer-Landau suggests four criteria that exclude morality from the natural subjects. First, morality does not seem to admit of "the kind of precision and susceptibility to quantification and cardinal measurement that the purportedly natural disciplines possess" (Shafer-Landau 2001: 58-59). Second, morality is unlike those disciplines insofar as it is "essentially concerned with practical evaluation and recommendation" (Ibid., 60). Third, morality does not appear to focus "on questions of prediction and causal efficacy" (Ibid.). And fourth, if there are fundamental moral truths they "seem to express metaphysical, rather than physical (or biological or chemical, etc.) necessities" (Ibid.).

Of course, one option here is to insist that the subject matter of morality is natural but suggest that it is nevertheless unlike any of the other natural subjects with which we are familiar. Perhaps morality is unique amongst all natural subjects. But this seems to fall out of some antecedent commitment to thinking that morality *must be* in the business of talking about natural properties rather than a response to our moral discourse and practice as we find it. Recall again the suggestion from above: if it is true that with each passing decade we find increasing evidence that naturalism is the way to go, then this might give us reason to assume that morality *must be* understood naturalistically. However, this obscures another option: that our moral discourse and practice is fundamentally in error (which is

what moral error theorists argue). So we should not necessarily commit ourselves to having to find a home for moral properties among the natural. We should be open to the idea that moral properties are not part of the natural world *at all*: they may, like unicorns, exist in some unearthly mystical dimension, or they might be illusory.

So if we have a sense of the methodology of the disciplines above — namely that they are empirical, which means that they are in the business of quantification and measurement, primarily concerned with description, and uncovering physical necessities — then we will need to identify a corresponding methodology of nonnatural disciplines. Empiricism is of course typically contrasted with rationalism. Empiricists are those who think that all knowledge comes via *a posteriori* reasoning, while rationalists are those think that at least some knowledge comes via *a priori* reasoning. In subsequent chapters I will explore how it is that we might come to have moral knowledge *a priori*. For now we need only say that, however we might come to have that knowledge, it is non-empirically.

To sum up, naturalism is the view that moral properties, if they exist, are ontologically and metaphysically unexciting (Enoch 2017). They are the sort of properties that we ordinarily use to explain the natural world. And they are the sort of properties that we are fairly accustomed to discovering via empirical methods. To be sure, naturalists need not hold that we have identified exactly which natural properties serve as the truthmakers of our moral judgments. Rather, naturalists need only say that, whatever properties serve as the truthmakers of our moral judgments will be discoverable *a posteriori* through the usual methods of uncovering properties — but strictly speaking they can remain silent about exactly which properties those are.

On the other hand, if moral properties are not natural properties then we cannot discover them via our ordinary empirical methods. And similarly, nonnaturalists need not hold that we have identified exactly which nonnatural properties serve as the truthmakers of our moral judgments. Rather, they need only say that, whatever properties serve as the truthmakers here will be discoverable *a priori*. But not any old *a priori* method will do. We should

note also that while morality, mathematics, and metaphysics are all plausible nonnatural subjects that rely on *a priori* reasoning, we come to know moral truths via moral thinking, mathematical truths via mathematical thinking, and metaphysical truth via metaphysical thinking. These subjects are in an important sense autonomous. And this is precisely why we need a convincing moral epistemology, a convincing story of how we can come to have moral knowledge. This is what the subsequent chapters of this thesis explore.

2.3.1 Argument for Nonnaturalism

A common objection to nonnaturalism, dating back even to its earliest form, is that it is ontologically or metaphysically extravagant, positing properties or relations (or both) that are over and above those that we ordinarily use to explain the natural world (Hurka 2014: 89). This has become platitudinous, with some saying that "the orthodox view in the meta-ethical debate of the second half of the twentieth century was that . . . a nonnatural unanalyzable property . . . was metaphysically adventurous" (Olson 2009: 164). Even defenders of nonnaturalism accept this charge. Derek Parfit, for one, accepts that

> Irreducible [moral] truths, if there are any, are most unusual. As many writers claim, it is not obvious how such truths fit into a scientific world-view. They are not empirically testable, or explicable by natural laws. Nor does there seem to be anything for such truths to be about. What can the property of badness *be*? Given these points, it is natural to doubt whether these alleged truths even make sense. If such truths are not empirical, or about features of the natural world, how do we ever come to understand them? (Parfit 2006: 330, emphasis his)

It is worth noting that not all versions of nonnaturalism are (or, rather, take themselves to be) open to this 'too adventurous' charge. There are at least two versions of nonnaturalism. One view suggests that nonnatural properties are what *make* true moral judgments true, while another view suggests that nonnatural properties do not play this role. This is roughly the difference between *robust realism* (e.g. Enoch 2011, Wielenberg 2014) and *quietist realism*

(e.g. Parfit 2011, Scanlon 2014). Robust realists say that nonnatural properties exist in the full-on sense of that word, insofar as positing them is an ontological commitment on our part. But quietest realists say that nonnatural properties 'exist' in a way that does not commit us to anything ontologically weighty.[26]

In what follows I will rely on a robust realist conception of nonnaturalism rather than a quietist conception. Insofar as I need to take a position, I think that robust realism more readily captures both our moral discourse and practice, but more importantly it better squares with our understanding of what it means to say that properties 'exist'.[27]

Because nonnaturalism (of a robust realist stripe) *does* seem to be ontological or metaphysically weighty, it is not a position that we should accept lightly. Nonnaturalism is difficult to square with a scientific world-view, as Parfit notes, but it also means that our perfected natural and social sciences would be explanatorily incomplete. This is a claim that we should not underestimate. Nonnaturalists do not merely suggest that the sciences merely currently provide an incomplete account of reality. Naturalists can (and should) agree to that. Rather, the nonnaturalist claim is that *by definition* our perfected natural sciences will leave out aspects of reality; in this case, *moral reality* is left unexplained. Thus, if we wish to talk about moral reality, we will be talking about a part of reality that is unable to be captured by our ordinary empirical methods.[28]

---

[26] I should stress that the so-called quietist realists do not generally self-describe as such. See McPherson (2011) for a characterization of those nonnaturalist positions that claim to avoid making ontological commitments. Ontological quietism is McPherson's label for these views, and it is a label that has been adopted largely by detractors of quietism. For their part, so-called quietists instead use terms like *non-metaphysical nonnatural cognitive realists* (Parfit 2011), or else just *realists* (Scanlon 2014).

[27] For recent critiques of quietism, see McPherson (2011), Enoch (2011), and Wodak (2017). The general thrust of these critiques is that quietism does not sufficiently distinguish itself from forms of fictionalism, and in that sense does not genuinely qualify as a form of realism.

[28] See also Plantinga (2011), who argues that we cannot easily move from methodological naturalism (i.e. the idea that our natural and social sciences study natural things) to

Compare this to a familiar debate between theists and atheists. Many atheists will insist that God does not exist because God is in principle incapable of being discovered and described by even our perfected sciences. Many theists will of course deny that this is a problem. "Of course God cannot be discovered and described by our natural sciences", they might say, "because our natural sciences can only discover and describe what is natural. God is *supernatural*, so our perfected sciences are simply not suited to the task." I think that many atheists are rightly suspicious of these sorts of claims, as it appears to license all sorts of fanciful entities (e.g. "Of course you can't see the teapot that's floating between the Earth and the Moon..."). When our usual ways of knowing are said to be inadequate, we should be wary of taking on board new claims.

So why should one accept nonnaturalism? One reason stems from something that we looked at above. Recall that we attempted to characterize the distinction between naturalism and nonnaturalism via two strategies: the *feature strategy* and the *subject matter strategy*. In availing ourselves of the *feature strategy* we cannot seem to find something in common to all and only natural things. Plausible candidates like MATERIALNESS, CAUSAL EFFICACIOUSNESS, DEPENDENCY ON HUMAN EXISTENCE, and so on admit of obvious counterexamples. So we concluded that the *feature strategy* was an ineffective way of distinguishing naturalism from nonnaturalism. That in turn led us to the *subject matter strategy*. Yet, one can coherently think that what distinguishes naturalism and nonnaturalism is a disciplinary approach, but that the latter disciplines simply do not capture anything that genuinely exists. Theology, for example, might be a nonnatural discipline but we might nevertheless maintain that its objects of study are illusory.

But suppose that we start from some sort of moral realist position, thinking that some things are morally good, others things are morally bad, some actions are permissible, other

---

metaphysical naturalism (i.e. the idea that methodological naturalism alone will fill up our ontology).

forbidden, and so on. Above we routinely used the example "gender discrimination is bad". If you think that it is, then what makes that judgment true? From the above sections, we should conclude that our belief is true when the action 'gender discrimination' actually has the property of 'badness'. And this property can only be of two types: either it is a natural property (discoverable just by *a posteriori* reasoning) or else it is nonnatural (discoverable just by *a priori* reasoning).

But if it is a natural property, which property is it? Earlier I said that naturalists do not have to answer this question, rather they need only insist that *whichever* property it is it will be a natural one. So we should not press the naturalist on this point. Rather, we can show that the naturalist project of showing a moral property to be identical with *any* particular natural property will ultimately fail. Here we will be relying on G.E. Moore's *Open Question Argument* (Moore 1903). With the Open Question Argument Moore was targeting the view that certain moral properties (like rightness or wrongness) are *identical* to nonmoral ones. This view is called *analytic reductionism*. What the open question argument provides is a general strategy for showing that we this purported identity relation will not hold.

Because we have relied so often on the toy moral judgment "gender discrimination is bad", we can home in on the property of moral *badness*. If we are trying to show that moral badness is identical to some natural property, then we might try slotting in various candidates. For example, if we are naturalists we might suggest that 'gender discrimination is bad' is identical to CAUSING PAIN. CAUSING PAIN is of course a perfectly respectable natural property. It is something that we can empirically discover, quantify, and so on. Or we might instead say 'gender discrimination is bad' is identical to REDUCING HAPPINESS. Or that it's identical to DECREASING SOCIAL TRUST or DECREASING SOCIAL STABILITY. Or that it's identical to PRODUCING INEQUALITY. Or maybe a combination of all of these things.

If we are nonnaturalists, and if we were modeling our approach here on our earlier undermining efforts regarding the *feature strategy*, then what we would do is find an action

that is morally bad that does not seem to (e.g.) cause pain, or reduce happiness, or produce inequality, etc. This piecemeal strategy might be effective for knocking down a particular proposal, but it would not itself constitute an argument in favour of nonnaturalism.

However, Moore's Open Question Argument is a handy tool for showing why moral properties are not identical non-moral properties. Take any moral property, for example, badness, and any non-moral property, for example CAUSING PAIN. As noted above, analytic reductionism is the view that a moral property can be said to be identical to the non-moral one. In this case, the purported identity relation is that moral badness is identical to the causing of pain. Now, what the Open Question Argument says is that we can test this sort of identity by asking the following question: "is causing pain bad?" If the identity here is accurate then 'causing pain' and 'bad' are identical, and the question is *closed*: it would be like asking "is a female fox a vixen?". But the former question about whether causing pain is bad sounds open, as if it can be sensibly answered (or, at least, that it admits of further explanation and examination), while the latter question about vixens sounds closed, as if it cannot be sensibly answer because it is tautological. So the identity relation here seems to fail.

Moore suggests that any such naturalistic reduction will generate *open questions*, and that what this demonstrates is that moral properties are not identical to natural properties. This gives us reason to think that the property is not a natural property. So it must be a nonnatural property.

And yet, if Moore's target is analytic reductionism, then it might be the case that slotting in a nonnatural property here is just as problematic. To see why, note that analytic reductionism is different than naturalism as we have described it above. This is because analytic reductionism does not say exactly what the nature of the non-moral properties in question might be. While they are usually taken to be natural properties, they might also be nonnatural (but non-moral) properties, or they might be supernatural properties. So there is some daylight between analytic reductionism and naturalism.

So we may be able to go further (at least, further than Moore) by suggesting that the Open Question Argument rules out identity between moral properties and natural properties but also identity between moral properties and supernatural properties. Take, for example, the supernational property BEING DISAPPROVED OF BY GOD. If we think that the moral and non-moral properties here are identical, then the following question should sound tautological: "is being disapproved of by God bad?" But this question instead sounds open.

What this may lead us to conclude is that not only are moral properties nonnatural they are also *sui generis*. Moral properties are nonnatural properties yet they are unique among all nonnatural properties. They are their own special class of properties — nonnatural, of course, but not just any old nonnatural properties.

As before, this is far from a conclusive argument in favour of nonnaturalism or against naturalism. For one, it is generally agreed upon that the Open Question Argument is controversial. One thing that it overlooks is the difference between the *meaning* of terms and the *identity* of properties (think about how the meaning of the terms 'water' and 'H2O' differ yet pick out the same property). For another, naturalists might be able to come up with a natural property and thus identity relation that genuinely sounds closed. If this is true, then the Open Question Argument is guilty of begging the question against naturalism.

But recall that this is not a dissertation canvassing the arguments in favour or against these views. I present this argument in favour of nonnaturalism in order to motivate my assumption that nonnaturalist cognitivist realism is the view to beat in metaethics. Nonnaturalist cognitivist realism of course requires nonnaturalism, and I have attempted to show that nonnaturalism is an independently plausible position regarding moral properties. Moreover, it is the view that one will have to depart from in order to be a naturalist. I do not think that we start out as naturalists, but rather we become them on the strength of one argument or another. So nonnaturalism is going to be the view to beat.

2.4 Third Cut: Success versus Failure

In section 2.1 I said that metaethics is the examination of our moral discourse and practice along four lines: linguistic, psychological, metaphysical/ontological, and epistemological. At this point we have canvassed three quarters of the metaethical landscape. In section 2.2 I looked at our moral judgments along linguistic and psychological lines, endorsing cognitivism. In section 2.3 I turned to metaphysical and ontological matters, endorsing a version of (robust) nonnaturalist about moral properties. The remainder of the thesis will concern epistemological questions. Nonnaturalist cognitivist realism requires us to accept both cognitivism about moral judgments and nonnaturalism about moral properties, and I think that the view is complete when we add to that the following claim: *success* regarding some of our moral judgments.

What does it mean for our moral judgments to be 'successful'? I think that a judgment like 'torturing people merely for is wrong' is successful when we know (or at a minimum are justified in believing) that the judgment is true. The important aspect here is the emphasis on *knowing* (or being *justified* in believing) that the judgment is true, rather than on the judgment merely *being* true. This is what makes the issue of success epistemic, rather than simply ontological or metaphysical. We should not accept a metaethical account that says our moral judgments are successful just in case there *are* moral facts that accord with those judgments. I think that a view like that is problematic for three reasons. First, it rules in what we can call *lucky* moral epistemologies. Second, and relatedly, it does not allow us to capture the way in which successful moral judgments are an accomplishment. And third, it does not properly distinguish nonnaturalist cognitivist realism from epistemic moral scepticism. In this last section I will quickly explain these problems, which shows why we should understand the success of moral judgments in epistemic terms.

The first two problems are related to each other. A purely ontological or metaphysical notion of successful moral judgments allows us to have 'lucky' correct moral judgments, which in turn erases the accomplishment that accompanies success. If our moral judgment

that 'torturing people merely for fun is wrong' is successful just in case there is a nonnatural moral fact *that torturing people merely for fun is wrong*, then it seems irrelevant how we actually come to form that judgment. Consider the following: coming to have the belief that 'torturing people merely for fun is wrong' after a strong blow to the head would be as good as coming to have it after reflection on the moral notions inherent to the situation. In this way, we are simply lucky that the blow to the head produced a belief that accorded with the moral fact of the matter. But something about this case is clearly off, and it is not obvious that a purely ontological or metaphysical notion of success can tell us why.

This in turn shows how a purely ontological or metaphysical notion of success erases the way in which successful moral judgments are actually an accomplishment. It seems to matter *how* it is that we come to have the moral judgments that we do. Ignoring this seems to be a mistake. We should want there to be *some* connection between the correct and successful moral judgments that we make and the moral facts. I think that this connection should not be a lucky one but instead *explanatory*: the moral facts should in some way figure into the full explanation of how our moral judgments are successful.[29] This, I take it, rules out lucky correct moral judgments, and requires that we come to *know* the moral facts rather than by chance hit upon them. Because (as we will see in Chapter 3) there is a gap of sorts between our moral judgments and the nonnatural moral facts, bridging that gap requires some sort of effort or at least explanation on our part. But if successful moral judgments are not understood in epistemic terms then the sense in which it is an accomplishment to bridge that gap is lost.

[29] While I intend this as a very weak claim, it is nevertheless a substantive thesis. Some moral philosophers deny that the moral facts do figure into the best explanation of our moral beliefs - see Harman (1977) and Thomson (1996, 1998). For defense of an explanatory requirement, see Sturgeon (1984) and Loeb (2007). This so-called *moral explanations* debate turns on whether the moral facts can actually explain our beliefs about them. But the explanatory requirement I am appealing is intended merely to rule out *lucky* or *concidental* moral knowledge.

I will call a non-lucky moral epistemology that includes this explanatory component a *plausible* moral epistemology. Whatever it turns out to be, this plausible moral epistemology is of course constrained by cognitivism about moral judgments and nonnaturalism about moral properties. The remaining chapters of the dissertation explore various accounts of how we can come to have moral knowledge given these two key desiderata.

The third problem is that a purely ontological or metaphysical notion of success brings with it a spectre of scepticism. As we saw in the previous chapter, moral realism is supposed to be a *non-sceptical* metaethical view. If we take that non-sceptical component seriously then we will have to understand the success of moral judgments in epistemic terms. Recall that moral realists say that there *are* moral facts, while moral sceptics say that whether or not there are moral facts we are not in a position to *know* them. Nonnaturalist cognitivist realism is of course a form of moral realism, and so the view should be readily distinguishable from moral scepticism. If so, then it seems we should accept an epistemic notion of the success of moral judgments. The epistemic moral sceptic can certainly agree that there *are* moral facts — they will simply insist that nevertheless we do not know anything about them. Because of that, if we have a purely ontological or metaphysical notion of successful moral judgments then there is nothing to distinguish the moral realist and the moral sceptic here. And that is why we should prefer an epistemic notion of success — that way a moral realist can then properly separate themselves from the moral sceptic by insisting that we have knowledge (or justified beliefs) about those facts.

Because of those issues, we should understand the success of moral judgments in epistemic terms. It is also worth noting that understanding the success of moral judgments in epistemic terms is not without precedent. Judith Jarvis Thomson has a success theory regarding our moral judgments. She proposes an account of moral objectivity that is first and foremost epistemological in nature (Thomson 1996). She suggests that moral facts are to be considered objective (i.e. independent of our thoughts and attitudes about them) if it is in principle possible for us to come to know truths about them. Notice here that

Thomson is describing what would ordinarily be an ontological claim (that moral facts are mind-independent) in specifically epistemic terms: moral facts are mind-independent only if we can come to *know* them. This contrasts with those philosophers who suggest that moral objectivity is fundamentally an ontological thesis: moral facts are objective if they exist in a mind- and attitude-independent way. So, according to Thomson, there are those who think that moral epistemology precedes moral metaphysics, and those who think that moral metaphysics precedes moral epistemology. She fits into the former camp.

The reason that Thomson puts moral epistemology ahead of moral metaphysics is because she says she takes seriously our moral discourse and practice. She argues that when we are attempting to figure out whether a particular moral judgment is true (like 'torturing people merely for fun is wrong'), we are engaged in so-called "moral assessment" (Thomson 1996: 67). Built into our ordinary moral discourse and practice is the idea that moral assessment is not a pointless endeavour, rather "it is possible to find out about some moral [judgments] that they are true" (Ibid. 68). Because the emphasis is on *finding out* whether a moral judgment is true, this in turn constrains how one can characterize the moral facts. Whatever ontological or metaphysical account is ultimately plausible will have to be one that allows us to see how it is that we can come to *know* whether a given moral judgment is true. Importantly, this rules out the sort of 'lucky' moral epistemologies I mentioned above. And for Thomson, this makes moral objectivity an epistemological matter, and not (as some might think) an ontological or metaphysical one.[30]

To sum up, if nonnaturalist cognitivist realism stands or falls only according to ontology or metaphysics, then successful moral judgments are simply those that are *in fact* true.

[30] To be clear, I am not taking a stand on whether moral epistemology precedes moral metaphysics, or vice versa. Rather, I am taking a more holistic approach. Nonnaturalist cognitivist realism can only be vindicated if it can provide a plausible account of how we come to have moral knowledge. This says nothing about the priority of epistemology or metaphysics. Rather, it ensures that whatever metaethical view is vindicated will have to say *something* substantial regarding each of the four areas within metaethics.

Our moral judgments would be successful depending only on whether or not there are the appropriate nonnatural moral facts. But I suggested that this is problematic because it rules in lucky moral epistemologies, does not adequately capture the way in which successful moral judgments are accomplishments, and does not allow nonnaturalist cognitivist realism to distinguish itself from epistemic moral scepticism. To remedy this, we should instead think that nonnaturalist cognitivist realism is vindicated only if we *know* that our moral judgments are true. And this does not depend merely on whether there *are* moral facts, but rather it crucially depends on whether there is a plausible account of how we know that there are such facts. So we should understand the success of moral judgments in epistemic terms. In doing so, we rule out mere luck, maintain the sense of accomplishment for our judgments, and keep nonnaturalist cognitivist realism and epistemic moral scepticism separated.

So we should keep these three criteria in mind: cognitivism about moral judgments, nonnaturalism about moral properties, and an understanding of successful judgments in epistemic terms. A plausible moral epistemology for nonnaturalist cognitivist realism will need to meet all three.

CHAPTER 3
Evolutionary Debunking and the Folk/Theoretical Distinction

3.1 What is the Evolutionary Debunking Argument?

In this chapter I address a sceptical challenge that has recently captured the attention of metaethicists: the evolutionary debunking argument (EDA).[1] One aim of the EDA is to show that, owing to some to-be-specified evolutionary influence, our moral judgments are unjustified. Characterizing the EDA is a matter of considerable debate, but in general the argument goes something like this: if evolution has influenced our moral faculties in a particular way, then we should not think that our moral judgments are justified; evolution *has* so influenced our moral faculties; so we should not think that our moral judgments are justified. Because the first premise is a claim about when justification is defeated or debunked it is an *epistemological* premise. Because the second premise is a claim about our evolutionary history it is an *empirical* premise. It is the combination of epistemological and empirical premises that has resulted in the EDA garnering so much attention. So-called 'debunkers' think that the evolutionary etiology of our moral judgments has this undermining effect, while 'anti-debunkers' do not.

[1] For an overview of this literature, see Vavova 2015. To my mind, the EDA literature has gone through a few distinct stages. Initial debunking arguments (Bedke 2009; Joyce 2001, 2006; Street 2005, 2008 — and an earlier, if widely disparaged, attempt from Ruse & Wilson 1986) promoted anti-debunking arguments (Brosnan 2011; Clarke-Doane 2012; Copp 2008; Enoch 2010; Nagel 2012; Parfit 2011; Shafer-Landau 2012). This has led to increasingly sophisticated responses, and the EDA literature has ballooned over the past decade (for much discussed contributions see Berker 2014; Bogardus 2016; Das 2016; FitzPatrick 2015; Fraser 2014; Mogensen 2015, 2016; Sinclair 2018; Vavova 2014, 2015).

When it comes to identifying which *vindicatory* metaethical theories are suspect, debunkers disagree amongst themselves over whether it is merely some (e.g. Street 2006) or else all (e.g. Joyce 2001; 2006). By a vindicatory metaethical theory I mean to differentiate those metaethical theories that suggest our judgments are in some way *successful* from those that suggest they are *not* (recall that this is the 'Third Cut' from the previous chapter). For example, while moral error theory is of course a metaethical theory, it is not a vindicatory theory in the sense of allowing us to think that we know the moral facts.

One way to home in on the target of the EDA is by asking *which* moral judgments in particular are suspect. A quick answer might be: those judgments that natural selection probably exerted significant pressure upon. But it makes a difference whether we are calling into question the ability to make any justified moral judgments at all, or whether some specific subset of our judgments is impugned. This is because if the EDA is understood as a *justification* debunking argument, then no matter how moral judgments are construed by any metaethical theory they will be unjustified - so whether we characterize our judgments per cognitivism or per noncognitivism, or per naturalism or per nonnaturalism, we will not be able to secure justification.[2] But if the EDA is understood as a *theory* debunking argument, then whether or not a moral judgment is impugned will depend upon how a particular

[2] For more on the difference between a theory debunking arguments versus a justification debunking argument, see Joyce 2014 and Sinclair 2018.

metaethical theory characterizes our moral judgments - for example, nonnaturalist cognitivist realism characterizes our moral judgments much differently than any noncognitivist or naturalist theories.[3]

This is why I think it is important to address the EDA. I take it that the empirical premise is roughly correct even while it is being increasingly refined.[4] I also take it that something in the ballpark of the epistemological premise is right.[5] Now, if the EDA is properly understood as a justification debunking argument, then the search for a plausible moral epistemology ends right here. It would entail that there can be no plausible moral epistemology *at all*, let alone for nonnaturalist cognitive realism. However, if the the EDA is better understood as a theory debunking argument, then we will need to answer another question: is nonnaturalist cognitive realism one of the targeted theories?

In addition to the distinction between theory debunking and justification debunking, recent work on the EDA has distinguished between the *contents* of our moral judgments

[3] On the one hand, Sharon Street's EDA is a paradigmatic theory debunking argument. In brief, she argues that moral realists are unable to square our evolutionary history with the purported connection between our moral judgments and the moral facts. This leaves open the possibility that certain metaethical *constructivist* views can win the day: these theories do not have to explain a connection between our judgments and the moral facts, rather these theories usually suggest that the moral facts are *constituted* or *constructed* by our moral judgments. Bedke's EDA is also a theory debunking argument, as it targets nonnaturalist versions of moral realism (Bedke 2009; 2014). On the other hand, Joyce's EDA is a justification debunking argument (2001; 2006). While Street thinks that our evolutionary history makes trouble for moral realist theories but not for constructivist one's, Joyce's EDA targets all attempts at justifying our moral judgments.

[4] See for example Bloom 2013; Churchland 2011; de Waal 1996, 2009; Gazzaniga 2005; Greene 2013; Hauser 2006; Joyce 2006; Kitcher 2011; Tomasello 2016; etc.

[5] The epistemological premise is much more controversial than the empirical premise. As a result, there is a lot of debate about how best to formulate it. For recent work that touches on this, see Hanson 2017; Isserow (forthcoming); Lutz (forthcoming); Moon 2017; Schechter 2018; Sterelny & Fraser 2016; Vavova 2015, 2018; etc.

and the *capacities* or *abilities* for making those judgments.[6] Suppose that natural selection favoured those ancestors who made some specific moral judgment j. If one group of early hominids that made only moral judgment j out-produced another group of early hominids who made only moral judgment j* then natural selection might have favoured the former hominids *because of* the specific content of their moral judgment. And if we are still disposed to make moral judgment j, that judgment may be suspect because it had been 'merely' advantageous and not necessarily truth-sensitive.[7] There would then be an undercutting defeater for thinking that j is true.[8] When glossed in this way, the EDA is a *content etiology argument* which attempts to impugn some of our moral judgments.

On the other hand, a *capacity etiology argument* seeds doubt concerning our ability to make moral judgments at all. Again, if one group of early hominids that was disposed to make any moral judgments at all out-produced another group of early hominids *without* that ability, then natural selection might have favoured the ability in the former hominids for making any moral judgments in the first place. So explained, the EDA seems to impugn all of our moral judgments.[9]

Combined with the distinction between theory debunking and justification debunking, casting the EDA's purported undermining effect in terms of content or else capacity helps to

[6] The content/capacity distinction can then correspond to one's characterizations of the EDA, which will make salient those respective features. See for example FitzPatrick 2015

[7] This of course supposes that we ought to understand moral judgments as truth-apt. I suggest in section 3.2 that the typical target of the EDA is a nonnaturalist cognitive realism, which as we know from Chapter 2 characterizes moral judgments in this way.

[8] An *undercutting* defeater, rather than a *rebutting* defeater. The EDA purportedly serves to undermine thinking 'j is true', but it does not directly support thinking that 'j is true' is false. Cf. Pollock (1986).

[9] Moreover, when evolutionary debunking arguments are glossed in terms of a causal premise, epistemic premise, and a conclusion, it is more plausible to read the causal premise of *capacity* rather than *content*. See Kahane 2011 and Sinclair 2018 for a discussion of the EDA structured in this way.

narrow our focus. But these distinctions alone cannot settle the question about *which* moral judgments are problematic. And there is a common anti-debunking strategy that exploits this lacuna. What we can call the *Not All Judgments* (NAJ) response begins by granting to the debunker that some to-be-specified evolutionary influence has probably played a distorting role on our moral judgments, whether that is understood as a distortion regarding the content of our judgments or else the very capacity to make them at all. But the NAJ response continues that, even accounting for such a distortion, we might nevertheless possess the ability to make justified moral judgments.

For example, Derek Parfit writes that, though our capacities for moral judgment "were partly produced by evolutionary forces... these abilities later ceased to be governed by these forces, and had their own effects" (Parfit 2011: 520). Some of our judgments are made owing to pushes from our past, while other judgments are made because our abilities for such judgments have been sufficiently modified. In a similar vein, William FitzPatrick says that "it's enough if natural selection has given us general cognitive capacities that we can now develop and deploy in rich cultural contexts, with training in relevant methodologies, so as to arrive at justified and accurate beliefs in that domain" (FitzPatrick 2015: 5-6).[10] In brief, and on a first pass, the NAJ response to the EDA says that while natural selection may have pushed us towards making some unjustified moral judgments, we have nevertheless ceased to be pushed by those same evolutionary forces and are instead able to make other judgments that are successful. Though evolution may be a defeater for some of our moral judgments, we can make other moral judgments that are successful — we know that they are true.

[10] See also Toner 2011, who suggests that our capacity for moral judgments has been "co-opted... for purposes other than natural selection" (Toner 2011: 529). Similarly, Copp 2008 and Brosnan 2011.

I think that we can identify three aspects of this response. First, as Parfit says, "[if] we can plausibly conclude that certain normative beliefs were greatly influenced by natural selection, this fact may give us strong reasons to doubt that these beliefs are true" (Parfit 2017: 279).[11] But while this influence is plausible for some of our moral beliefs, like "the belief that people can deserve to suffer... [or] that we have reasons to care more about good or bad experiences when these experiences are in the future rather than the past", it is not plausible for other beliefs, like the belief "that everyone's well-being matters equally, and the belief that all suffering is bad" (Ibid.). This makes salient the first aspect of the NAJ response: there are different moral judgments that we can make, some of which might be unjustified but others which are justified.

The remaining two aspects are brought out most clearly by FitzPatrick. He writes that:

> the basic mental capacities that enable us to sit around worrying about things like metaphysical modality are part of our evolutionary heritage: they didn't appear by chance and they weren't designed by God; they evolved through natural selection. But natural selection did not design our cognitive capacities to track truths about metaphysical necessity, which were as irrelevant to the reproductive success of Pleistocene hunter-gatherers as the realist's moral truths would be... Still, despite the irrelevance of obscure metaphysical facts to the natural selection etiology of our basic mental capacities, we're able to deploy those capacities, in the cultural context of philosophical training, to think intelligently and often accurately about things like metaphysical necessity or countless other arcane topics such as differential geometry and relativistic quantum theory, the

[11] Note that Parfit says that evolution can give *us* a reason to doubt that some moral judgments are true. In order to eliminate this pernicious influence, we would need to be able to say that we have reason to think that our moral judgments *are* true. This, again, in an epistemic criterion for a successful moral judgment.

> facts of which are equally irrelevant to the etiology of the capacities we use in thinking about them (FitzPatrick 2015: 887).

The second aspect of the NAJ is this: we should focus on the basic mental *capacities* for certain judgments, as opposed to the *contents* of those judgments. So the NAJ response requires that we understand the EDA as a capacity etiology argument. And the third aspect of the NAJ response is this: philosophical training (or some other cultural phenomenon) can allow us to make successful moral judgments using our evolved capacities.

Combining these three aspects gives the following way of characterizing the NAJ response: though evolutionary influences are responsible for our having the capacities needed to make moral judgments (and thus we have a defeater for thinking we can make successful moral judgments), these capacities eventually cease to be governed by evolution and can later be put to use, in combination with philosophical training, in issuing successful moral judgments.

If you have them, set aside questions about whether our capacities or abilities for moral judgment could truly 'cease to be governed' by evolutionary forces in the relevant way. The question that I will focus on is this: what does it mean to say that evolutionary forces impugn the credibility of some of our moral judgments but not others? How can we, despite evolutionary influences, make successful moral judgments?

Proponents of the NAJ response have not spelled this out. To remedy this, I suggest (in section 3.2) a way of characterizing these two types of moral judgments. On the one hand, we have *folk moral judgments*, which are the sorts of judgments that evolutionary pressures plausibly would have exerted significant distorting influence upon. And on the other hand we have *theoretical metaethical judgments*, which are not. This folk/theoretical distinction may be familiar to us from the natural sciences, and I argue that it can help us flesh out the NAJ response.

Moreover, availing ourselves of this distinction helps show that the EDA should not be understood as a justification debunking argument, but rather as a theory debunking argument. What I call (in section 3.3) **Inference Debunking** constitutes a potent defeater, albeit one with a very specific scope — it undermines the inference *from* our folk judgments *to* our theoretical judgments, but leaves our folk judgments and theoretical judgments themselves untouched. Lastly, since the EDA is best understood as a theory debunking argument, is nonnaturalist cognitive realism one of the targeted theories? I claim that it is not: the view does not merely infer theoretical metaethical judgments from folk moral judgments. In light of this, nonnaturalist cognitivist realism remains the view to beat in metaethics, and the search for a plausible moral epistemology can go on.

3.2 The Folk/Theoretical Distinction

The NAJ response says that, though our capacities for moral judgment evolved, and are subject to the EDA's undermining defeater, when combined with philosophical training those capacities can be used to arrive at justified moral beliefs. This section spells out this response. To do so I suggest that we can draw a distinction between our *folk moral judgments* and our *theoretical moral judgments*. The folk/theoretical distinction here mirrors the perhaps more familiar one from the natural sciences. There it is more common to draw a distinction between pre-theoretical or folk judgments and theoretical or formal judgments. We now readily draw a distinction between, for example, folk judgments about the movement of heavenly bodies and formal astrophysical judgments, between folk judgments about animals and theoretical judgments about biology, and so on. I think metaethicists should also avail themselves of this common distinction in order to characterize different types of moral judgments.

In the natural sciences, folk judgments are typically characterized as "aris[ing] more informally and not as direct reflections of formal instruction in scientific principles" (Keil 2010: 826; also Carey 1988). Similarly, folk sciences are said to be those that "without

explicit instruction in such areas, [lead] people ... to develop domain-specific ways of thinking about relatively bounded sets of phenomena" (Keil 2010: 826). We can capture these 'ways of thinking' about certain phenomena by referring to the folk theories or folk sciences from which our folk judgments issue. We have domain-specific ways of thinking about "the behavior of solid objects, living kinds, and the minds of others" (Keil 2010: 826), or what we would commonly refer to as folk physics, folk biology, and folk psychology, respectively. Similarly, Medin & Atran note that "to get along in the world, people need to be able to understand and predict the general properties and behaviours of physical objects and substances (physics), the more specific properties of plants and animals (biology), and the particular properties of their fellow human beings (psychology)" (Medin & Atran 2004: 926). In addition, work has been done to uncover other folk judgments issuing from folk chemistry (Au 1994), folk cosmology (Siegel et al. 2004), folk economics (Laksminaranan et al. 2008), etc. So the folk sciences are those ways of thinking about physics, biology, psychology, and so on, without explicit instruction in those areas, while our folk judgments are just specific instances of thinking about those bounded sets of phenomena.

The existence of folk sciences and their accompanying folk judgments is well-established. There is an "emerging consensus about the existence of many folk sciences across all cultures that lead both to real successes at understanding the world and to misconceptions" (Keil 2010: 828). Since by and large people are not trained physicists, biologists, or psychologists (let alone polymaths trained in all three disciplines), we do not often rely on our own robust theoretical knowledge of physical objects, plants and animals, and so on when navigating our environments. Instead, we rely on folk science in order to make judgments about these phenomena.[12]

[12] While there is no exact label for this type of thinking, I will continue to use the phrase folk science. Others refer to "intuitive theories" or "naive theories" (see Carey 1985; Carey & Spelke 1996; Slaughter & Gopnik 1996; etc.) Also common within the literature are references to "framework theories" (see Band, Medin, & Atrain 2007).

For example, every human society appears to think about plants and animals in the same intuitive way — which is folk biology. In particular, all societies appear to use roughly the same folk biological taxonomy system: a system of grouping plants and animals into "hierarchically-organized groups, such as white oak/oak/tree or mountain robin/robin/bird" (Medin & Atran 2004: 962). This tendency for taxonomizing plants and animals in this hierarchically organized way appears to be universal (Berlin, Breedlove, & Raven 1973, 1974). And this cross-cultural evidence speaks in favour of the existence of folk biology.

These domain-specific ways of thinking without explicit instruction are helpful for navigating the world, allowing us to infer that if something it is a 'white oak' it is also an 'oak' and therefore also a 'tree'. But folk sciences and the judgments that issue from them also contain numerous misconceptions about the relevant domain. A common example from folk biology: "4- and 5-year-olds typically do not realize that all animals have babies, indicating a concept animal without reproduction as a core property, and a concept baby not tied to the young of each animal species" (Carey 1988: 167). Here the temptation is to say that children's folk biological judgments are mistaken about the domain in question, which shows that formal or theoretical sciences need not simply conserve or preserve the folk judgments. Indeed, sometimes folk judgments are in conflict with each other, as when people judge that on the one hand that all birds can fly, but on the other that penguins are birds.[13] While folk theories often get a number of things right (indeed, they would be poor navigational tools if they didn't get at least *some* things right), they also often get things (sometimes spectacularly) wrong. More accurately, folk theories can lead to *both* "real successes at understanding the world *and* to misconceptions" (Keil 2010: 827, emphasis added).

[13] For more on mistakes in folk biology, see Shtulman & Schulz 2008; mistakes in folk psychology see Malle, Knobe & Nelson 2007; mistakes in folk physics see Bertaminini et al 2004.

What folk science and formal or theoretical science both share is "a common goal of explaining real-world phenomena and of making predictions" (Keil 2010: 834). But where folk sciences and folk judgments struggle is when it comes to articulation. Keil notes that people "are frequently unable to come up with complete explanations of mechanisms, even for surprisingly simple systems" (Keil 2010: 829). Regarding folk mechanics, when pressed we are often unable to explain the mechanical workings of bicycles, toilets, and other basic systems that we may believe ourselves to have an intuitive sense of (Rozenblit & Keil 2010). So the judgments within folk science are often "plagued with the problems of gaps, inconsistencies, and contradictions" (Keil 2010: 830).

However, identifying the borders between folk sciences and formal or theoretical sciences proves to be a difficult task. It is not enough to identify that folk science is incomplete and inconsistent, as formal sciences may also share those features. One thing we can say, however, is that "if there are gaps in formal science, there will be *much larger* gaps in folk science. If there are hunches in science, there be *much wilder* guesses in folk sciences. And if there is a need to defer to others in science, that need may be *larger* in folk science as well" (Keil 2010: 834, emphasis added). So we may distinguish folk sciences from formal or theoretical sciences by noting that the latter are concerted efforts at reducing the gaps, taming the guesses, and minimizing deference within the former.

In part owing to the likelihood of folk sciences being gappy and relying on large hunches there is motivation to move on and articulate a formal theory for certain domains — hence, formal or theoretical sciences. But in what sense does this count as 'moving on' from the folk science? David Braddon-Mitchell suggests that folk theories "are not theories that agents can necessarily write down. Rather, we have "evidence for the [folk or] tacit theory from the *practice* of agents in making their judgments and navigating around the world" (Braddon-Mitchell 2004: 278, emphasis added). The divergence, then, between a folk science and a formal or theoretical science can be marked by an explicit acknowledgment of the commitments of the formal science. Folk sciences are tacit or implicit about their commitments,

since they are just ways of thinking of certain phenomena without explicit characterization about the domains in which these phenomena occur. Formal or theoretical sciences, on the other hand, are much more formalized.

So *implicitness* can be taken as typical of a folk science. Making commitments *explicit* would then represent a move towards a formal science. Why think this represents progress? For one, because when a theory is made explicit the judgments and commitments of that theory can be scrutinized. Do some of our judgments about that domain cohere with our other judgments about that domain? For example, does our folk judgment about bird flight conflict with our folk judgments about which animals (e.g. penguins) are birds? Or, do our judgments fail to capture relevant phenomena? Subjecting these claims to scrutiny assists in our updating and correcting for the common errors of our folk judgments and theories. And the more explicit the commitments are, the more readily such scrutiny can take place.

Though our folk judgments about some domain are not the result of explicit instruction, formal judgments about that domain do in fact require such instruction. And while a folk science that allows us to issue folk judgments is merely implicit, a formal science that allows us to issue formal judgments is necessarily explicit. Owing to the theory being explicit, it is characteristic of a formal theory that its core properties and concepts are available for analysis and negotiation, while the implicit nature of folk sciences renders their accompanying folk judgments unavailable for scrutiny. When they are scrutinized it is an indication of moving from the folk science to a formal or theoretical science. Lastly, for a bounded-set of phenomena, we typically mark the distinction not in terms of referring to the folk science or formal science, but rather by dropping talk of 'formal' entirely: folk biology is contrasted with biology, folk physics with physics, and so on.

My aim here is of not to make the case for the folk/theoretical distinction in general, or for folk biology or folk chemistry (as compared to formal biology and formal chemistry) in particular. Rather, I wish to draw our attention to an extant distinction in the natural sciences that I suggest can be of help in moral philosophy. With this in mind, if there is an

emerging consensus about the existence folk biology, folk chemistry, and so on, might there also be *folk morality*? And if so, what would be the difference between folk morality and its accompanying folk moral judgments and a more *theoretical metaethics* and its accompanying theoretical metaethical judgments?

There has been some limited attention given to the folk/theoretical distinction being serviceable to philosophy in general, and to moral philosophy in particular. Frank Jackson, for one, offers a distinction between *folk morality* and what he calls *mature folk morality* (Jackson 1998). For Jackson, 'folk morality' marks something like the characterization of folk science above: it regards those domain-specific judgments that we implicitly make absent explicit instruction. An example of a folk moral judgment might be of the form 'it would be right to ϕ in circumstance c'. And this folk judgment can have a particular conceptual profile, perhaps something like 'the rightness of ϕ-ing can be true or false'. So we might make a folk moral judgment and also implicitly think that it is truth-apt — just like we might make the folk biological judgment that 'this is a white oak tree' and the accompanying folk biological taxonomic judgment that white oak trees are a type of oak tree, which are a type of tree. In the moral case, both the folk judgment itself and the conceptual profile for that judgment would be 'ways of thinking' about right and wrong action without training in moral philosophy.

On the other hand, 'mature folk morality' would represent an advance on those folk moral judgments, because it springs up after reflection and negotiation regarding our folk judgments. In particular, mature folk morality would be departure from folk morality insofar as it works towards settling the facts (or at least making explicit the commitments) regarding folk moral judgments. In our toy case, mature folk morality is a maturation of folk morality insofar as it is able to explain *whether* or *whether not* moral judgments are truth-apt.[14]

[14] Jackson also offers a more familiar distinction between a folk theory and a mature folk theory: people commonly attribute beliefs and desires to other people, so this is a *prima*

This would of course require deliberate reflection about our moral discourse and practice, in much the same way that a formal or theoretical science is a deliberate examination of the relevant bounded-set of phenomena.

It is important to recognize that it would be a serious mistake to assume that mature folk morality will simply vindicate folk morality. If we have the folk moral judgment that 'the rightness of ϕ-ing can be true or false' we should not think and should not expect that mature folk morality will necessarily confirm that moral judgments are *in fact* truth-apt. To do so would I think radically underestimate the way in which mature folk morality improves upon folk morality. Recall that the move from folk science to formal or theoretical science involves a concerted effort to eliminate gaps in understanding, ironing out or reducing hunches, and so on. If that is right, then we should expect some daylight between folk science and formal or theoretical science. As we saw above, there can be notable inconsistencies among our folk judgments — like the inconsistency between 'all birds can fly' and 'penguins are birds'. To expect formal or theoretical biology to confirm those folk judgments, thus preserving the inconsistency, underestimates both the point and power of the theoretical sciences. Similarly, in the moral case, we should expect that some aspects of folk morality will not be preserved by mature folk morality — and it may be that the supposed truth-aptness of our moral judgments is one such hunch that gets left behind.

For his part, Jackson does not appear to be drawing upon the common folk/theoretical distinction within the natural sciences. And while he doesn't spell out how we should understand folk morality in particular (in terms of, say, the necessity of folk judgments being implicit and absent explicit characterization for the relevant domain), he is nevertheless committed to the existence of folk moral judgments. Indeed, for Jackson such judgments

---

facie plausible commitment of so-called *folk psychology*. However, *formal* or *theoretical neuroscience* (or neuroscience *simpliciter*) would vindicate this folk commitment if it establishes that people really do have beliefs and desires.

are necessary in order to engage in moral philosophy at all. He suggests that our folk moral judgments are those pre-philosophical judgments that are "part and parcel of having a sense of what is right and wrong, and of being able to engage in meaningful debate about what ought to be done" (Jackson 2000: 130). Note how this echoes the earlier characterization of folk sciences, which are said to arise informally and yet are heavily relied up for navigating our environments.

When it comes to folk morality, suppose that we have the pre-philosophical or folk judgment 'that hurting people merely for fun is morally wrong'. This would be a folk moral judgment if we did not have any specific moral theory in mind when holding this judgment (e.g. we did not have it because we thought hurting people reduces overall utility, or violates the inherent dignity of persons, etc.). It would be a folk moral judgment if we did not actually know that there are different moral theories that seek to explain why this judgment is justified (e.g. utilitarianism, deontology, etc.). We would simply have this judgment independent of explicit instruction in moral philosophy. Such pre-philosophical beliefs about right and wrong "make up the core we need to share in order to count as speaking a common moral language" (Jackson 2000: 132).

I suggest that by engaging in moral philosophy we are attempting to bring clarity and rigor to these folk moral judgments. Because they constitute the shared basis for moral discussion, even amongst non-philosophers, our folk moral judgments are what we often appeal to in debating moral matters. We ask and attempt to resolve questions about them. Are any of our folk moral judgments in tension? Which judgments, if any, should we refrain from making? What follows from our moral judgments? And so on. Notice again the similarity between this reflection and negotiation and that of formal or theoretical biology: we start with the bounded set of phenomena (say, plants and animals) and attempt to bring rigor and clarity to our thoughts about that domain by applying the scientific method. And yet, even if we perhaps by necessity appeal to our folk moral judgments when negotiating our lives with others, that does not render our folk judgments incapable of being revised —

indeed, of even radical revision. This is because, due to our engaging in moral philosophy, "folk morality is currently under negotiation: its basic principles, and even many of its derived ones, are a matter of debate and are evolving as we argue about what to do" (Jackson 2000: 132). That is to say, moral philosophy can attempt to offer a corrective or update on our folk moral judgments, just as easily it might attempt to justify and support those judgments. Which, again, neatly parallels the move from folk science to formal or theoretical science.

Thus, the account for folk moral judgments looks similar to the account of folk judgments more generally. However, because Jackson doesn't mark the distinction between folk and theoretical judgments, some work needs to be done to extend the folk/theoretical distinction from the natural sciences to moral philosophy. Here we are helped by the initial characterization seen above. While our folk moral judgments are domain-specific judgments absent explicit instruction, engaging in moral philosophy is itself a way of coming to make theoretical metaethical judgments — and that can only take place after explicit instruction. While our folk moral theories are implicit, our theoretical metaethical theories are necessarily explicit. Because of this explicitness, the core precepts of morality are readily available to reflection, analysis, and updating or correction. In this way, folk moral theories and theoretical metaethical theories share a commonality with folk theories and formal theories more generally. There is no reason to suspect that all aspects of a folk theory will be preserved by the more theoretical one.

And if the folk/theoretical distinction can indeed be extended to moral philosophy, then it may help to explain the NAJ response. The next section explores this possibility.

3.3 Diagnosing the EDA

Here is how the distinction between folk judgments and theoretical judgments can help fill in the details of the NAJ response. Recall that the EDA is supposed to be a threat for our moral judgments. The challenge it presents is to explain why we should think that

any of our moral judgments are true if those beliefs are likely to have been evolutionarily advantageous. If such advantage tracks evolutionary fitness and not necessarily truth, then we have an undercutting defeater for thinking that our moral judgments are in fact true. We would then not be justified in thinking that any of our moral judgments are *successful*, and as a consequence not justified in accepting nonnaturalist cognitivist realism. In response, anti-debunkers have urged that while *some* of our moral judgments are susceptible to this undermining argument, *other* moral judgments are not. This was the point suggested by Parfit, FitzPatrick, and others: through philosophical training we are able to use our evolutionarily developed capacities for moral judgment to have successful moral judgments.

In light of the folk/theoretical distinction, one way to disambiguate our moral judgments is as follows. We could refer to our folk moral judgments, which are those judgments that the EDA most easily targets, or we could refer to our theoretical metaethical judgments, which seem at least in principle capable of bypassing the EDA. The NAJ response might be able to take advantage of the folk/theoretical distinction in the following way: though our folk moral judgments are susceptible to the undermining influence of the EDA, our theoretical metaethical judgments are not. This seems to be what Parfit, FitzPatrick, et al. have in mind when they suggest that philosophical training allows us to use our evolutionarily developed capacities to think about moral matters. With this in mind we can canvass various conclusions of the EDA, and see how the folk/theoretical distinction allows us to make sense of the NAJ anti-debunking strategy.

First, some anti-debunkers may be interpreting the EDA in just this way, as targeting only our folk moral judgments, and as a result they respond by saying that though our folk moral judgments might be suspect our theoretical metaethical judgments are not. If this is how some anti-debunkers are characterizing the EDA then we can call their interpretation:

Folk Judgment Debunking Our folk moral judgments (but not our theoretical metaethical judgments) are suspect because of evolutionary considerations

If this is the conclusion of the EDA then the NAJ response seems effective. Its proponents suggest that our folk moral judgments are not of central importance for vindicating nonnaturalist cognitivist realism. Though some of our moral judgments might be unjustified, insofar as we lack justification for thinking that they respond to the moral facts, we may have other judgments that are not similarly impugned. Since our theoretical moral judgments need not merely be our folk moral judgments (as we saw above), there is a gap between impugning our folk moral judgments and impugning nonnaturalist cognitivist realism itself. Moreover, since our theoretical metaethical judgments are what nonnaturalist cognitivist realism characterizes, characterizing the EDA as **Folk Judgment Debunking** does not undermine nonnaturalist cognitivist realism.

But there is room for the debunker to respond to **Folk Judgment Debunking**. While one can accept that the EDA is *supposed* to target our folk moral judgments, and that it would in principle leave untouched our theoretical moral judgments, this ends up creating the following problem: even if the undermining effect is restricted to our folk moral judgments this would still be significant enough to undercut the plausibility of nonnaturalist cognitivist realism. This is because nonnaturalist cognitivist realism might be an impoverished view if it cannot rely on *any* of our folk moral judgments. Recall Parfit's examples: while evolutionary considerations do seem to undermine some folk judgments, like 'that we have reason to care more about good or bad experiences when these experiences are in the future rather than the past', those same considerations do not undermine other judgments like 'that everyone's well-being matters equally'. However, this leads to an impoverished view of morality. Even if the EDA is best understood along the lines of **Folk Judgment Debunking**, nonnaturalist cognitivist realists would have the following problems: they would not be able to say, for example, that we have reason to care more about future experiences than past ones, nor would they be able to take for granted *any* of our folk moral judgments. In giving up our folk moral judgments nonnaturalist cognitivist realists would be giving up too much. It would lead the moral realist here to have an emaciated view of morality and of our moral

judgments. They could only justifiably accept some claims, such as 'everyone's well-being matters equally', but not seemingly basic claims, such as 'we have more reason to care about future experiences than past ones'.

To this the anti-debunker might respond in the following way: the folk/theoretical distinction does not necessarily rule out the *content* of our folk moral judgments, but rather the way in which one *arrives* at those judgments. Though the EDA might undermine our folk moral judgments, precisely because they are the sort of judgments that evolutionary influences plausibly have exerted a distorting influence upon, our theoretical metaethical judgments might arrive at judgments that are *similar* or *identical* in a justified way. Consider: if we come to learn that a belief is the result of a thoroughly unreliable formation process, that gives us a reason to be sceptical of the truth of that belief. But it of course doesn't show the falsity of that belief. Nor does it preclude our using a more reliable belief formation process to arrive at a similar yet justified belief.

For example, if you learn that that the only reason you believe that Mercury is the closest planet to the Sun was because you were hypnotized to think so, you should not think that your belief is true. You should instead think that your belief is unjustified. But if you later learn, via a more truth-sensitive belief formation process, that Mercury *is* the closest planet to Sun, then you should think that your belief is justified. Your initial unreliability would not necessarily preclude your later reliability.

Nevertheless, the anti-debunker owes us an explanation for why our theoretical moral judgments just so happen to be the same judgments that we would have had anyway (and which were targeted by the EDA). Absent this story the anti-debunker response is not fallacious but it is deeply suspicious. The upshot for the anti-debunker, however, is that if the EDA is characterized as **Folk Judgment Debunking** it is a weak argument, and one easily met by distinguishing between folk and theoretical judgments.

What if, instead, the EDA was best characterized not as targeting our folk moral judgments but instead as targeting our theoretical moral judgments? This would get us the following:

Theoretical Judgment Debunking Our theoretical metaethical judgments (but not our folk moral judgments) are made suspect because of evolutionary considerations

I think we can see that this way of characterizing the EDA would not be of help to a would-be debunker. **Theoretical Judgment Debunking** is at best too quick, and at worst false. Recall FitzPatrick's claim: we are able to "think intelligently and often accurately about things like metaphysical necessity ... the facts of which are equally irrelevant to the etiology of the capacities we use in thinking about them" (FitzPatrick 2015: 887). So it is less plausible that our theoretical moral judgments are the sorts of judgments thought to be distorted by evolutionary influence. And without that influence getting into the picture no damage could be done to nonnaturalist cognitivist realism itself. So, the would-be debunker shouldn't characterize the EDA in this way.

A similarly poor option would be to say that it is not the folk moral judgments or else the theoretical metaethical judgments that are problematic, but rather both. This would lead to the following view:

Folk and Theoretical Judgment Debunking Both our folk moral judgments and our theoretical metaethical judgments are made suspect because of evolutionary considerations

This characterization seems to account for everything, but it begs the question against the NAJ response. The EDA is so far silent on why we should not think that any of our moral judgments are in principle capable of being justified, rather than merely some of our judgments being justified. Distinguishing folk moral judgments from theoretical metaethical judgments via the folk/theoretical distinction has the advantage of allowing for evolutionary distortion, yet preserving the possibility of moral justification. Rolling the two sorts of

judgments together seems to elide that distinction altogether, and we would need a principled reason for doing so.

Canvassing these three weak options seems to show that either the EDA is a fundamentally bad argument, or else its strength lies elsewhere. Anti-debunkers suspect the former. But I wish to explore the latter.

While the folk/theoretical distinction is helpful for understanding the strength of the NAJ response, it also has the unintended but beneficial effect of revealing what I think is the appropriate way to characterize the EDA. Rather than it targeting our folk judgments, or our theoretical judgments, or both, the EDA is better understood as targeting the *inference from* our folk moral judgments *to* our theoretical metaethical judgments. This gives us the following:

Inference Debunking If our theoretical metaethical judgments are inferred merely from our folk moral judgments then those theoretical metaethical judgments are made suspect because of evolutionary considerations

Characterizing the EDA in terms of **Inference Debunking** has the following advantage. It allows would-be debunkers to account for the merits of the NAJ response, while also allowing that there may be metaethical theories that are defeated by the EDA. This would not require a giving up on the EDA, but rather require us to offer a recharacterization of the argument.

Accepting **Inference Debunking** also makes salient that there should be a measure of independence between our folk moral judgments and our theoretical metaethical judgments. This is exactly as things should be. Incorporating the folk/theoretical distinction from the natural sciences shows us that thinking otherwise would underestimate that difference between our folk theories and our more formal or theoretical ones. Just as it would limit the epistemic power of our scientific theories if they were expected to merely confirm all and only our folk judgments about biology, chemistry, etc., we should also think that it limits our theoretical metaethical theories if they merely and only preserve our folk moral

judgments. As we saw above, folk theories often get things wrong, which is part of the reason for developing more formal or theoretical theories in the first place. So if we incorporate the folk/theoretical distinction into moral philosophy then we will have to allow for the possibility that some of our folk judgments are mistaken.

When the EDA is understood in terms of **Inference Debunking** it has the following upshot. Recall that there are at least two ways of understanding the EDA. When it is understood as a *theory debunking argument*, the EDA purports to show that a particular theory fails to justify our moral judgments. One way to do this is to show how a realist interpretation of our moral beliefs is incompatible with our evolutionary history. I noted that Street's EDA is a paradigmatic theory debunking argument (Street 2006; 2008). She does not deny that our moral judgments are justified, rather she denies that the moral realist can offer a plausible account of their justification that is compatible with our evolutionary history.

Conversely, if the EDA is a *justification debunking argument* then it is an argument that aims to show that all attempts at justifying our moral beliefs fail.[15] Joyce's debunking argument is best understood in this way. His EDA concludes that all of our moral beliefs are unjustified, not just that they are so according to a particular account (Joyce 2001; 2006). Justification debunking arguments thus have a wider aim than theory debunking arguments. And because they have fundamentally different sights, and thus prospects for success, which way the EDA ought to be characterized makes an important difference.

According to my recharacterization, the EDA is a very specific type of theory debunking argument. The reason it is a theory debunking argument and not a justification debunking argument is because **Inference Debunking** does not rule out justification of our moral

[15] Importantly, a justification debunking argument differs from a moral error theory because an error theory say that our moral judgments are systematically false. Justification debunking theories can instead grant that our moral judgments might be true, but then stress that even on that assumption such judgments are unjustified.

judgments. Nor does it preclude justification according to some sort of realist account. Because the EDA (as an instance of **Inference Debunking**) is in principle open to some form of justification it cannot be understood as a justification debunking argument. Instead, it purports to undermine a specific way of justifying moral judgments. What it rules out are those metaethical theories whose judgments are justified by inference only from our folk moral judgments.

I hope to have identified a particular type of verboten inference, and in the subsequent chapters we will see if there are metaethical theories that violate it. And while I think which metaethical theories are targeted is an open question, here is an example of the kind of inference that is ruled out by **Inference Debunking**. Suppose that one has the folk moral judgment that 'gender discrimination is wrong', and that accompanying that judgment is the belief that even if no one believed it to be so gender discrimination would *still* be wrong.[16] A metaethicist might point to these judgments and conclude that, therefore, moral facts are mind- and attitude-independent. And it is *this* sort of inference that is ruled out by **Inference Debunking**. The inference *from* the folk moral judgments *to* the theoretical moral judgment is what is ruled out.

Of course, this recharacterization of the EDA does not imply that there are no justified metaethical theories. Nor does it imply that there are no justified moral realist theories. Rather, I hope it is clear that it only blocks *one way* of supporting a metaethical theory in general and a moral realist theory in particular: namely, by straightforwardly inferring the theoretical moral judgments from our folk moral judgments.

[16] There is some evidence from the social sciences that suggests non-philosophers understand moral judgments in this way. See Goodwin & Darley (2008).

3.4 Conclusion

The above offers a way of filling out one common anti-debunker response to the EDA. The NAJ response suggests that, while some of our moral judgments are subject to the EDA and its undermining defeater, we have other moral judgments are not. I have argued that a plausible way of characterizing the difference between which judgments are impugned and which are not is by distinguishing between our folk moral judgments and our theoretical metaethical judgments. Relying on this folk/theoretical distinction also has the upshot of showing EDA to be a very specific type of theory debunking argument. It prohibits just those metaethical theories that rely on theoretical metaethical judgments inferred merely from our folk moral judgments.

The reason for exploring and recharacterizing the EDA is because it is typically thought to be a strong sceptical challenge to any plausible moral epistemology. But, as I noted, it makes a difference whether the EDA is a justification debunking argument or a theory debunking argument. I have argued that it is a theory debunking argument. The remaining question is whether nonnaturalist cognitivist realism is one of those theories. And the simple and perhaps unsatisfying answer is: not yet.

Recall that nonnaturalist cognitive realism is committed only to three claims: cognitivism, nonnaturalism, and success. But it tell us nothing about *which* moral judgments are successful, nor *how* they are successful. Because of this, nonnaturalist cogntive realism cannot be assumed to be one of the theories targeted by the EDA. Rather, in order to answer the *which* question we will need to answer the *how* question: how are our moral judgments justified? And to do so, I address different options in the next three chapters: moral intuitionism, reflective equilibrium, and metaethical Mooreanism.

What I will show is that moral intuitionism and reflective equilibrium, if successful, would not be subject to the **Inference Debunking**. But nevertheless, I will explain why there are problems with both of this views that result in them being of no help to nonnaturalist cognitivist realists. On the other hand, I will argue that metaethical Mooreanism

is particularly vulnerable to the EDA when it is understood as an instance of **Inference Debunking**. Owing to structural features of the position, metaethical Mooreanism is often guilty of inferring formal moral judgments from folk moral judgments. For this reason, I will return to the EDA in Chapter 6.

CHAPTER 4
A Dilemma for Moral Intuitionism

I ended Chapter 3 by claiming that nonnaturalist cognitivist realism is not automatically guilty of inferring theoretical moral judgments from folk moral judgments. Rather, whether or not the view is suspectible to **Inference Debunking** depends on how we think it characterizes our *successful* moral judgments. That is, it depends on how it says that we *know* the moral facts. That means the details of nonnaturalist cognitivist realism's accompanying moral epistemology are critical.

Moral intuitionism is very a prominent moral epistemology. It is an often-discussed account of how we have moral knowledge, with a long and rich history in moral philosophy. It may therefore be interesting to explore the view in it's own right. But moral intuitionism also offers a unique answer to the question of *how* we have moral knowledge — an answer, as we shall see, which is highly congenial to nonnaturalist cognitivist realism's two desiderata: *cognitivism* about moral judgments, and *nonnaturalism* about moral properties.

As I unpack the details of moral intuitionism below it will become clear that the view does not appear to infer theoretical moral judgments from folk moral judgments. This means that it is not subject to **Inference Debunking**. However, showing the details of the view will allow us to see other problems with intuitionism. And it is these problems, I argue, which make the view untenable.

The structure of the chapter is as follows. In section 4.1 I explain why moral intuitionism looks like an appealing moral epistemology. The view promises to explain how our moral judgments connect to the nonnatural moral facts. It thus seems to offer an account of how our moral judgments are successful. In section 4.2 I unpack intuitionism by drawing upon the 19th and 20th century British Ethical Theorists, showing how the view's strength is

its foundationalism. In section 4.3 I highlight a well-known problem with this form of 'old intuitionism', namely that the view is *dogmatic* about which moral propositions are self-evidently justified. Because of this, in section 4.4 I turn to the so-called 'new intuitionists' who have sought to recharacterize the view. They have argued that self-evident justification does not rule out other forms of justification. This *epistemological overdetermination* is supposed to help prevent intuitionism from lapsing into dogmatism.

But in section 4.5 I argue that embracing epistemological overdetermination has a significant cost: it causes new intuitionism to lose its status as a foundationalist epistemology. New intuitionists therefore face a dilemma: without stressing overdetermination, intuitionism cannot explain self-evident propositions non-dogmatically, but in stressing overdetermination, intuitionism loses its claim to foundationalism (and therefore cannot explain our successful moral judgments). Though it escapes the undermining effect of **Inference Debunking**, I conclude that without resolving this dilemma intuitionism is not in a position to support nonnaturalist cognitivist realism. So our search for a plausible moral epistemology goes on.

4.1 An Explanatory Challenge

As a view in moral epistemology, intuitionism can be characterized in two ways: as a *conditional* account, or as a *vindicatory* account. As a conditional account intuitionism is best understood as saying that *if* we have moral knowledge then it is in some way owing to our moral intuitions. But this conditional account is of course compatible without our not having any moral knowledge at all. On the other hand, as a vindicatory account intuitionism is the view that we *do* have moral knowledge because of our intuitions.

It makes a difference which of these characterizations we have in mind. If we are seeking to explain our moral knowledge, then we are attempting to give a vindicatory account. But in what follows, I am not presuming that we in fact have moral knowledge, and so I am not examining intuitionism as such an account. Recall that Chapter 2 gave us two important desiderata: *cognitivism* about moral judgments, and *nonnaturalism* about moral properties.

But from these desiderata alone we do not know whether our judgments are successful or whether they fail. We are not yet convinced of non-sceptical moral realism or epistemic moral scepticism. For this reason, I think we should proceed to examine intuitionism as a conditional account.

Because we are looking at candidate moral epistemologies as conditional accounts, our two desiderata present us with a peculiar explanatory challenge. How do we have manage to know about moral properties if moral properties are nonnatural? In particular, knowing about something nonnatural seems to be unlike our ordinary ways of knowing about natural things (which are readily available to our senses). Though this challenge is peculiar, it is not entirely unique to moral properties. There is a long-standing worry about the supposed link between, on the one hand, our judgments or beliefs, and, on the other, abstracta. We see this cashed out in various areas. Influential ideas from Paul Benacerraf helped to give this challenge its contemporary shape (Benacerraf 1973). For that reason, I will quickly summarize the *Benacerraf Challenge* in the philosophy of mathematics before moving on to a structurally similar challenge within metaethics. It is this challenge that we will want a plausible moral epistemology to meet.

Suppose that certain mathematical entities (like numbers) exist independently of us: that is, they exist apart from our thoughts, speech, and practices about or with them. This is what mathematical platonists believe.[1] Mathematical platonism shares something important with nonnaturalist cognitivist realism, because both are committed to the existence of abstract entities — mathematical entities on the one hand, moral properties on the other. Philosophers of mathematics have long been interested in how it is that we can come to know about mathematical entities, if these entities ought to be construed as the platonist says. This again is unlike our ordinary ways of knowing about natural things. Benacerraf says that

[1] Some prominent defenders of mathematical platonism: Frege 1884 and 1893; Godel 1947; Quine 1948; Wright 1983; Hale 2001; Colyvan 2001; etc.

"something must be said to bridge the chasm, created by . . . [a] realistic . . . interpretation of mathematical propositions, between the entities that form the subject matter of mathematics and the human knower" (Benacerraf 1973: 675).

The reason there is a chasm here is because of the unusual profile of mathematical entities. Because mathematical platonists say these entities are abstract they inherit the following conceptual profile: abstract entities (and thus mathematical entities) are *causally inert* or *causally inefficacious*. Indeed, 'causal inertness' or 'causal inefficaciousness' is the standard way of drawing the abstracta/concreta distinction.[2] In a loose sense, to say that something is causally inert is to say that it cannot make anything happen. To get a grip on this notion, compare abstract entities to concrete ones like chairs. Chairs *do* have causal powers and they *can* make things happen: they can support you when you sit on them, they can physically stand in your way, they can be stacked with other chairs, they make sounds and noises when they move, and so on. And because chairs are like this, we can *know* about them. We can bump into them, we can see them, we can feel them — in short, we can interact with them.

But abstract entities in general, and mathematical entities in particular, are not like this at all. We cannot see an abstract object like '2' — rather, we can see only particular *instantiations* of it (e.g. two chairs). We cannot bump into numbers, or feel them, or smell them. They are, instead, causally *inert*. So how do we manage to *know* anything about them? Something must bridge the chasm here, between abstract objects and our knowledge of them. In light of this the following question arises: how is it that our beliefs about numbers even remotely resemble what numbers are actually like? Note that we cannot say that numbers *just are* whatever we take them to be. Mathematical platonists deny this, because they take mathematical entities to be importantly *independent* of us. So the challenge is for

[2] David Lewis (1986) calls this "the Way of Negation", which says that an object is abstract if and only if it is causally inefficacious. See also Burgess & Rosen (1997).

mathematical platonists "to provide an account of the mechanisms that explain how our beliefs about [mathematical objects] can so well reflect facts about them" (Field 1989: 26).

My interest of course is not with mathematical platonism and the Benacerraf Challenge, but with a structurally similar challenge in metaethics. Chapter 2 set *cognitivism* and *nonnaturalism* as two crucial desiderata for any plausible moral epistemology, so if we have moral knowledge then it is because we have successful judgments about nonnatural moral facts. But how? The same chasm opens up here. If moral facts are nonnatural how do we manage to know anything about them?

Metaethicists have recently clued-in to this challenge.[3] Given our two desiderata, the challenge for nonnaturalist cognitivist realism is "to explain how it would be anything more than chance if [one's] moral beliefs were true, given that [one cannot interact] with moral properties" (Huemer 2005: 123). Similarly, "if normative beliefs are about an independent order of normative facts, how is it that we know anything about them?" (Enoch 2010: 143).

We could understand this as an ontological challenge. That is, if we cannot know anything about the purported moral facts, then we should not think that they exist. But as I mentioned in Chapter 1 am not exploring ontological challenges here. For one, I have little to say regarding the price of admission (as it were) into one's ontology. Maybe we should accept something like the *Eleatic Principle*, which "justifies belief in only those entities to which causal power can be attributed" (Colyvan 2001: 39). This is an epistemological principle of sorts (since it refers to 'justified belief'), but it nevertheless has an ontological aim: which things are we justified in believing *exist*. But the Eleatic Principle requires a defense which I am not in a position to provide.[4] So I set aside ontological matters in what follows.

[3] See for example Clarke-Doane 2015 and 2017; Crisp 2006; Enoch 2009; Leibowitz and Sinclair 2016; Street 2016; etc.

[4] For defenders of the Eleatic Principle, see Armstrong 1978; Ellis 1990; Field 1989; etc.

Rather, I will understand the challenge as an epistemological one. Let's assume that there are moral facts. Moreover, let's assume that these facts should be characterized non-naturalistically. Now, why are we justified in thinking that our moral beliefs link up to the moral facts? How could we know that this is the case? If we have an answer here then we have an explanation for why we can think that our moral judgments are sometimes *successful*. What we are looking for here is a plausible moral epistemology, rather than a lucky one. And if we can answer this question then we can help vindicate nonnaturalist cognitivist realism on epistemic grounds.[5]

4.2 What is Intuitionism?

Intuitionism is one way of meeting the above challenge. With the renewed interest in nonnaturalist cognitivist realism we are seeing a renewed interest in moral intuitionism. The two are importantly distinct. Nonnaturalism is about moral metaphysics and ontology — it's about characterizing objects that exist, while intuitionism is a moral epistemology — it's about how we know about those objects. However, when shopping for a moral epistemology many nonnaturalists tend to favour moral intuitionism. This is because intuitionism promises to bridge the chasm between our moral beliefs and the nonnatural moral facts. And if moral intuitionism succeeds here then the plausibility of nonnaturalist cognitivist realism waxes. This would allow us to justifiably think that some of our moral judgments are successful. But if moral intuitionism does not succeed, then the plausibility of nonnaturalist cognitivist realism wanes. This again would not show that the moral facts are not out there, but it would give us reason to think that they may be inaccessible to us. And, in turn, epistemic moral scepticism would get a boost.

[5] But note that if we cannot answer these questions it does not follow that we should accept moral error theory. After all, there may in fact be moral facts, but we do not seem to be in a position to know anything about them. This is why failing to find a plausible moral epistemology results in epistemic moral scepticism rather than moral nihilism.

Ideas of an intuitionist bent have been around for a while.[6] So much so that it is difficult to speak with any sort of authority on the issue. But as a distinct epistemological view in moral philosophy intuitionism begins to take shape in late 19th and early 20th century Britain. It's there that we see a clearly delineated intuitionist epistemology starting with Henry Sidgwick (1874), and then taken up by G.E. Moore (1903, 1912), W.D. Ross (1930, 1939), H.A. Prichard (1912), and others. Call views from this general era *old intuitionism*. Recent developments of intuitionism have come from Robert Audi (2004, 2008), Michael Huemer (2005), Peter Railton (2014), Russ Shafer-Landau (2003), and Philip Stratton-Lake (2002). Call these latter views *new intuitionism*.

Both old and new intuitionism are accounts of when and how our moral beliefs are justified. Because intuitionism is an account of epistemic justification we need to be clear about the *type* of justification on offer. In what follows there are going to be two important senses of justification. The first is *propositional* justification. This is justification that one can have for a proposition. A moral proposition like 'torturing people merely for fun is wrong' is propositionally justified when an agent would have good reasons, grounds, or evidence for believing that proposition. One thing to note here is that it is irrelevant whether any actual agent *in fact* believes the proposition. Instead, propositional justification is about the reasons or evidence for a particular claim, which comes apart from anybody actually

[6] As a moral epistemology intuitionism stretches at least as far back as Plato, where the allegory of the cave suggests that what we are able to see by daylight is not as reliable as what we can come to see by the light of reason (*The Republic* Bk. VII). Much later, Spinoza gives a detailed treatment of intuitive knowledge (*Ethics* Bks. II and V). Descartes characterizes intuition as 'the natural light', saying that "things that are revealed by the natural light ... are not open to any doubt, because no other faculty that might show them to be false could be as trustworthy as the natural light" (*Meditation* III).

believing the claim.[7] So part of what we are after is whether there are good reasons to believe any moral propositions.

Now suppose we establish that one has propositional justification for some moral proposition — say, that 'torturing people merely for fun is wrong'. Even if there are good reasons or evidence for believing that claim, that *alone* would not entail that any *particular* person is thereby justified in believing it — just that one *could be* justified in believing it. If we want to know whether or not someone in particular is justified in believing a given proposition then we are typically looking for *doxastic* justification. Unlike propositional justification, doxastic justification is justification that concerns agents and their particular beliefs (rather than the reasons or evidence for certain propositions themselves). We can say that a person is justified in believing a moral proposition if that person is in the right sort of relation to the evidence for that proposition. So, how did the person come to form their belief? Does their belief stand in the right of epistemic relation to their evidence or to their other beliefs? These are questions that illuminate whether one is doxastically justified in believing a moral proposition.[8]

In addition, the *relationship* between propositional and doxastic justification is relevant. Epistemologists typically think a person can only be doxastically justified in believing a proposition if one can be propositionally justified in believing it. This is the orthodox view, which says that "if p is propositionally justified for S in virtue of S's having reason(s) R, and S believes p on the basis of R, then S's belief that p is doxastically justified" (Turri 2010: 313).[9]

---

[7] For more, see Pryor 2004: 351 — 355. Pryor writes that "justification is a quality that *hypotheses* can possess for you: they can be hypotheses you have *justification to believe*. It doesn't matter whether *you do* believe them. This is sometimes called 'propositional justification'." (Pryor 2004: 352).

[8] Exactly what these right states or relations are depends on your preferred theory of doxastic justification. Think here of debates between *internalists* and *externalists* about doxastic justification. See Goldman 2009.

[9] See also Feldman 2002; Korcz 2000; Kvanvig 2003; Pollock 1986; etc.

So that makes doxastic justification *dependent* on propositional justification. However, John Turri argues against the orthodox view (Turri 2010). Giorgio Volpe in turn argues against Turri (Volpe 2017). The actual relationship between propositional and doxastic justification is something that we can set aside. Instead, we will look at how intuitionists *themselves* think of these issues.

While intuitions are the stock-in-trade of philosophy, there is a lot of disagreement about what they are.[10] Accounts of moral intuitions inherit that disagreement. Nevertheless, two contemporary rival accounts have emerged: doxastic (or self-evidence) accounts and perceptual (or seeming state) accounts. On the one hand, doxastic accounts say that our moral intuitions are beliefs or adequate understandings of self-evident propositions (Audi 2004, 2008). On the other hand, perceptual accounts say that our moral intuitions are intellectual seemings or quasi-perceptual states (Huemer 2005). The difference is roughly that doxastic accounts say that intuitions are beliefs, while perceptual accounts say that they are not beliefs. To get a grip on this, think of a non-moral case, like the experience of looking at a red apple. Some people might think that by looking at the red apple you *believe* that the apple is red. Others might say that you are in a state where it *seems* to you that the apple is red, but you don't quite believe it. Same goes for the moral case: when we have a moral intuition (like in a trolley case), doxastic theorists think you are thereby having a moral belief, while perceptual theorists think you are having an intellectual seeming that isn't quite a belief.

What I will go on to say does not depend on whether intuitions are beliefs or not. Rather, my interest is in how intuitionists think that our intuitions *justify*. And both the doxastic and perceptual accounts are united in thinking that intuitions justify in a *non-inferential* way. This makes intuitionists *foundationalists* regarding epistemic justification. As a view

[10] See Bealer 1998; Bengson 2015; Chudnoff 2011; De Cruz 2007; Goldman 2007; Nagel 2007; Sosa 2007; Williamson 2004.

in epistemology, foundationalism says that all our justified beliefs ultimately ground out in non-inferentially justified beliefs. So all beliefs are justified either *inferentially* (via other beliefs) or else they are justified *non-inferentially* (via themselves). And the key claim of foundationalism is that any inferentially justified belief will in fact have an epistemic route back to a non-inferentially justified belief. So if you think that we have some beliefs that are non-inferentially justified, and if all other inferentially justified beliefs depend on the non-inferentially justified ones, then you are a foundationalist. Intuitionists *do* think this (as we'll see below). So intuitionists are epistemic foundationalists.

A helpful way of talking about the distinction between inferential and non-inferential justification is in terms of *transmission*. For inferential justification, we can think of the way justification can transmit as if justification were like electricity. A laptop might have power because it is connected to its power cord, and that power cord might have power because its connected to an extension cord, and the extension cord might have power because its connected to an outlet, and so on. In this way, the power is transmitted from an initial source (whatever that is) to the laptop. A similar story holds for inferential justification. Justification *moves through* beliefs or propositions the way electricity moves through outlets and extension cords to the laptop — from some source to some final point.

But non-inferential justification is a much stranger beast. It would be like your laptop having power *all by itself*, and not because that power was transmitted from somewhere else. Similarly, a belief in a proposition would be justified not because that justification is transmitted from elsewhere (like some *other* belief), but rather because the proposition is justified all by itself. What makes foundationalism unique is its insistence that propositions can in fact be justified in this 'all by itself' way. These are our foundationally justified beliefs.

Non-inferentialism gets us part of the way to foundationalism. We complete the journey by supposing that *every* inferentially justified belief needs some foundational non-inferential justification. Moral intuitionism, then, is a version of foundationalism because it holds that belief in some moral propositions can be non-inferentially justified, and that even if there

are inferentially justified beliefs in moral propositions this justification has a route back to some non-inferentially justified belief.

Two questions arise here. First, why is foundationalism a feature rather than a bug? And second, how do intuitionists characterize non-inferential justification? I address these in turn.

First, foundationalism is appealing because having a non-inferential basis of justification seems necessary for avoiding certain epistemic problems. I have in mind here problems like infinite regresses, 'mere' appeals to coherency, and so on. Regress problems generate in the following way: if our moral beliefs are justified only in virtue of some more basic belief, then we may never be able to find justificatory bedrock. Consider that, if justification is indeed transmitted then it must transmit from *somewhere*. But if there is no compelling story about how the justificatory ball gets rolling in the first place then we seem left with a partial picture at best. Without a basis of non-inferential justification it is difficult to make sense of that transmission getting started. Similarly, in regards to 'mere' coherency worries, without non-inferential justification our moral beliefs would have to be justified inferentially (if at all). The worry here is that this makes our moral beliefs 'merely' consistent with each other, but does not give us reason to think that a coherent set of beliefs is justified. Worse, that coherency does not seem to promise any connection to the moral facts. So in light of these worries, accepting foundationalism looks appealing.[11]

We can pause here to note that appealing to a non-inferential basis of justification has the following upshot: intuitionism is not guilty of inferring theoretical moral judgments from folk moral judgments, and so is not susceptible to **Inference Debunking**. Indeed, as we will see below, the judgments that the view *does* rely on are not folk judgments at all, but

[11] This treatment of coherentism is admittedly truncated, but I turn to coherentist accounts of justification in the next chapter.

rather judgments that have a unique and peculiar characteristic. This characteristic takes us to our second question: how do moral intuitionists characterize non-inferential justification?

Sidgwick says that some moral propositions "present themselves as self-evident" (Sidgwick 1874: 383). Similarly, Ross says that we can see that some propositions "are evident without any need of proof, or of evidence beyond [themselves]" (Ross 1930: 29). But what does it mean for something to be *self-evident*? Explaining self-evidence requires a more extended discussion.

Old intuitionists often draw a similarity between moral propositions and mathematical propositions.[12] By comparing these we can get a grip on what self-evidence amounts to for intuitionists. Here is one example of a purportedly self-evident moral proposition:

Future I ought not to prefer a present lesser good to a future greater good

Sidgwick says that **Future** is self-evident. This proposition is self-evident because it is capable of being understood merely in terms of its constitutive components. If you know what a present lesser good is, and you know what a future greater good is, then you can thereby know that you ought not to prefer a present lesser good to a future greater one.

Compare a moral proposition like **Future** to the following mathematical proposition:

Equal If equals be added to equals then the wholes are equal

For Sidgwick, the similarity between **Future** and **Equal** is that both are self-evident. They do not appear to call for justification in virtue of anything beyond themselves. Rather, we can apparently see that these propositions are justified by their own lights: if you know what

[12] However, even though he thinks there is a similarity between moral and mathematical propositions, Sidgwick nevertheless criticizes Samuel Clarke's "anxiety to exhibit the parallelism between ethical and mathematical truth" (Sidgwick 1874: 358). Sidgwick's complaint here is that insisting on the similarity between the two leads Clarke to "extravagancies and absurdities" (Ibid.). So we should not read Sidgwick as saying that moral and mathematical propositions (even of the self-evident stripe) are *identical*, just that they are *similar*.

equals and wholes are, and if you know what addition is, then you know why if equals are added to equals the wholes are equal.

We should be clear that it is not the *consistency* or *coherency* of these (or similar) propositions that intuitionists find compelling. It is not the fact that **Future** or **Equals** might get along with other things that we happen to believe about morality or mathematics, respectively. Justification does not get into the picture in that way. Rather, what makes these propositions unique is that they have justification in virtue of themselves.

Moreover, Sidgwick stresses that while both **Future** and **Equal** seem self-evident, they are both *in fact* self-evident.[13] We need to draw this distinction out. A proposition can *seem* to us self-evident without *actually* being self-evident. Sidgwick calls these "sham-axioms" (1874: 374). Here is an example of one such sham-axiom:

Speak I ought to speak the truth.

This moral proposition might look to us as self-evident as **Future**. That is, it might strike us that **Speak** requires no evidence beyond itself. But Sidgwick does not think that this proposition is self-evident. Why not? Because when we attend to this proposition we recognize that we cannot think that **Speak** is evident merely in terms of its constitutive components. Knowing what speaking is, and knowing what truth is, will not thereby clue us in to why we ought to speak the truth. So if belief in the proposition is indeed justified (as it might in fact be), that justification will have to come from elsewhere. And that's just another way of saying that **Speak** is not self-evidently justified, but rather is inferentially justified.

So the difference, then, between actually self-evident propositions like **Future** and **Equal**, and a sham-axiom like **Speak** is this: in order to be justified, only the latter is

[13] Of course, this implies that Sidgwick is a non-inferentialist about both morality and mathematics. For a discussion of the relationship between one's philosophical commitments concerning morality and mathematics, see Justin Clarke-Doane (2014, forthcoming).

going to require "'rational justification' of some kind" (Sidgwick 1874: 383). If we ultimately accept **Speak** then it will have to be because it admits of rational justification, and not because it is self-evidently justified. It might instead be justified by means of inference from something else we believe, or via coherence between that proposition and others, or via the instrumentality of believing that proposition, and so on. There are of course many forms of justification, and the old intuitionists are not suggesting that all justified beliefs are self-evident. Rather, they are arguing that *some* propositions are justified because they are self-evident, in particular key moral ones.

The takeaway from this is not that we *should not* accept **Speak** (nor, conversely, that we *should* accept it), but instead that we cannot accept it without our providing justification for it. That makes it importantly different from a proposition like **Future**, which provides its own justification. Self-evident justification is justification that can be had independently of inference, or coherency, or instrumentality, and so on. More precisely, self-evident justification can be characterized as meeting the following three conditions:

Non-inferential the justification is not inferred from other propositions
Non-derivative the justification is not derived from arguments
Direct the justification issues directly from the proposition in question

And just as we asked this about foundationalism, we can ask why self-evidence is a feature rather than a bug. Why should we think that some moral propositions can be plausible characterized in this way, as capable of **Non-inferential**, **Non-derivative**, and **Direct** justification?

There are two reasons here. First, the intuitionist stresses that we should think these things because some moral propositions seem knowable (and our belief in them justified) merely by understanding the constituent components of the proposition itself. Consider that if we can appreciate how a mathematical proposition like **Equal** admits of self-evident justification, then we are in a good position to appreciate why some moral propositions

could be capable of similar justification. This was the reason that Sidgwick expressly compares mathematical and moral propositions, in order to motivate our thinking that moral propositions can indeed be self-evidently justified.

More importantly, the second reason that intuitionists find self-evidence attractive is this because it allows them to address the challenge identified at the outset of this chapter. If the moral facts are causally inert, and so our ordinary ways of knowing about them (e.g. by bumping into them, as it were) are unavailable, how do we manage to know anything about them? Appealing to a notion of self-evidence appears to dissolve this challenge. Moral intuitionists explain that we can have justified moral beliefs because some moral propositions present themselves to our awareness as self-evident. This means that a proposition does not need justification or evidence beyond itself. These non-inferentially justified propositions serve as the foundation of some of our moral beliefs. Thus, by intuiting these self-evidently justified moral propositions we can have justified moral beliefs.

4.3 Problems for Intuitionism

While intuitionism promises to explain how our moral beliefs link up or connect to the moral facts, nonnaturalistically construed, the view is not without its own unique problems. Some of these problems I believe we can set aside. I will mention and dismiss a few problems for intuitionsim before moving on to a much more serious challenge.

Throughout the history of metaethics there are worries expressed about the mysteriousness of 'intuition'. It has sometimes been said that intuition, as a doxastic faculty, is not like "our ordinary counts of sensory perception" (Mackie 1977: 38). This purportedly amounts to a challenge to intuitionism, and it is usually chalked up to J.L. Mackie and other moral nihilists who claim that there are no moral facts. The worry is that if the only way we can know about moral facts is via some mysterious faculty of intuition, then we might as well dismiss both the moral facts and the mysterious faculty by which we purportedly come to know about them.

But there are a few problems with this 'mysteriousness' worry. First, it will not come as a surprise to intuitionists that intuition is not an 'ordinary' mode of sensory perception. After all, the view is specifically contrasted *against* sensory perception. It therefore begs the question if we discount intuitionism on that basis alone. While some intuitionists say that we can 'see' that a proposition is self-evidently justified, this parallel with sensory perception is best understood as a rhetorical flourish — which is why that view of intuitions is often called the *quasi-perceptual* account. The important point for intuitionists is just that some justification is said to be **Direct** — i.e. issuing straight from the proposition in question.

Second, Mackie (for one) is clear that he is not picking on any 'intuitive' mental faculty in general. Rather, he points out that "there is no need to think of this 'special faculty' [i.e. intuition] as a separate organ of the mind ... [rather] all that is required is that the understanding, the aspect of our minds which can discover objective truths, has the ability to discover, among others, this particular sort of necessary truth" (Mackie 1980: 147). So for Mackie there is nothing 'mysterious' about intuition as a doxastic faculty. The 'faculty' of intuition is simply that mental ability that we have for arriving at objective truths in general.[14]

And third, even if it is not Mackie's line of argument, worries about 'mysterious' intuitive faculties are not persuasive in their own right. Recall that intuitionists are keen to point to similarities between self-evident moral propositions and self-evident mathematical propositions. Because of that, it seems that intuitionists need only posit a faculty as 'mysterious' as that which allows us to arrive at certain mathematical truths. If one grants that we *can* arrive at such truths, then the case for specifically moral intuitionism seems to receive a boost.

[14] Note the similarity between what Mackie says here and the faculties that FitzPatrick (in Chapter 3) describes regarding our ability to know metaphysical or modal truths.

A separate worry that one might have is that moral intuitionism is unable to establish that there *are* moral facts. The thought might be that the intuitionist has said nothing about why we should think there are such facts. But this line of argument is also confused, for two reasons. First, it mistakes intuitionism, which is an epistemological account, for some ontological or metaphysical account. Recall that intuitionism is simply supposed to explain how we can *know* about moral facts, not how we can demonstrate that they exist. And second, recall that I proceeded here by taking intuitionism as a *conditional* account: i.e. *if* we have moral knowledge, etc.

While none of the those lines of criticism seem to me persuasive, there is I think a deep problem for old intuitionism. What if you think that a proposition is self-evident whereas I do not? For example, you and I might disagree over whether **Future** is self-evident. The problem here is not necessarily that one of us *accepts* **Future** while the other does not. Rather, the problem is that we disagree with each other about the way in **Future** gains its unique justificatory status. Is it self-evident or not? Disagreements of this sort are easy to multiply. What if we both disagree with Sidgwick that **Future** is self-evident? Or what if you insist that **Speak** is in fact self-evident and not a sham-axiom?

This is the problem of *dogmatism* for moral intuitionism — it is a problem regarding whether or not some moral claim is self-evidently justified. Recall that self-evident justification is essential to intuitionism. If some moral propositions are justified in this way — i.e. non-inferentially — then intuitionism has the conceptual resources to bridge the chasm between our moral beliefs and the nonnatural moral facts. If nonnatural moral facts are causally inefficacious then it is not clear how we can know about them. It cannot be due to our ordinary modes of sensory perception, which are inferential. Intuitionists appeal to non-inferential justification in general and self-evident justification in particular to meet this challenge. So intuitionism is in trouble if cannot adequately explain which propositions are self-evident. Being dogmatic and insisting that some propositions just are self-evident does not help meet the challenge.

Old intuitionists attempted to meet this challenge. Some will simply double-down on the appeal to self-evidence — some claims simply *are* self-evident (e.g. Prichard 1912). Others, like Ross, are more forthcoming, saying that once one has "reached sufficient mental maturity" and has "given sufficient attention to the proposition" then he or she will be able to see that a proposition is "evident without any need of proof, or of evidence beyond itself" (Ross 1930: 29). This of course shifts the explanatory burden to the notions of 'sufficient mental maturity' and 'sufficient attention'. But, unfortunately, Ross does not expand on what he means by these notions. We can ask: is sufficient mental maturity marked by apprehending that some moral proposition is self-evident? If so, that still does not allow us to explain in a non-dogmatic way how it is that we know that our moral judgments are self-evidently justified. So dogmatism remains a problem. It is noteworthy that even defenders of intuitionism recognize this charge. Audi writes that "if self-evident truths are conceived as not only not in need of external evidence but also necessary, it is understandable that intuitionism might be considered dogmatic" (Audi 2004: 66).

Often times the problem of dogmatism is actually a misstated objection regarding disagreement. For example, J.R. Lucas asks "If I was arguing with a man, and he did not allow that causing pain was a reason for an action's being wrong, that is, he did not see the relevance of the fact that the action caused pain, I think I should break off the argument with him ... the problem is what weight to attach to such considerations ... sadists ... need to be cured rather than convinced." (Lucas 1971: 9 - 10).[15] Critics of intuitionism might hear such responses and think of them as dogmatic. But the problem here is just disagreement (or dogmatism, if you prefer) over some first-order moral claim, and not disagreement or

[15] See also Audi 2001: "[M]any philosophers regard intuitionism as dogmatic or consider it inadequate because it lacks a comprehensive moral theory as a basis for its disparate principles. It also shares with virtue theories - and arguably with other rule theories - great difficulty in providing a way to resolve conflicts of duties, those 'knotty points' in ethics, as Mill called them, that are so salient in practical ethics." (Audi 2001: 601 - 602)

dogmatism about whether such claims are self-evidently justified. I suspect that *all* moral epistemologies will have to contend with real or imagined first-order moral disagreement. So that should not be viewed as a strike against intuitionism in particular.

That said, the problem of dogmatism about self-evident justification and the problem of disagreement over first-order moral claims might be related. If you are disagreeing with someone regarding some first-order moral claim then you are disagreeing about whether that claim is *true*. But if the explanation for why one should think that it is true is that it is *self-evidently justified* then disagreement and dogmatism are related. Without a way of adjudicating disagreement, a view like intuitionism can appear to be dogmatic.

Nevertheless, the problem of dogmatism is more pressing for present purposes. Because self-evident justification is what allows intuitions to bridge the chasm between our moral beliefs and the nonnatural moral facts, if they cannot adequately explain *which* moral propositions are self-evident then their view is of no help to nonnaturalist cognitivist realism. If we cannot account for how we know that a moral proposition is self-evidently justified, then we cannot establish that we know the moral facts. That is, we cannot claim that we have *successful* moral judgments. And it is here that old intuitionism struggles. From the above, we cannot construct a plausible way to adjudicate which claims are self-evident. For this reason, intuitionism was often thought to be an antiquated view most often associated with the 19th and 20th century British Ethical Theorists. But in the next section, we will see that the view is far from dead. Indeed, it may have all along contained the conceptual resources to answers charges of dogmatism.

4.4 New Intuitionism

Recent defenses of intuitionism have taken up its general strategy and brought newfound rigour to the view. What we can call *new intuitionism* is an account which accepts the bulk of old intuitionist claims, but which also buttresses it with new conceptual and argumentative

resources. Like old intuitionists, new intuitionists avow that the view is *foundationalist*. And likewise, they claim that certain moral propositions are *self-evident*.

Regarding foundationalism, Roger Crisp says that moral intuitions are "non-inferential, self-evident beliefs that certain actions, rules, or whatever are right or reasonable" (Crisp 2002: 65). And Michael Huemer claims that "an intuition that p is a state of its seeming to one that p that is not dependent on inference from other beliefs and that results from thinking about p, as opposed to perceiving, remembering, or introspecting" (Huemer 2005: 102). The appeal to non-inferential justification is paradigmatically foundationalist.

Similarly, like their ancestors, new intuitionists say that certain moral propositions are self-evident. Perhaps the most influential voice among new intuitionists, Audi writes that intuitionists believe a proposition is self-evident when, if one understands it then "in virtue of having that understanding, one is justified in believing that proposition" (Audi 1999: 206). Phillip Stratton-Lake characterizes intuitionism as the view "that certain basic moral propositions are self-evident, and thus can be known directly by intuition" (Stratton-Lake 2002: 18). And Simon Kirchin says that intuitionism is the view "that certain ethical beliefs can be known, and are known, self-evidently" and that "merely by reflecting on the belief at issue and understanding it, one can gain ethical knowledge" (Kirchin 2015: 168).

To be sure, all of these people have distinct views.[16] However, taken as a homogeneous group, these new intuitionists accept the key features of old intuitionism, namely its foundationalism and its notion of self-evidence. But they also do more than merely dust off an old moral epistemology. They have also taken steps to polish the view, by making explicit something that had previously been implicit (at best) or absent (at worst) in old intuitionism. In particular, new intuitionists take seriously the problem of dogmatism. And they seek to recharacterize the view in order to avoid that charge.

[16] As noted above, some new intuitionsts provided doxastic accounts, while others provide perceptual accounts.

Towards that end, new intuitionists are careful to highlight an important feature of moral propositions. As we saw, old intuitionists argued that some moral propositions are self-evident. And it was precisely this feature that generates the aforementioned problems (i.e. how does one *know* that a proposition is self-evident?). New intuitionists also claim that some moral propositions are self-evident - indeed, they must do so in order to count themselves as intuitionists. Yet they go on to clarify the nature of self-evidence so as to meet charges of dogmatism.

New intutionists stress that belief in some moral propositions can be justified in multiple ways. In particular, even those moral propositions which are self-evident can admit of alternative or supplemental justification. Stratton-Lake says that "there is nothing about the fact that a proposition can be known solely on the basis of an understanding of it that rules out the possibility that it can be known in some other ways as well" (Stratton-Lake 2002: 21). Similarly, Kirchin says that "intuitionists should not think that ethical beliefs that can be justified self-evidently can be justified only self-evidently" (Kirchin 2015: 169). Rather, intuitionists should commit only to saying "that self-evident justification is the *prime way of justifying*, at least when considering those ethical beliefs that sit at the centre or which form the foundation of our moral outlook" (Ibid: 170). This is a critical development.

It is important to be clear about this new intuitionist claim. The point they are making is not that, of all moral propositions, *some* can be justified self-evidently while *others* admit of alternative forms of justification (e.g. inferential or derivative justification). As we saw above, that much was recognized by old intuitionists, who said that self-evident justification was simply the prime way justifying moral claims. Rather, the new intuitionists are claiming the following: of those propositions that are justified self-evidently, *even those* propositions can be justified in other ways. That is an important development on old intuitionism, and it is what makes new intuitionism 'new'.

Stratton-Lake calls this feature of self-evident moral propositions *epistemological overdetermination* (Stratton-Lake 2002). The novel new intuitionism claim is that in principle there

can be multiple ways of justifying a proposition that is also taken to be justified self-evidently. If so, then a proposition that is non-inferentially justified in that way is not thereby precluded from also being inferentially justified. As an example, if **Future** is justified self-evidently — that is, in a **Non-inferential**, **Non-derivative**, and **Direct** way — then epistemological overdetermination entails that it can also be alternatively or supplementally justified - say, in an inferential, derivative, and indirect way. While we might be justified in believing **Future** because it is self-evident, we also might be justified in believing it for non-self-evident reasons.

The new intuitionist appeal to epistemological overdetermination is an attempt to answer charges of dogmatism. Recall the problem: what can we say to an interlocutor who just doesn't agree that a certain moral proposition is self-evident? New intuitionists argue that one need not be dogmatic here, because there are alternative forms of justification available even for moral propositions that are self-evidently justified. That is precisely what it means for a proposition to be epistemologically *overdetermined*. So it is not the case that there is nothing to be said to a stubborn interlocutor who does not appreciate the apparent self-evident justification of a moral proposition. One can instead lead that person through an alternative account of justification.

As we saw above, Sidgwick says that a self-evident proposition does not *require* rational justification. Similarly, Ross says that a self-evident proposition does not *require* evidence beyond itself. So one might be tempted to read the old intuitionists as endorsing a notion of epistemological overdetermination. However, what is clear is that the old intuitionists did not explicitly develop that notion. But new intuitionists do just that. They explain how, even though self-evident propositions do not *require* rational justification of further evidence, we can still *provide* rational justification or further evidence for them anyway. If this is right, then new intuitionists can adequately address charges of dogmatism. Where the appeal to self-evident justification fails to convince someone, one can turn to alternative (i.e. inferential, derivative, etc.) forms of justification.

It is worth looking at at this process in action. Stratton-Lake shows that Ross himself wavered over whether the following proposition was self-evident:

Pleasure Pleasure is intrinsically good

In *The Right and the Good* (1930) Ross accepts **Pleasure**. More importantly, he took this moral proposition to be self-evident. But less than a decade later in *Foundations of Ethics* (1939) he ends up denying both: Ross no longer accepts that **Pleasure** is self-evident, nor does he accept it *at all*. So why the change?

Remember that Ross claims we can come to see that a proposition is self-evidently justified once we have "given sufficient attention to the proposition" (Ross 1930: 29). The same appeal to 'sufficient attention' can also reveal that a proposition is *not* self-evidently justified. That is precisely why Ross later came to reject **Pleasure**. For Ross, denying **Pleasure** years after accepting it was a result of him giving increased scrutiny to that proposition. Thus, 'sufficient mental attention', which I had above characterized as too vague to be helpful, now seems best characterized as *doing moral philosophy*. It involves one thinking through moral questions, subjecting them to continued and increased scrutiny, forming arguments for positions, and so on.

Now, this shows that one can *reject* that a moral proposition is self-evident by closely scrutinizing the proposition. But that is not exactly what new intuitionists have promised us. The key addition from new intuitionists was to stress the epistemological overdetermination of self-evidently justified moral propositions. An account of how we can *reject* certain propositions as self-evident is beside that point, and (more importantly) it does not help answer the charge of dogmatism. So why have new intuitionists stressed epistemological overdetermination?

Showing that there are good reasons to reject that a proposition is self-evident is, I think, essential to new intuitionism. If new intuitionists can show that one can do more than dogmatically assert that a proposition is *not* self-evident, than perhaps one can do

more than dogmatically assert that a proposition *is* self-evident. To see this move, let's turn to Stratton-Lake's reconstruction of the dialectic between Ross's *The Right and the Good* and *Foundations of Ethics*.

Stratton-Lake argues that even though Ross initially accepted **Pleasure** as self-evident, he was nevertheless able to provide supplementary justification for that proposition. Importantly, this supplementary justification was inferential, derivative, and non-direct. Which means, if there is this supplementary justification, then the justification for **Pleasure** is overdetermined: that proposition admits of both self-evident justification, and alternative justification.

Ross accepts **Pleasure** in *The Right and the Good* for four reasons. First, he of course claims that **Pleasure** is self-evident — if one understands the constituent parts, has reached sufficient mental maturity, and given sufficient mental attention to that propositions, one can come to see that it requires no evidence beyond itself. But there is more one can say, which is what the next three reasons demonstrate.

Second, Ross says that if we imagine "two states of the universe", one state with widespread pleasure and the other state with widespread pain, then we cannot help but conclude that the state with widespread pleasure is the better of the two (Ross 1930: 135). Ross says that recognizing this gives us reason to accept **Pleasure** — to be clear, it does not give us a reason to think that **Pleasure** is *self-evident*, but rather a reason to think that **Pleasure** is *true*. Third, if we think that 'kindness' is the giving of pleasure, and that 'cruelty' is the giving of pain, then our widespread approval of kindness and our widespread disapproval of cruelty is best explained by assuming that **Pleasure** is true (Ross 1930: 134). And fourth, we can see that happiness (which is how Ross also understands pleasure) is intrinsically good. If we think that virtue ought to be rewarded with happiness (i.e. pleasure) and that vice ought to be rewarded with unhappiness (i.e. pain, or else the absence of pleasure) then Ross says that this shows that we think "deserved happiness" (i.e pleasure)

is good in itself (Ross 1930: 136). On the basis of these four arguments, Ross concludes that we should accept **Pleasure**.

I do not offer any assessment of these arguments here. They are presented only to demonstrate that Ross clearly *supplements* the claim that **Pleasure** is self-evidently justified with other arguments. The justification for the proposition here is overdetermined. So it is appropriate to read Ross in the following way: a moral proposition that can be justified self-evidently need not *only* be justified self-evidently. This is the key new intuitionist claim, and Ross's arguments in favour of **Pleasure** in *The Right and the Good* seem to show that.

But if **Pleasure** is self-evidently justified, and one can also provide at least three additional arguments for that claim, then why did Ross later come to reject it? Moreover, why did he both reject that the claim is self-evident *and* reject that the claim is true? In *The Right and the Good* (1930) Ross had already recognized a potential problem. If **Pleasure** is true, then it follows that Ross himself should think there is a *prima facie* duty to promote pleasure.[17] Toward that end, Stratton-Lake says that Ross gives a "half-hearted" defense of a duty to promote pleasure (Stratton-Lake 2002: 128). But the seeds of doubt were already planted. Upon further reflection, Ross came to appreciate "a number of other difficulties with the thesis that pleasure is intrinsically good" (Stratton-Lake 2002: 129).

As before, we can find three inferential arguments in *Foundations of Ethics* (1939), this time aimed at showing that **Pleasure** is not self-evident and that it is false. First, Ross says that "there is nothing admirable or commendable in the mere feeling of pleasure" (Ross 1939: 271). This is incompatible with accepting that **Pleasure** is true. For Ross, we should be able to replace 'good' with 'admirable' or 'commendable'. But Ross does not think that 'pleasure is commendable' is equivalent to 'pleasure is good'. Second, intrinsic goods should

[17] Ross says that "some [duties] rest on the mere fact that there are other beings in the world whose condition we can make better in respect of virtue, or of intelligence, or of *pleasure*. These are the duties of beneficence" (Ross 1930: 21, emphasis added).

reflect positively on a person that has them. But Ross says that someone is not evaluated positively "in respect of the mere fact of feeling pleasure" (Ross 1939: 271). And third, Ross notes that we do not have a *prima facie* duty to promote *vicious* pleasures, and we would if **Pleasure** was true (Ross 1939: 274). On the basis of these three arguments, Ross concludes that not only is **Pleasure** not self-evident, it is also false.

Again, I offer no assessment of these arguments. They are noted here to show that one can do more than dogmatically assert that **Pleasure** is false. Moreover, Ross's change in view between *The Right and the Good* and *Foundations of Ethics* is important for understanding new intuitionism. We can see above how even self-evidently justified propositions can be epistemologically overdetermined. Old intuitionists can be accused of being dogmatic about which moral propositions are self-evident. But if we appreciate how even those self-evidently justified propositions can admit of alternative justification, then intuitionists seem capable of addressing that criticism. By stressing that self-evident propositions (which are **Non-inferential**, **Non-derivative**, and **Direct**) can also be justified in ways that are inferential, derivative, and non-direct, new intuitionists can do much more than dogmatically assert that certain claims are or are not self-evident. When the appeal to self-evidence proves unpersuasive, an intuitionist can lead their interlocutor through a more typical reasoning process.

As I see it, this is a welcome insight and development from the new intuitionists. These types of epistemic supplements are important because what may seem self-evident to us may not seem so to others. We saw how Ross seemed to implicitly appreciate as much, which is why in addition to relying on self-evidence he provided further arguments in favour of those same propositions. New intuitionists have done important work to fill in the details, seeming to recognize the problems of disagreement and dogmatism. To sum up the guiding thought behind epistemological overdetermination, even if self-evident moral propositions do not *require* rational justification or evidence beyond themselves, "it is nevertheless good if [those things] are available . . . [because] anything that gives us further, i.e. extrinsic, reason

to think some apparently self-evident moral proposition is true is useful" (Stratton-Lake 2002: 118).

4.5 A Dilemma

Let us now return to what was supposed to be appealing about moral intuitionism in the first place. Recall that nonnaturalist cognitivist realism faces an important explanatory challenge: why think that our moral beliefs link up to the moral facts? That is, why think that our moral judgments are sometimes *successful* given the two desiderata of *cognitivism* and *nonnaturalism*? What bridges the gap between our moral beliefs and the moral facts?

Moral intuitionists meet this challenge by saying that we have the ability to *intuit* some moral propositions: we can know about them in a **Non-inferential**, **Non-derivative**, and **Direct** way. That is, we can sometimes intuit that certain moral propositions are self-evidently justified. We come to know that they are justified merely by understanding their constituent parts. And because of this, we do not need to be in causal contact with the moral facts, nonnaturalistically construed. Moral intuitionism is therefore appealing because it provides an epistemological account amenable to nonnaturalist cognitivist realism.

As we saw, one consequence or implication of thinking that some moral propositions are self-evidently justified is that moral intuitionism is best characterized as a foundationalist theory of moral knowledge — it holds that belief in some moral propositions can be self-evidently justified, and that even if there are non-inferentially justified beliefs in moral propositions this justification has a route back to some non-inferentially justified belief. These foundationally justified beliefs are self-evidently justified. But it was precisely this appeal to self-evidence that exposed old intuitionists to charges of dogmatism. New intuitionists addressed these charges by appealing to the notion of epistemological overdetermination. Since even self-evidently justified moral propositions can be justified in other ways, the view is not dogmatic. A stubborn interlocutor can instead be lead through an inferential and derivative reasoning process and come to accept a given moral proposition.

Though this is a welcome development, I think this raises a new and important (and so-far unseen) problem for new intuitionists. How can new intuitionists reconcile and square the foundationalism of old intuitionism with their novel emphasis on epistemological overdetermination? As we saw, the new intuitionists appear to deftly avoid the charge of dogmatism. However, I think that the way they do so comes at a significant cost. By embracing epistemological overdetermination, new intuitionists inadvertently undermine intuitionism's foundationalism. New intuitionists rightly claim that a proposition that is self-evidently justified is not necessarily precluded from being inferentially justified. But it seems to me that epistemological overdetermination actually *undermines* intuitionism's foundationalism: if a proposition can be inferentially justified, then we have no recourse for thinking that it can *also* be self-evidently justified. To be sure, it doesn't show that a proposition *cannot* be self-evidently justified. Rather, self-evident justification seems to drop out of the picture entirely. And if it has, then so too has the foundationalism that was essential to intuitionism.

Now, why is it a problem if intuitionism ends up being non-foundationalist? Because that was what allowed intuitionism to address the explanatory challenge mentioned at the outset — namely, why we should think that our moral judgments link up to the nonnatural moral facts. The appeal to self-evidence provided the important link here. Rather than relying on some causal account, intuitionists say that we can know some moral propositions in a **Non-inferential**, **Non-derivative**, and **Direct** way. That is what allows us to bridge the gap between our moral beliefs and the causally inefficacious moral facts. But if inferential justification is now doing the epistemic heavy-lifting — and more importantly, the explanatory work about which moral propositions are true — then intuitionism seems better characterized as a non-foundationalist moral epistemology.

Consider an exchange between someone who accepts **Future** and someone who does not. Suppose Henry is attempting to convince you that **Future** is a moral proposition that you should accept. He tells you that **Future** is self-evidently justified. You've reached sufficient mental maturity, thought about the issue deeply, but you're ultimately not convinced that

Future is true (and so not convinced that you should accept it). Henry stresses that you should accept **Future** because it is justified in virtue of its constituent components. But you still are not convinced, and you accuse him of dogmatism. Wanting to both avoid this charge and convince you that **Future** is true, Henry changes strategies. He drops the appeal to self-evidence. Instead, he walks you through a reasoning process. He gives you other reasons and arguments for thinking that you shouldn't prefer a present lesser good to a future greater one. Maybe he shows how you should accept this because it follows from things you already believe. Or perhaps it is the result of a clever argument. In any event, Henry's strategy works! Lo and behold you come to accept **Future**.

I think we can agree that Henry has made important strides, and for that we can thank the new intuitionists. Where before the conversation may have ended with the charge of dogmatism, now Henry can remind himself that self-evident propositions are epistemologically overdetermined. However, the moral epistemology that Henry relies on to convince you is not intuitionist, nor is it foundationalist — the justification here is instead inferential, derivative, and non-direct. It's also important to note that *you* are not intuitionist or a foundationalist, since you still do not think that **Future** is self-evidently justified. So it looks like the introduction of epistemological overdetermination has caused foundationalism to drop out of the picture. Rather than it being (as Stratton-Lake says) a good thing that an epistemic supplement is available, it instead seems *necessary* for getting an interlocutor to accept **Future**.

This reveals the problem for intuitionism, which takes the form of a dilemma. Without stressing the possibility of epistemological overdetermination, intuitionists are open to (and unable to address) charges of dogmatism. But when intuitionists do stress epistemological overdetermination, intuitionism loses its grip on the necessity of self-evident justification, loses its status as a foundationalist moral epistemology, and is unable to explain how we come to know the nonnatural moral facts. So either intuitionism is foundationalist and dogmatic, or else it is non-dogmatic but non-foundationalist. And since foundationalism is

necessary for answering the explanatory challenge and bridging the gap between our moral beliefs and the moral facts, giving up on foundationalism concedes too much. Intuititionism may be able to answer charges of dogmatism, but at too great a cost. For that reason, the view in untenable.

4.6 Conclusion

I ended Chapter 3 by claiming that nonnaturalist cognitivist realism does not infer theoretical moral judgments from folk moral judgments. Rather, whether the view is susceptible to **Inference Debunking** or not so depends on how it ultimately characterizes our successful moral judgments. Intuitionism is one such way of explaining how it is that we can know the nonnatural moral facts. While this account does not simply infer theoretical judgments from folk judgments, there is another problem. Either the view is dogmatic, or else it cannot actually explain how we know the moral facts. I conclude that without resolving this dilemma new intuitionism is not in a position to support nonnaturalist cognitivist realism.

CHAPTER 5
The Circularity of Reflective Equilibrium

In this chapter I move on to reflective equilibrium. I take this to be a natural transition because reflective equilibrium has traditionally been presented and defended in opposition to moral intuitionism.[1] In Chapter 4 we saw that the British Ethical Theorists understood moral intuitionism as a *foundationalist* epistemology. And we also saw that the new intuitionists added to the view, but at the expense of undermining intuitionism's claim to foundationalism. In light of that development, one response we might have is to celebrate the demise of intuitionism precisely *because* it is foundationalist. We might think that we should look for, and prefer, a non-foundationalist moral epistemology.

It is here that reflective equilibrium stands out. Most philosophers understand reflective equilibrium as a paradigmatically non-foundationalist moral epistemology.[2] It offers instead a *coherentist* moral epistemology. So those suspicious of intuitionism and its foundationalism might welcome reflective equilibrium and its coherentism.

---

[1] "In the moral case, [reflective equilibrium] entails rejecting rational intuitionism, or the view that moral theories may be founded on a privileged set moral intuitions" (Little 1984: 376). "One of the main appeals of reflective equilibrium as a moral methodology lies in its ambition to provide an account of moral objectivity, without having to bear the controversial epistemological and metaphysical burdens of alternative, 'foundationalist' methods of moral justification, such as rational intuitionism, which claim that certain moral beliefs can somehow be known non-inferentially" (de Maagt 2017: 444).

[2] See Cath (2016) and Tersman (2018) for recent overviews. Cath writes that "the method of [reflective equilibrium] is standardly interpreted as relying on a coherentist theory of justification and, indeed, is often referred to as simply being the 'coherence method"' (Cath 2016: 216). However, a minority of philosophers disagree, suggesting instead that the view is foundationalist. For the most notable of these latter arguments, see Bonjour (1985). In this chapter I follow the standard view of reflective equilibrium.

But here is a worry you might have. Recall that in the previous chapter I argued that intuitionism's foundationalism (and in particular the appeal to self-evidence) was its *strength*. Because the view is foundationalist and relies on the notion of self-evidence, intuitionism seems to be able to bridge the gap between our moral beliefs and the nonnatural moral facts (at least in principle). Indeed, abandoning foundationalism is what generates a serious problem for new intuitionism, since it can no longer meet that explanatory challenge. Now, in light of that, why think that a coherentist moral epistemology (which is non-foundationalist) is going to be helpful for nonnaturalist cognitivist realism? How will a coherentist view close the chasm between our moral judgments and the nonnatural moral facts? The worry is that the present chapter is in tension with the previous one.

I think any tension here is illusory. To see why, let's remind ourselves of the challenge from the previous chapter. If we have moral knowledge then it is because we have successful judgments about nonnatural moral facts. But how? If moral facts are nonnatural how can we know anything about them? Intuitionists thought they could bridge this gap by appealing to self-evident justification, which appears to close the chasm between our moral judgments and the nonnatural facts. Intuitionists say that some of our moral beliefs are non-inferentially justified because some moral propositions are self-evident. As such, no causal contact with the moral facts is needed. But I argued that for new intuitionists self-evidence becomes an idle wheel, and so the view is still unable to explain how we come to know about nonnatural moral facts.

As we can see, the claim was not that a view *must be* foundationalist in order to explain our knowledge of nonnatural moral facts. Rather, it was that one particular foundationalist view thought it could do so. And this view, I argued, faced a dilemma: either the view is dogmatic or else it is non-foundationalist. Because *by it's own lights* intuitionism is supposed to be foundationalist, abandoning that position is what creates the problem. But this of course does not rule out some other non-foundationalist view being plausible. For that reason, I think the previous and present chapters are not in tension. We have no antecedent

reason to think that a non-foundationalist view cannot bridge the gap between our moral judgments and the moral facts. As always, we will have to look at how the view says we can come to know those facts. Thus, looking at a non-foundationalist view like reflective equilibrium is a legitimate next move.

As I unpack the details of reflective equilibrium below, it will become clear that (like moral intuitionism) the view does not appear to infer theoretical moral judgments from folk moral judgments. So reflective equilibrium is not subject to the EDA and **Inference Debunking**. However (again like moral intuitionism), showing the details of the view will allow us to see other notable problems with reflective equilibrium. And these problems, I will argue, make the view untenable.

Though intuitionism and reflective equilibrium are longstanding antagonists, there is also a more local reason to turn to reflective equilibrium. As I have said, nonnaturalist cognitivist realists have not offered a plausible moral epistemology for their metaethical view. While one may take this as a notable weakness, some nonnaturalist cognitivist realists see the situation in a different light. David Enoch, for one, suggests that we should look for the "the philosophical theory that is best as a theory *overall*" (Enoch 2011: 14, emphasis added). He says that we should accept the theory that winds up with the most so-called 'plausibility points', rather than the theory which is best in every particular area. By exploring this claim about plausibility points we will see how it becomes important to examine reflective equilibrium.

For the sake of convenience, suppose that there are just two components of a metaethical theory: its metaphysics and its epistemology. If we take a plausibility points approach, we might think that a metaethical view can score and lose points in these two areas. For example, a metaethical view might score poorly on epistemic matters (by offering an implausible account, or not offering an account at all), yet score highly when it comes to metaphysics (by offering a parsimonious account). If we take an overall view on these issues, then we might prefer that metaethical theory to one that has middling scores in both areas. In this

respect, the lack of a plausible moral epistemology is not a decisive blow against nonnaturalist cognitivist realism. A defender of the view might say that, given everything else that the view seems to get right we should not abandon it simply because of its weaknesses on the epistemological front. That is Enoch's strategy. He argues that, even if nonnaturalist cognitivist realism does not have a worked out epistemic account, this alone does not sink the view. This is because he thinks the view gains plausibility in other areas so that it can withstand an attack on the epistemic front. But more importantly, Enoch's strategy is not simply to let nonnaturalist cognitivist realism face the full force of epistemic objections, but rather to defend against those worries. How he does this takes us to reflective equilibrium.

5.1 Enoch's Conjecture

Enoch argues that even if the nonnaturalist cognitivist realist cannot offer a worked out moral epistemology, they can offer something in that vicinity: they can offer a sketch of why we should think that it is not *unreasonable* to believe that we have moral knowledge. He presents an argument for thinking that we *might* have moral beliefs that are not so off-track such that we cannot hope "our reasoning processes will in general get us closer to rather than further from the normative truths" (Enoch 2010: 428). He claims that:

> given a starting point of normative beliefs that are not too far-off, presumably some reasoning mechanisms (and perhaps some other mechanisms as well) can get us increasingly closer to the truth by eliminating inconsistencies, increasing overall coherence, eliminating arbitrary distinctions, drawing analogies, ruling out initially justified beliefs whose justificatory status has been defeated later on, etc. (Enoch 2010: 428).

Note the two important aspects here: first, the starting point normative beliefs; and second, the reasoning mechanisms. As we will see, these two aspects capture the process of reflective equilibrium in everything but name. Though Enoch does not spell out the details of this process, he nevertheless appears to think that *something* like reflective equilibrium can

plausibly serve as a bulwark against sceptical arguments. So let's call this **Enoch's Conjecture**: reflective equilibrium (or something like it) can get our moral judgments closer to the moral facts. If this claim is correct then we have a good candidate for a plausible moral epistemology, namely reflective equilibrium.

In this chapter I use **Enoch's Conjecture** as a way of assessing the plausibility of reflective equilibrium. In the rest of 5.1 I explain why Enoch thinks that something like reflective equilibrium will be helpful to the nonnaturalist cognitivist realist. In section 5.2 I carefully describe the method of reflective equilibrium, first drawing upon descriptions of the method from Nelson Goodman and John Rawls, and then moving on to a recent defense from T.M. Scanlon. But in section 5.3 I will explain why we should not accept **Enoch's Conjecture**. Once we look at the details of reflective equilibrium it becomes clear that the view cannot actually help the nonnnaturalist cognitivist realist. In particular, there is a lacuna regarding the so-called *support relation* between moral judgments and the principles that account for them. Attempts to fill in this lacuna result in reflective equilibrium being circular, or so I will argue. For this reason, reflective equilibrium cannot explain why we have successful moral judgments. I conclude that it is not a plausible moral epistemology, and so it is of no help to the nonnaturalist cognitivist realist. I wrap up in section 5.4

5.1.1 Survival is (*pro tanto*) Good

What are our starting point normative beliefs? Enoch's argument for why we might have normative beliefs that are not too off-track relies on what is now called a *third-factor argument*.[3] Such arguments attempt to explain the correlation between two points, say A

[3] For prominent third-factor arguments see Enoch (2010; 2011) and Wielenberg (2010; 2014). See also Brosnan (2011); Schafer (2010); and Skarsaune (2011). For a recent analysis of third-factor arguments, see Tersman (2017). Third-factor arguments also share much in common with so-called *common cause explanations* familiar from philosophy of science. See Sober (1984).

and B, by identifying some third point, C. Rather than explaining the correlation between A and B *in terms of each other*, we can explain their connection in terms of another factor, C.

We can see this by looking at a non-moral example. Enoch asks us to consider the correlation between, on the one hand "giving rise to strongly affectionate feelings", and on the other hand, "having a poor sense of time" (Enoch 2010: 431). Suppose that whenever we see that someone gives rise to strongly affectionate feelings they also almost always have a poor sense of time. If so, what could explain this correlation?

One thing we could say is that giving rise to strongly affectionate feelings *causes* one to have a poor sense of time. But how this would be the case is mysterious. How exactly does the former cause the latter? Without a more detailed explanation we should be suspicious that we have adequately explained the correlation here. Similarly, we might say that having a poor sense of time *just is* the sort of thing that gives rise to strongly affectionate feelings. But this too seems unhelpful. Why should we think that the two are identical when they seem clearly distinct? Another thing we could say is that there is *no correlation* between the two. After further reflection, perhaps we should give up thinking that the two are correlated at all. That is, we could be sceptics about the correlation.

However, there is another way to explain the correlation. The two factors here — 'giving rise to strongly affection feelings' and 'having a poor sense of time' — might be captured by a third factor. This third factor would be something *besides* the two points that nevertheless explains the correlation. For our two factors here, the third factor that might explain them is this: 'the fact that young children are cute'. This fact could help explain the correlation between 'giving rise to strongly affectionate feelings' and 'having a poor sense of time' in the following way. On the one hand, being a young child *explains* having a poor sense of time (children being notoriously poor time keepers). And on the other hand, being cute *explains* giving rise to strongly affectionate feelings (being cute *just is* the sort of thing that gives rise to those feelings). So 'the fact that young children are cute' could be a third factor that

explains the correlation between 'giving rise to strongly affectionate feelings' and 'having a poor sense of time'. According to Enoch, we might say that 'the fact that young children are cute' "pre-establishes the harmony between [the two points]" (Enoch 2010: 431). Indeed, given that particular third factor, we should *expect* to see a correlation between 'giving rise to strongly affectionate feelings' and 'having a poor sense of time'.

It is not essential that the example *does* in fact work. For present purposes, what's important is that an explanation like this *could* work.[4] If we can sometimes appeal to third factors in order to explain apparent correlations then we have an argumentative strategy that might be of help to the nonnaturalist cognitivist realist. So if we understand how third factor accounts work in general then we can move on to a moral example.

We can ask: what explains the correlation between our moral judgments, on the one hand, and the moral facts, on the other? This was the challenge we explored in the previous chapter. Here again we might try answering that the points are explained *in terms of each other*. We might say that we have the moral judgments that we do precisely because the moral facts are thus-and-so. But it is again mysterious why this should be the case. Do the moral facts somehow reach out to us and compel us to have the sort of moral judgments that we do? Arguably not, since we should accept *nonnaturalism* and with it the causal inertness of abstracta. Nor can we say that the moral facts are thus-and-so *because* of our moral judgments. This seems to get the order of explanation reversed: our moral judgments are supposed to *link up* to the moral facts, not explain what the moral facts are.[5] Another

[4] Enoch says that "I have no idea whether this explanation actually works (or whether the phenomenon it is supposed to explain is actually a real phenomenon). All that is crucial for me is that it *could* work, and that its *structure* is exactly similar to that of the explanation I am really after" (Enoch 2010: 431, emphasis his).

[5] However, there are metaethical accounts that argue for something like this. So-called *constructivist* accounts claim that our moral judgments are correct when they follow some to-be-specified procedure, usually either a rational or deliberative procedure — i.e. Kantian constructivism (see O'Neill (1989) and Kosrgaard (2003), or else a procedure relativized to

option would be to say that any apparent correlation is illusory. But, of course, that sceptical conclusion does not help the nonnaturalist cognitivist realist.

However, nonnaturalist cognitivist realist can avail themselves of a third-factor argument. Just as we can explain the correlation between the non-moral factors by appealing to another fact (namely, 'that children are cute'), we can explain the correlation between the moral factors by appealing to the following fact: 'that survival is (*pro tanto*) good'. According to Enoch, the fact 'that survival is (*pro tanto*) good' explains why we have the moral judgments that we do. Many of our moral judgments are about survival — e.g. 'that someone has treated me poorly is a reason to distrust them in the future', or 'that I have children is a reason to take care of them', etc. And that fact also explains why our judgments are not too far off from the moral facts: because the aim (in a loose sense) of evolution, survival, is (*pro tanto*) good. So like before, we can say that the fact 'that survival is (*pro tanto*) good' pre-establishes the harmony between our moral judgments and the moral facts. Given this, Enoch says that we should expect our moral judgments to (at least somewhat) resemble the moral facts.

5.1.2 The Importance of Reasoning Mechanisms

My primary concern is not with Enoch's third-factor argument. As I said, this argument needs to be presented in order to understand the claim that our starting point normative beliefs might not be too far off-track. So I will take up the question of the *reasoning mechanisms* that our purportedly 'not too far-off' normative beliefs are subject to. After all, per **Enoch's Conjecture**, these mechanisms justify thinking that our normative can resemble the nonnatural moral facts. So I will examine why we should think that the mechanisms that

---

one's own practical standpoint — i.e. Humean constructivism (see Street 2008). On these sorts of constructivist accounts, the moral facts are not 'out there' to be discovered, but rather are in an important sense *constituted by* our moral judgments themselves. But since these constructivist accounts deny *nonnaturalism*, I do not focus on them.

Enoch mentions — i.e. eliminating inconsistencies, increasing overall coherence, eliminating arbitrary distinctions, drawing analogies, ruling out initially justified beliefs who justificatory status has been defeated later on, and so on — will be of any help to the nonnaturalist cognitivist realist in bridging the gap between our moral judgments and the nonatural moral facts.

This is important because Enoch's third-factor argument does not assume a *strong* correlation between our moral judgments and the moral facts, rather only a *weak* one. It is not the case that survival is *all things considered* good, or *always* good, just that it is at least *sometimes* good. So the third factor here gives us reason to think that our moral judgments are not too far-off from the moral facts. There is some correlation between our moral judgments and the moral facts, and if we start from this position then reflective equilibrium can perhaps help tighten up the correlation. In turn, this is why it is important to take a close look at reflective equilibrium. If Enoch's third-factor argument is right, then our moral judgments might not be too far off track — in which case, the process of reflective equilibrium might be able to bring our moral judgments closer to resembling the the moral facts. And the process of reflective equilibrium supposedly does this by eliminating inconsistencies, increasing overall coherence, and so on.

Of course, *how* it can achieve this depends on the details of reflective equilibrium. And it's these details that occupy the next two sections.

5.2 The Standard View

Contemporary discussion of reflective equilibrium is indebted to both Nelson Goodman and John Rawls.[6] Working in different traditions (Goodman in logic, Rawls in moral

[6] Rawls rejects intuitionism, but he is clear that he is not targeting the claims of the British ethical theorists but rather the 'intuitionism' of Barry, Brandt, and Rescher. This latter intuitionism is committed to just the following claims: the notion that there is a plurality of first principles which can conflict when applied to cases, and that we can strike

and political philosophy), and towards different ends (justification of logical inference versus justification of moral and political principles), they are largely responsible for the current understanding of reflective equilibrium's unique form of coherentist justification. It is noteworthy that reflective equilibrium is apparently perfectly general, insofar as it able to capture justification in both logic and moral philosophy. I take it that this generality is a virtue of the method. Turning then to this general account, or what I call the *standard view* of reflective equilibrium, will allow us to see the epistemic strengths of which nonnaturalist cognitivist realists hope to avail themselves.

5.2.1 Goodman

It was Goodman (1955) who presented the first clear articulation of the standard view of reflective equilibrium. Offering an account of the justification of logical inference, there are two key components to Goodman's view:

(1) rules of inference (e.g. double-negation elimination)

(2) particular inferences (e.g. ¬¬p, so p)

Because his focus is on justification in particular, these two components give rise to two questions. First, when are rules of inference justified? And second, when are particular inferences justified?

While these two different questions would seem to invite two different explanations, Goodman's account highlights an important relationship between (1) and (2). He says that (1) and (2) are justified when they are "brought into agreement with each other" (Goodman 1955: 64). More specifically, (1) and (2) are justified not by something else (say, a third factor), but rather in *terms of each other*. This is very schematic, but note that this tells us

---

a balance here by 'intuition', that is, "by what seems to us most nearly right" (Rawls 1971: 34).

how both (1) *and* (2) can be justified. Just as reflective equilibrium offers a single justificatory account for both logic and moral philosophy, the view also provides a single account for the justification of rules of inference *and* for particular inferences.

With the two components on the table, what we need to examine next is exactly what it means for (1) and (2) to be justified in terms of each other. This requires us to get clearer about the locution 'brought into agreement'. Goodman's story is rather simple, and it appeals to the way in which rules of inference and particular inferences are used within our logical practices. He says that a rule of inference and a particular inference are 'brought into agreement' when each is compatible with the other *and* when we are unwilling to amend either of them. For example, if the rule of double-negation elimination is compatible with the particular inference '¬¬p, so p', and on rational reflection we are unwilling to give up neither the rule nor the inference, then the two are in agreement. And when there is such an agreement then both are justified.

As we can see, this agreement is what makes reflective equilibrium a *coherentist* form of justification. It isn't the case that (1) and (2) can be independently justified, for example by some special foundation or because each is self-evident. Rather, Goodman argues that (1) and (2) are justified because they *cohere* with each other in the relevant way. And it's important to note that the *coherency* of (1) and (2) is much different than their *consistency*. If all that was required for a rule of inference and a particular inference to be justified was that each was *consistent* with the other, then a rule of inference which was unrelated to a particular inference could be justified too easily. For example, a rule of inference like double-negation elimination is consistent with a particular inference like 'if p then p'. But these are unrelated to each other. They are consistent with each other, but they do not cohere.

For Goodman, it is critical that rules of inference and particular inferences agree with each other. We can get more grip on this form of justification by flipping the situation. What would it mean to say that a rule of inference and a particular inference are *not in agreement*? Consider the following situations. When a rule of inference results in an inference that "we

are unwilling to accept", then the two are not in agreement (Goodman 1955: 64). Or else when a particular inference "violates a rule we are unwilling to amend", the two are also not in agreement (Ibid.). We should not accept a rule of inference that was incompatible with what we genuinely think are acceptable particular inferences. And we should not think a particular inference is acceptable if it is incompatible with a rule of inference that we believe is happily in agreement with other inferences. In the absence of either, we can say that a rule of inference and a particular inference are in agreement with each other.

Because rules of inference and particular inferences are said to be justified in terms of each other, reflective equilibrium requires an openness to updating our assessment of both (1) and (2). It might be the case that we have a rule of inference that we are confident generates sound particular inferences. But suppose we come to learn that it results in an inference that strikes us as suspicious in some way. One thing we might say is that a rule that generates that inference must be mistaken; the rule is not as sound as we imagined. Another thing we might say is that the generated inference must be mistaken: the inference looks plausible but on reflection is not one we should accept. What we will need to do in this situation is reflect more on these two components. Are we more confident in the rule or the inference? And how do these two components cohere with the rest of the rules and inferences we currently accept? What we are doing in this situation is attempting to bring rules and inferences into equilibrium.

This is the characteristic 'rocking back and forth' between rules and inferences that is unique to the standard view of reflective equilibrium. The process requires us to revisit and reevaluate our acceptance of both (1) and (2). Unless we think that rules and inferences are in agreement, we cannot say that they are justified. However, Goodman notes that justification of (1) and (2) *in terms of each other* can look "flagrantly circular" (Goodman 1955: 64). On the heels of intuitionism's appeal to self-evident justification, reflective equilibrium appears to trade solid foundationalism and linearity for nebulous coherentism and circularity.

But Goodman suggests that the circularity of reflective equilibrium is virtuous rather than vicious. He suggests this because he understands the justificatory process as the delicate process of "making mutual adjustments between rules and accepted inferences" (Ibid.). Whether this successfully neutralizes the charge of circularity remains to be seen. But so it goes with coherentist accounts of justification. Absent special foundations there is in principle nothing to tether justification to a fixed point. Objecting to the lack of one is to assume that any plausible epistemology must be foundationalist.

Since reflective equilibrium is a general justificatory process, we are not wedded to thinking that the method only works for logical rules of inference and particular logical inferences. Instead, we can identify elements of the method that are able to capture reflective equilibrium as it might work in any area of discourse or practice. There are, I think, three such elements:

(i) our judgments about particular cases

(ii) our judgments about what accounts for those cases

(iii) our 'rocking back and forth' between (i) and (ii) in order to find agreement between them

As we saw, the particular cases that Goodman focuses on are those concerning logical inferences, that which accounts for those inferences is said to be logical rules, and the agreement sought and found is what confers justification on both (1) and (2). If we understand how the method works in general than we are well-positioned to see its relevance in moral philosophy.

Given the above setup the transition can be fairly quick. In a sentence, if we replace (1) 'particular inferences' with 'considered judgments', and (2) 'rules of inference' with 'principles', then we have a simple summary of Rawls's account of reflective equilibrium. The fundamental structure of the method stays fixed between Goodman and Rawls, since justification arises when our considered judgments and the principles that account for them

are in agreement. Besides applying the method to moral theorizing, what Rawls adds to Goodman's picture is richer detail concerning (i) and (iii). Let's look at those details now.

5.2.2 Rawls

When discussing reflective equilibrium in moral philosophy, our judgments about particular cases are typically referred to as 'considered judgments'. These are an important subset of the judgments that we can or do make about a particular domain. And when it comes to morality, of all the moral judgments that we could make only a select portion of them will be considered judgments. It is these considered moral judgments for which we seek 'agreement' with a moral principle. So it is important that we are able to clearly mark the distinction between moral judgments *as a whole* and considered moral judgments. But how to do this?

Rawls says that our considered judgments are "those judgments in which our moral capacities are most likely to be displayed without distortion" (Rawls 1971: 47). He identifies a number of distortions that can push us away from the ideal circumstances for making moral judgments. He says that we can exclude from the class of considered judgments "those judgments made with hesitation, or in which we have little confidence ... [or] those given when we are upset or frightened, or when we stand to gain one way or the other", and so on (Ibid.). If, for example, we are hesitant when we judge that 'torturing people merely for fun is wrong', Rawls thinks that this judgment should not enter into the privileged class of *considered* judgments. And because it is not a considered judgment, it is not the sort of judgment for which we need to find agreement with a moral principle. Ditto if we are not confident when judge that 'gender discrimination is bad'.[7]

[7] We will have to set aside concerns about whether standing to gain from a judgment being true should really be thought of as a distortion. We might think that if *any* of our judgments are true then we stand to gain something (e.g. epistemic or moral credibility, or some other virtue). Also, it seems that those who stand to gain from some moral judgments

So far, all we have is the claim that under certain ideal conditions we are more likely to issue considered moral judgments. If we wanted to know what marks the distinction between moral judgments and *considered* moral judgments, the answer is simply that the latter are made under certain ideal conditions. The notion of 'ideal conditions' plays a critical role in the method of reflective equilibrium. Note that Rawls is not saying that we *only* make moral judgments under these ideal conditions. Rather, he is suggesting that under those ideal conditions the judgments we make are *considered* judgments. It is *these* particular judgments for which we will attempt to find principles. And when these judgments and principles agree then they are in reflective equilibrium. This makes it important that we understand how and when we are in the conditions necessary for producing considered judgments in the first place.

On that point, Rawls says that those ideal conditions are

> conditions favorable to the exercise of the sense of justice, and therefore [we are] in instances where the more common excuses and explanations for making a mistake do not obtain. The person making the judgment is presumed, then, to have the ability, the opportunity, and the desire to reach a correct decision (or at least, not the desire not to) ... [T]he relevant judgments are those given under conditions favorable for deliberation and judgment in general (Rawls 1971: 47-48).

Though Rawls is talking about conditions favourable to exercising one's 'sense of justice', we can replace the political domain with pretty much any other (again, this is because reflective equilibrium is supposed to be a perfectly general method). The ideal conditions for making

---

being considered true are precisely those in a position to recognize *why* they are true. E.g. someone subject to gender discrimination will both be in a position to recognize it's badness, but also stand to gain from that judgment being taken seriously. And it's not clear why the latter is sufficient to rule out the judgment being classified as a considered judgment.

considered moral judgments are those where we are in conditions favourable to exercising our 'sense of morality'. What exactly these conditions are supposed to be is something I will take up in the critique below. But for now, note that Rawls is suggesting that these conditions are perhaps no different than those needed for good deliberation and judgment *in general*. We need to think in good faith and be trying to ascertain which of our moral judgments are correct. Moreover, we need to ability and the opportunity to think about the subject matter.

I take it that Rawls's account adds to Goodman's in the following two ways. First, Rawls describes what is involved in our *making* considered judgments. These judgments have a particular conceptual profile, namely that they are free from certain distortions like hesitation, lack of confidence, etc. And second, Rawls notes that there are certain ideal conditions for moral thinking (conditions that parallel those required for good deliberation in general). Recall the three elements of reflective equilibrium:

(i) our judgments about particular cases
(ii) our judgments about what accounts for those cases
(iii) our 'rocking back and forth' between (i) and (ii) in order to find agreement between them

We can see that Rawls adds to (i). It is not merely *any* judgments about particular cases, but rather our *considered* judgments about particular cases. And understanding what marks that distinction allows us to introduce the notion of the *ideal conditions* for making moral judgments. The ideal conditions for making considered moral judgments are those where our sense of morality can be freely and appropriately exercised.

Rawls also adds to (iii), the characteristic 'rocking back and forth' of reflective equilibrium. Here Rawls's account is similar to Goodman's, but he helps by applying the thought squarely to first-order moral theorizing. He says that "moral philosophy is Socratic: we may want to change our present considered judgments once their regulative principles are

brought to light . . . [and] knowledge of these principles may suggest further reflections that lead us to revise our judgments" (Rawls 1971: 49). Here is what this looks like in practice. Suppose we have a set of considered judgments {j1, j2, . . . }. We should then attempt to find principles that capture those judgments {p1, p2, . . . }. Reflective equilibrium is so named because it identifies the ideal outcome of having our considered judgments and the principles that capture them in a state of harmonious balance. Maybe p1 captures j1 *but not* j2, or p2 captures *both* j1 and j2. Maybe p1 captures j1 and we would be willing to jettison j2. The combinatorial possibilities are endless, but it is harmony that we seek here.

As an example, one principle that could capture the judgment 'torturing people merely for fun is wrong' is the following: 'it is wrong to cause unnecessary suffering'. If causing unnecessary suffering is wrong, then torturing people merely for fun is wrong — indeed, the former explains why the latter is wrong. We can then continue to ask if there are principles that capture *that* principle. Maybe we think that 'acts are wrong in proportion as they tend to produce the reverse of happiness'. Or maybe some other principles is needed. Again, we cannot say antecedently, rather we will have to engage in the method of reflective equilibrium itself. And as we find more and more agreement between judgments and principles that account for them, then we move closer to reflective equilibrium. In turn, per **Enoch's Conjecture**, our moral judgments are moving closer to the moral facts.

5.2.3 Problems with the Standard View

Note that when we describe a general principle as accounting for a particular inference we do so via a judgment: it is a judgment *about* whether a moral principle captures a considered judgment. This opens up a potential complication. Recall that the process of reflective equilibrium is said to work by trying to find principles that capture all and only our considered judgments. But what about our judgments *about* judgments? That is, what about our judgments that general principles account for particular judgments — do these need to be considered judgments as well?

So far as I understand the method of reflective equilibrium, I think that these judgments *about* judgments must be considered judgments as well. That is, they are subject to the same form of justification as our moral judgments themselves. This is because (as I've stressed) reflective equilibrium is a perfectly general theory of justification. It aims to explain justification in logic, justification in moral philosophy, and even justification in philosophical methodology itself. So when are we justified in thinking that a general philosophical principle captures a particular philosophical inference? More to the point: what is the *support relation* between moral judgments and the principles that purportedly account for them? An explanation of this support relation is missing in the accounts of reflective equilibrium from both Goodman and Rawls.

In light of this, we can help the case for reflective equilibrium by saying that the process of reflective equilibrium is comprised of judgments *all the way down*. According to the view, there is no fundamental difference between our judgments about particular inferences (e.g. about which acts are right or wrong) and our judgments about whether our general principles *account for* or are *in agreement with* those initial judgments. Because judgments (of both kinds) play such a critical role in the method, it is crucial that we understand the conceptual profile of these judgments and when we can identify them. As we saw, Rawls says that considered judgments are those that we issue in confidence, without hesitation, when we don't stand to gain from them, and so on. This much is easy to understand. But we also saw that Rawls identifies certain ideal conditions for issuing those judgments. He says that the ideal conditions are those necessary for good deliberation and judgment in general. So if we can identify these conditions then we will be in a position to issue considered moral judgments.

Unfortunately, there is much less detail from Rawls about these ideal conditions then there is about the conceptual profile of considered judgments. As such, the standard view of reflective equilibrium is not complete. Without an explanation of both the support relation between moral judgments and the principles that account for them, on the one hand, and the

ideal conditions for making considered judgments, on the other, we are not yet in a position to say whether reflective equilibrium is a workable view. Likewise, we cannot yet explain when our moral judgments can be said correspond to the nonnatural moral facts.

It is worth reminding ourselves of the challenge here. Some nonnaturalist cognitivist realists think reflective equilibrium is a process by which our moral judgments can come to resemble the nonnatural moral facts. This was **Enoch's Conjecture**. And what we have so far is an account of when our moral judgments might be plausible candidates for *starting point* normative beliefs. The starting point normative beliefs are just what proponents of reflective equilibrium call 'considered judgments'. But considered judgments are not *themselves* justified. Rather, justification arises when our considered judgments and the general principles that account for them are in agreement. But identifying when we are justified in thinking *that* is the case requires both a judgment and another general principle. And on it goes.

As we can see, the Standard View of reflective equilibrium has not yet explained when it is that we are in a position to say that our moral judgments resemble the moral facts. To put it another way: why should think that when our considered judgments and general principles are in agreement we are somehow closer to the moral facts? The Standard View gives us no reason to think so. Indeed, it is difficult to see why we should think that our moral judgments are justified at all. Perhaps if we had more detail about the supposedly ideal conditions for making considered judgments we could answer these questions. Rawls, however, does not provide them for us. Fortunately, a recent defense of reflective equilibrium does. And it is to that account that I will now turn.

5.2.4 Supplementing the Standard View

In recent work Scanlon builds upon the standard view of reflective equilibrium. His own account shares many of the details that we saw in Goodman and Rawls. Here again we see

that there are three components: (i) considered judgments, (ii) general principles, and (iii) the characteristic rocking back and forth between them in order to find agreement.

According to Scanlon, the first step for the method of reflective equilibrium is to identify a set of considered judgments about a subject matter. He says that these are judgments that seem to us "clearly to be correct and seem so under conditions that are conducive to making good judgments of the relevant kind about this subject matter" (Scanlon 2014: 77). More specifically, they are judgments about which we reach a preliminary conclusion that they do "not have any implausible implications or presumptions" (Ibid., 84). For example, some types of considered judgments might be "judgments about the rightness or wrongness of particular actions, general principles, or judgments about the kind of considerations that are relevant to determining the rightness of action" (Ibid., 77). To use our familiar example, we might say that 'torturing people merely for fun is wrong' is one such judgment. This judgment seems clearly correct, it doesn't seem to lead to anything implausible, it's a judgment we can arrive at when thinking about moral matters, and so on.

Note that Scanlon suggests that these are judgments that seem to us correct when we are in those *conditions* which are conducive to making good judgments of that sort. These are the 'ideal conditions' that we heard about from Rawls. But like Rawls, Scanlon is also not explicit here. While there is no positive account given, there are nevertheless some negative proposals about what *does not* count as a good condition under to which make considered judgments. For example, sitting in your armchair and thinking that there is a rock on the moon with your name on it is not the sort of condition under which good judgments of that sort are made (Scanlon 2014: 83). This implicitly gives us a clue about the right sort of conditions for such judgments.

Here is one attempt at unpacking this thought. It seems that our *a priori* judgment about a rock on the moon is inadequate because what we should rely on instead is an *a posteriori* judgment. Let's say that we can specify a particular domain that captures 'names on rocks on the moon'. The conditions under which we can make good judgments about

this domain are clearly not 'sitting in one's armchair'. If we understand what is involved in making judgments about this domain then we should understand why armchair thinking alone is inadequate. We can infer from this that being under conditions where we simply rely on *a priori* judgments is not conducive to making good judgments about the 'names on rocks on the moon' domain. Instead, it seems that we should instead be under conditions which we rely on *a posteriori* judgments. If something like this is right, then we can easily make up for Scanlon's lack of detail regarding the appropriate conditions for issuing considered judgments. From his negative proposal about what *would not* work we can infer what would work.

This should resonate with what we heard in Chapter 2. Recall that we characterized the distinction between naturalism and nonnaturalism using the *subject matter* strategy: some subjects are about natural things, other subjects are about nonnatural things. And the type of reasoning we engage in regarding natural subject matters is empirical, while that for nonnatural subject matters is rationalist. Another way to say this is that natural subjects require *a posteriori* judgments while nonnatural subjects require *a priori* judgments. We can avail ourselves of these distinctions here to help the case for reflective equilibrium. The appropriate conditions under which to make good judgments about a given domain will either be those conducive to issuing *a priori* judgments or else those conducive to issuing *a posteriori* judgments. And (per Chapter 2) if we should accept nonnaturalism about moral properties, then our moral judgments would have to be *a priori* judgments. Therefore, the ideal conditions for making moral judgments are those conducive to making *a priori* judgments in general.

After we have identified our considered judgments, and the ideal conditions needed for making them, the next step is to "formulate general principles that would 'account for' these judgments" (Scanlon 2014: 77). What does 'accounting for' mean? Here Scanlon sticks close to Rawls and Goodman. He says that a general principle accounts for a considered judgment if the following is true: had you simply been trying to apply that principle, you would have

been led to that judgment. For example, if you were trying to apply the general principle 'all and only immoral actions cause someone to suffer', it would seem to result in the judgments 'this action causes so-and-so to suffer, so this action is immoral'.

And from there third step is that we should work "back and forth between principles and judgments" (Ibid., 77). So Scanlon avows all three features of reflective equilibrium. On his account, we begin with considered judgments (i), then attempt to find general principles that account for those judgments (ii), and we engage in the characteristic rocking back and forth between (i) and (ii) in order to find agreement between them (iii).

Just as Rawls adds needed detail to Goodman's picture, Scanlon adds needed detail to Rawls's picture. In particular, Scanlon's account allows us to infer what the ideal conditions for issuing considered moral judgments are. Because moral judgments are about nonnatural facts, the ideal conditions are those conducive to making *a priori* judgments.

Scanlon also provides more detail about the way justification arises in the process of reflective equilibrium. This is important because it can help the nonnaturalist cognitivist realist explain why we should think that judgments that survive to equilibrium somehow correspond to the nonnatural moral facts (i.e. **Enoch's Conjecture**). Scanlon says that the "justificatory force, if any ... must lie in the details of how the equilibrium is reached" (Scanlon 2014: 79). This does not depart from the Standard View, but it is a helpful clarification. Justification arises not because there are special foundations for considered judgments or general principles, but because of the way in which we arrive at the reflective equilibrium between them.

However, as I explain in the next section, spelling out the implications of Scanlon's claim reveals a serious problem for reflective equilibrium in general, and for the nonnaturalist who rely on it in particular.

5.3 The Circularity

The beliefs that survive the process of reflective equilibrium are said to be justified because of "the substantive merits of the judgments we make along the way" (Scanlon 2014: 82). What does this mean? It does *not* mean that a moral judgment is justified simply because it is confidently held, or because we are inclined to make that judgment without hesitation, or because we do not stand to gain in any way from that judgment being correct. These are necessary conditions for considered judgments, but they are not themselves sufficient for thinking that those considered judgments are *justified*. Rather, as Scanlon stress, what matters is that these beliefs seem "to be clearly true when I am thinking about the matter under good conditions for arriving at judgments of the kind in question" (Ibid.). This, I take it, is the critical point of Scanlon's account.

Above I raised a problem for the standard view: what justifies our judgments about whether general principles account for particular judgments? The only plausible answer here is that we will know once we have engaged in the process of reflective equilibrium. I take it that is why Scanlon says that our judgments are justified by the merits of the judgments we make within the process of reflective equilibrium. Since justification does not get into the picture from *outside* the method, it must come from within. So if our moral judgments are justified it is because they are in agreement with general principles. And the reason we are licensed to think *that* is because of our judgments about whether certain principles account for particular judgments. So these judgments and principles will also have to survive the process of reflective equilibrium.

But this reveals a serious problem: reflective equilibrium appears to justify our moral judgments only if we already accept that reflective equilibrium justifies our judgments. As we know, reflective equilibrium is a view about when moral judgments are justified. The view says that a moral judgment is justified if it is in agreement with a general moral principle. And as I explained above, whether a general moral principle is in agreement with a moral judgment is *itself* a judgment. Therefore, this judgment too must be justified — and whether

it *is* justified depends on whether there is principle with which *it* agrees. But here is the problem: this means that reflective equilibrium is already doing the justificatory work at this critical step. And that means that we can only conclude that a moral judgment is justified if we accept that reflective equilibrium justifies our judgments. This, I submit, is flagrantly circular.

The worry can be expressed in two ways. One way is to say that reflective equilibrium does not actually provide justification for our moral beliefs, rather it pushes the problem back *ad infinitum*. Because justification for a moral judgment relies on agreement with a principle, and whether that agreement is justified is a judgment itself, the problem of justification is never answered. Another way to express the worry is in terms of circularity. Reflective equilibrium ultimately says that it justifies itself. A judgment is justified when it agrees with a principle, and whether *that* is the case is decided by reflective equilibrium. So the much needed support relation between judgments and principles that account for them turns out to be decided by reflective equilibrium itself — hence the circularity.

At the beginning of the chapter we saw Enoch's argument for thinking that we might have moral beliefs that are not so off-track such that we cannot hope that "our reasoning processes will in general get us closer to rather than further from the normative truths" (Enoch 2010: 428). What we called **Enoch's Conjecture** says that if we have starting point moral beliefs that are not too off-track, then something like reflective equilibrium can get us closer to the moral truths. By eliminating inconsistencies, increasing overall coherence, eliminating arbitrary distinctions, and so on, we are justified in thinking that our moral beliefs are more likely to reflect the moral facts.

But we should not accept this claim. Reflective equilibrium will not help the nonnaturalist cognitivist realist, because the former view is is circular at a critical juncture. In order to conclude that a moral judgment is justified we must explain why we are justified in thinking that a particular judgment is in agreement with a general principle. That is, we must explain the support relation between our moral judgments and the principles that

account for them. And *that* is itself a judgment, which in turn will need to be in agreement with another general principle. So this means that reflective equilibrium justifies our moral judgments only if we assume that our judgments are justified by reflective equilibrium. And that is flagrantly circular. So we should not accept the view. It follows that we should not accept **Enoch's Conjecture**.

5.4 Conclusion

I ended Chapter 4 by claiming that moral intuitionism faced a dilemma, and because of this was unable to explain our successful moral judgments. This invited us to explore reflective equilibrium, which is the traditional opponent to intuitionism. Like intuitionism, reflective equilibrium is not necessarily guilty of inferring theoretical moral judgments from folk moral judgments. But (again like intuitionism) reflective equilibrium faces its unique own problems. In particular, reflective equilibrium turns out to be circular. According to the view, our moral judgments are justified when they are in agreement with general moral principles. But whether there is such agreement is a judgment itself, and so it too must be in agreement with a general principle. And that means that a moral judgment is justified only if we accept that reflective equilibrium justifies our judgments. But this is circular. Because of this, reflective equilibrium is not in a position to support nonnaturalist cognitivist realism.

CHAPTER 6
Using and Abusing Moorean Arguments

In Chapter 3 I argued that the folk/theoretical distinction is helpful for understanding the Not All Judgments (NAJ) response to the evolutionary debunking argument (EDA). This particular distinction allows us to flesh out the NAJ response and identify *which* moral judgments are threatened by our evolutionary history. We also saw that availing ourselves of the folk/theoretical distinction has the benefit of revealing the proper scope of the EDA. Rather than undermining our folk or theoretical moral judgments themselves, the EDA is better understood in the following way:

Inference Debunking If our theoretical moral judgments are inferred merely from our folk moral judgments then those theoretical moral judgments are made suspect because of evolutionary considerations.

When it is characterized in this way, the EDA has a very specific scope. It does not preclude all ways of justifying our moral judgments, nor even justification compatible with nonnaturalist cognitivist realism (as debunkers usually argue). Rather, it undermines a specific way of justifying our moral judgments. It rules out those moral epistemologies in which our theoretical moral judgments are justified merely by inferring them from our folk moral judgments.

In Chapters 4 and 5 we saw how moral intuitionism and reflective equilibrium were not susceptible to **Inference Debunking**. If we carefully and charitably characterize these two views then it becomes clear that neither view merely infers theoretical moral judgments from folk ones. Nevertheless, both of these views succumbed to other problems, and so both were unable to help explain how we can have successful moral judgments. But if neither moral intuitionism nor reflective equilibrium are guilty of this fallacious inference then we should

be left wondering which views are. In this chapter I answer that question. I will argue that one view in particular, namely metaethical Mooreanism, is squarely targeted by **Inference Debunking**. Owing to necessary features of the position, metaethical Mooreanism is guilty of inferring theoretical moral judgments from folk moral judgments.

In section 6.1 I explain metaethical Mooreanism and set it up in opposition to its usual opponent: moral error theory. While I do not defend the plausibility of moral error theory, if the Moorean argument against it fails then epistemic moral scepticism (my preferred conclusion here) is thereby bolstered.[1] In the next two sections I dig deeper into the Moorean position in order to show its necessary features. I separate Moorean *claims* (section 6.2) from Moorean *arguments* (section 6.3). And in section 6.4 I show why metaethical Mooreanism violates **Inference Debunking**. To do so, I draw a distinction between two types of reasoning: *presumptive* and *dialogical*. I will show that the key Moorean claim relies on presumptive reasoning, but in order to fend off moral error theory the Moorean argument requires dialogical reasoning. Unpacking the reason why shows the way in which Mooreanism is subject to my construal of the EDA.

From this I draw two conclusions (section 6.5). First, like moral intuitionism and reflective equilibrium, metaethical Mooreanism cannot adequately explain *how* we have successful moral judgments. But since Mooreanism is not best understood as attempting to offer such an account in the first place this conclusion is fairly weak. However, the second conclusion is much stronger: Mooreanism cannot adequately explain *that* we do in fact have successful moral judgments. Since this is what the view sets out to do, metaethical Mooreanism fails by its own lights. It is thus no help to nonnaturalist cognitivist realism.

[1] Again, I have defended the claim that we do not have successful moral judgments, which is a sceptical argument about moral knowledge, rather than the claim that there are no moral facts, which is a form of moral nihilism. The latter position entails the former, but not vice versa.

6.1 What is Metaethical Mooreanism?

Like both intuitionism and reflective equilibrium, Mooreanism is not limited to moral matters. All three are perfectly general epistemic accounts. But unlike intuitionism and reflective equilibrium, Mooreanism is not best understood as a *positive* epistemic account. That is, Mooreanism does not offer an account of how we know that some particular claim is true. As I understand it, Mooreanism is chiefly a *response* to sceptical arguments. It is a view about why, in spite of arguments to the contrary, we can be confident that we *do* in fact know a particular claim. So stated, metaethical Mooreanism is not a view about how we know that 'torturing people merely for fun is wrong' but rather a view about why sceptical arguments notwithstanding we can be sure that we do in fact know it.

Though it is is not a positive epistemic account, Mooreanism is amenable to nonnaturalist cognitivist realism. While the view does not explain how we have successful moral judgments, it offers assurance that we do in fact have them. It is therefore helpful for bolstering the overall nonnaturalist account. If we should understand moral judgments according to *cognitivism* and *nonnaturalism*, and if we can be sure that we have moral knowledge (per Mooreanism), then we have successful moral judgments even if moral facts are nonnatural. Therefore, the chasm between our moral judgments and the moral facts evaporates. Of course, it would be preferable to give a positive epistemic account explaining exactly how we have moral knowledge, but failing that nonnaturalist cognitivist realist may content themselves with an argument concluding that we *do* have moral knowledge.

Moorean arguments have a readily identifiable structure. They typically take the following form:[2] if a revisionary claim (RC) is true then a Moorean claim (MC) is false, but since that Moorean claim is true the revisionary claim must be false. For simplicity we can schematize this argument as such:

[2] For a number of different schematized Moorean arguments, see McPherson (2009).

If RC then not-MC

But MC

∴ Not-RC

Why call an argument like this *Moorean*? Because we see this type of argument first used by G.E. Moore, who (in)famously argued against external world scepticism in the following way: if external world scepticism is true then I don't know that I have hands, but since I do know that I have hands (i.e. here is one hand, and here is another), external world scepticism is false (Moore 1939).

Despite being a forerunner to contemporary nonnaturalist cognitivist realists, it is noteworthy that Moore did not use a similar argument against revisionary metaethical theses (like moral nihilism). Rather, his nonnaturalism was a result of the Open Question Argument that we explored in Chapter 2. So we will depart from Moore in our analysis of Moorean arguments here. However, other philosophers have used 'Moorean' arguments to bolster certain metaethical views, and it is to these defenses that I turn.

The particular Moorean argument that I am interested in is one often used against a particular revisionary metaethical thesis the moral error theory.[3]

Moorean Argument Against Moral Error Theory (MAAMET)

If the moral error theory is true then it's false that some things are morally wrong

But some things are morally wrong

∴ The moral error theory is false

[3] See Dworkin (1996), Huemer (2005), Nagel (1997). For a similar argument, see Keller (2017). Keller argues that if moral error theory is true then there are no normative reasons based on love, but since there are normative reasons based on love the moral error theory is false. This shares a similar structure to MAAMET but it is importantly different. Since Keller offers arguments in favour of there being normative reasons based on love his argument against moral error theory should not be classified as Moorean. As Ill explain below, Keller is not merely relying on *presumptive* reasoning in his defense, but is giving *dialogical* reasoning for his key claim.

Note that MAAMET retains the same structure as the schematic version above. Note too that it structurally similar to Moore's argument against external world scepticism.[4] It can strike some people that a Moorean argument like MAAMET is fishy. But identifying exactly what the problem is, if indeed there is one, can be difficult. We can call someone who defends this sort of argument in one domain or another a *Moorean*, and the suspicious interlocutor a *Revisionist*.[5]

Revisionists sometimes say that Moorean arguments violate the philosophical spirit of following the argument wherever it leads. They also say that Moorean arguments are question-begging and flat-footed. Keith Lehrer (1971) says that the 'commonsensical' nature of a proposition (like 'some things are morally wrong') does not confer it with any special or positive epistemic status. Barry Stroud (1979) says that, whatever else we want to say about it, a Moorean argument should not be considered a 'philosophical' answer to the Revisionist's challenge. And Peter Unger (1974) says that even 'commonsense' beliefs cannot be held in the face of overwhelming evidence, such as that provided by the Revisionists deductively valid arguments.[6]

However, Mooreans have suggested that these arguments are justified (at least in part) by their widespread use in philosophy. If you don't think yourself a Moorean, consider the

[4] However, you can see that the consequent of the first premise in Moore's argument concerns *knowledge*, whereas that of MAAMET concerns *falsity*. The reason for the difference is because of the target revisionary view. Moore's target was external world *scepticism*, which is a view about knowledge of the external world. But the target of MAAMET is moral error theory, which is a view about moral ontology (whether there are moral facts, and not whether we *know* that there are moral facts). The difference, then, is warranted.

[5] These labels are not standard. Sometimes the debate is cast as between *Dogmatists* and *Sceptics*, as in Pryor (2000) and White (2006). Other times it is cast between *Mooreans* and *Sceptics*, as in Kelly (2005). Sometimes *Revisionists* are instead called the *Aw, Come On! crowd*, as in DeRose (1999). But I prefer the term *Moorean* to the term *Dogmatist* in order to avoid the latters pejorative connotations. And I prefer the term *Revisionist* to the term *Sceptic* because rejecting a Moorean argument need not amount to scepticism.

[6] These responses are typical of what William Lycan (2001) calls 'new skepticism'.

following argument: if Zeno's arguments against motion are true then it's false that things move, but things *do* move, so Zeno's arguments are false. Now, I suspect that you *do* think that things move. And I suspect that you are unable to say *exactly* what is wrong with Zeno's arguments against motion. So you're probably a Moorean after all, at least in this limited domain. David Armstrong says that "perhaps we have still not, after two and half thousand years, got to the full bottom of Zeno's brilliant arguments against the existence of motion ... But certainly Zeno should not persuade us that things do not move. Neither should anybody else" (Armstrong 1999: 79). In essence, the Moorean says that despite Zeno's arguments to the contrary we should continue to believe that things move. Even if we are unable to locate exactly what is wrong with Zeno's arguments, we are not (and should not) be persuaded by its conclusions. So giving a Moorean argument in response to Zeno's arguments against motion seems to be a legitimate philosophical response.

It is arguments like this that lead some Mooreans to accuse Revisionists of metaphilosophical inconsistency (Kelly 2005). Because Moorean arguments are employed widely in philosophy, Mooreans say that it is unfair to complain about their usage in principle. If we are happy to accept a Moorean argument against Zenos arguments against motion, then on what grounds can we complain about MAAMET? While it can be legitimate to complain about a *particular* Moorean argument (for example, one with a false premise), Revisionists at least have to accept that the Moorean argument *form* is legitimate. And Revisionists should indeed accept this, because the Moorean arguments above are just examples of *modus tollens*. Because they are deductively valid arguments, complaints can only be issued with respect to soundness.

So is MAAMET sound? Even if we are suspicious of the argument, in order to reject it we will need to be able to identify a false (or at least highly dubious) premise. The obvious target is the second premise: 'but some things are morally wrong'. Yet this seems like a perfectly legitimate claim. Who could possible disagree?

For one, moral error theorists, who deny that there are moral facts, and so deny that anything is in fact morally right or wrong.[7] There are many such accounts, importantly different in their details, but united by something like the following:

Argument for Moral Error Theory (AMET)

Our moral discourse and practice commits us to the existence of nonnatural moral properties

But there are no such moral properties

∴ Our moral discourse and practice is in error

There are two things to note here. One is that the first premise is similar (but not identical) to *nonnaturalism*.[8] The other is that the conclusion of this argument implies that the claim 'but some things are morally wrong' is false. Take a typical moral judgment, like 'torturing people merely for fun is wrong'. We saw in Chapter 2 why we should characterize this judgment according to *cognitivism* and *nonnaturalism*. Moral error theorists typically accept both claims. More precisely, while they usually accept that our moral discourse and practice is *committed* to the existence of nonnatural properties, they in turn deny that the universe *actually* contains any such properties. That is the 'error' according to moral error theory: our moral discourse and practice commits us to nonexisting entities, in the same way that our discourse about God or witches is also about something illusory.

Let the conclusion of AMET, or more accurately an *implication* of its conclusion, be the revisionary claim (or RC). It is *this* claim that MAAMET targets. As we saw above, the Moorean argument begins with a conditional: if RC then some particular Moorean claim (or MC) is false. In our case, MC is the claim 'but some things are morally wrong'. MAAMET

[7] For examples, see Garner (1994), Mackie (1977), Joyce (2001), Olson (2014), and Streumer (2017).

[8] Recall that *nonnaturalism* was about the irreducibility of moral properties to natural properties, rather than a claim about what our moral discourse and practice commits us to.

uses MC as its second premise. And by *modus tollens* it follows that RC (and, therefore, moral error theory) is false.

We can now explain why this Moorean argument might seem fishy. MAAMET centrally relies on a premise that explicitly conflicts with AMET. And moral error theorists presumably take themselves to have adequately dispensed with that central claim. After all, the conclusion of AMET implies that nothing is morally wrong. So when the Moorean not only reintroduces the claim 'but some things are morally wrong' but also uses it as their starring premise, it can easily elicit charges of question-begging from moral error theorists.

So why do Mooreans reintroduce that premise? I suggest that Mooreans think of it as a *folk judgment*. In Chapter 3 we saw that folk judgments are often characterized in the following way: "without explicit instruction in such areas, people seem to develop domain-specific ways of thinking about relatively bounded sets of phenomena" (Kiel 2010: 826). And it certainly seems that we have developed a specific way of thinking about our actions, namely that they are capable of being morally right or morally wrong. This way of characterizing our actions is not something that agents can (or need to) necessarily make explicit, but rather is something for which we can find evidence "from the practice of agents in making their judgments and navigating around the world" (Braddon-Mitchell 2004: 278). Of course, many folk are quite at ease making moral judgments, and they do not need philosophical training in order to do so.

I suggest that Mooreans think of these folks judgments as enjoying what we can call *presumptive support*. Presumptive support is a form of epistemic support. It can be thought of as *prima facie* support, or *pro tanto* support, or something similar. As we will see, there are commonsensical epistemic principles that explain why some claims (like the key Moorean one) should be thought of as having presumptive support. When Mooreans appeal again to the key premise ('but some things are morally wrong') it is because they think of that premise as having the required epistemic support. So it is not as if Mooreans are introducing some heretofore epistemically unsupported premise

As we have seen, if a judgment like 'but some things are morally wrong' is a folk judgment then it is not undermined by **Inference Debunking**. But if we appreciate that there is a *dialectic* between Mooreans and Revisionists, that there is a certain call and response occurring between these two camps, then it becomes clear that a judgment that begins as a folk judgment can turn into something else over time. In Chapter 3 we saw how *theoretical judgments* are in an important sense a development on and out of folk judgments. Whereas folk judgments can be tacit and implicit, theoretical judgments are characterized by their explicit acknowledgment and acceptance. Folk judgments simply help us navigate the world, and so they can be implicit. Theoretical judgments might also be navigational tools, but their distinguishing characteristic is that they are not *merely* implicit. Rather, they are judgments that can be analysed and scrutinized, which in turn makes them open to negotiation and, if need be, finetuning.

However, while the claim 'but some things are morally wrong' might begin as a folk judgment and enjoy presumptive support, it need not (and I will argue *does not*) forever stay as such nor retain that support. Indeed, I think that Mooreans must contend with a subtle shift to their key claim's epistemic status. When they reintroduce the premise as a *response* to the moral error theory they are no longer entitled to use it as a folk judgment that enjoys presumptive support. Instead, because they are reintroducing it as a response to a sceptical argument, they are now using it as a claim that needs to have dialogical support. Like presumptive support, dialogical support is a form of epistemic support. But unlike presumptive support, dialogical support is not *prima facie* support. If a claim has dialogical support, then it is has that support *in the face of* arguments to the contrary. By this I mean arguments that have been made, and not merely possible arguments to the contrary. This means that dialogical support is not something that a proposition simply has or doesn't haven't, but instead is something that is dependent on the context. Which context? One where a given counterargument has been made. That is what makes dialogical support importantly different than presumptive support.

But it is unclear that Moorean claims *can* enjoy such support. At a minimum, Mooreans have not explained that they do. And this is ultimately what is wrong with MAAMET. While MC can initially be justified presumptively, if MC is used within MAAMET it must be justified dialogically. Absent this latter justification, MC (within MAAMET) is unjustified. This is why MAAMET will fare poorly against the moral error theory and also fail by its own lights.

Whether or not a claim undergoes a shift in epistemic status, from requiring presumptive support to requiring dialogical support, is subtle but important. I think this shift in the required support has taken place within the dialectic between metaethical Mooreans and Revisionists, and that is precisely what renders the Moorean argument here unjustified. At one point, MC is justified it it has presumptive support. At a subsequent point, MC requires dialogical support to be justified. So in order to make this shift clear, I will spend some time unpacking the structural features of metaethical Mooreanism. In the next two sections I will clarify the key features of MAAMET by separating Moorean claims from Moorean arguments. This will allow us to see the role that the former plays in the latter. And, in turn, we will be able to more clearly see the shift that makes MAAMET subject to **Inference Debunking**.

6.2 Moorean Claims

There is an important yet under-appreciated difference between a Moorean *argument* and a Moorean *claim*.[9] The two are often not disambiguated, probably because their relationship is straightforward: a Moorean claim is simply one of the premises within a Moorean argument. To wit, the second premise of MAAMET ('but some things are morally wrong') is a Moorean claim.

[9] See Lewis (1999) and Kelly (2005) for a discussion of what they call Moorean *Facts*. See McPherson (2009) for a discussion of Moorean *Arguments*. See also Wittgenstein (1969) for a discussion of what he calls Moorean *propositions*, or *truisms*.

Separating the claims from the arguments serves us well. There are characteristics of each that are in need of examination. Let's first deal with Moorean claims. I think that they have two key features: (i) they are basic beliefs, and (ii) they play a critical role in constraining arguments.

6.2.1 Basic Beliefs

Let's focus on the following two Moorean claims:

Hands I have hands

Motion Things move

Moorean claims are typically characterized by their purported *basicness*. Whatever else we believe, we should at least believe **Hands** and **Motion**. These function as basic beliefs. To say that a belief is basic is to say that it is not inferred from other beliefs but is instead "self evident, incorrigible, or evident to the senses" (Plantinga 2000: 84).

We can understand basic beliefs as the expected default presumptions that we all share, as claims that are as evident to us as we expect them to be to others.[10] These basic beliefs are difficult to shake, but they are in principle capable of being changed. Even so, we should not draw an equivalence between basic beliefs and Moorean claims. While all Moorean claims are basic beliefs, not all basic beliefs are Moorean claims.

Why think that Moorean claims are basic beliefs in the above sense? Take **Hands**. One reason to think that this is a basic belief is because it is evident to the senses and incorrigible (in the weak sense of 'exceedingly difficult to shake'), and one reason to think that it is justified is because we gain that belief via our perceptual experiences (which I take

[10] This understanding of basic beliefs departs slightly from Plantinga's. He often discusses a belief as being properly basic *for a particular person*. For example, a belief in God can be properly basic for Alvin, even if it is not for me. I do not index basicness to an individual, but instead talk of basicness *simpliciter*.

to provide defeasibly justified beliefs). Remember that Moore's 'proof of the external world' involves him pointing his hands at each other in order to demonstrate their existence: here is one hand, and here is another. Moore seems to think that this licenses the belief **Hands**.

Now take **Motion**. This too is a belief that is not inferred from another belief, rather it is an incorrigible belief that is evident from our experiences. Both **Hands** and **Motion** are beliefs that are evident to us, and that we reasonably expect to be evident to everyone else as well. Even when confronted with counterarguments, like external world scepticism or Zeno's arguments against motion, it seems that we cannot shake **Hands** and **Motion**.

We do not have to wade into debates on whether (and how) our perceptual experiences can contribute to justified belief.[11] We can for now adopt the following commonsense stance towards our perceptual appearances: for claims like **Hands** and **Motion** "there is a natural presumption in favor of taking the appearance at face value" (White 2006: 527). The notion that we have a 'presumption in favour of taking appearances at face value" will be important going forward, and we can characterize this presumption in the following way:

Appearance It appears to S that P

Justification If **Appearance**, then S is (defeasibly) justified in believing P

Why accept either **Appearance** or **Justification**? For the following reasons. First, it is uncontroversial that there are things that appear to us. It appears to me that there is a roof over my head; it appears to me (as I expect it does to you) that the universe is more than five minutes old; etc. That much is straightforward. What **Justification** says is that if it appears to us that the universe is more than five minutes old, then we are (defeasibly) justified in believing that it is. Not *always* justified, or *all things considered* justified, but just *defeasibly* justified. From the mere appearance that p we can gain a minimal sense of

[11] But see, for example, Gupta (2006), who argues that perceptual experiences can generate an *entitlement* to form a justified belief (and eventually knowledge). See Berker (2011) for a critical response.

justification for believing that p. If we *didn't* accept something like **Justification** then we would really struggle to have *any* justified beliefs.[12]

The move from **Appearance** to **Justification** provides what we can call *presumptive* support. As said, this is a form of epistemic support. What makes the support here 'presumptive' is that it is support that we have *in the absence of further information* — i.e., in the absence of epistemic defeaters for that appearance. As I said, we might think of presumptive support as *ceteris paribus* support. We might also think of the justification that presumptive support can provide as *prima facie* justification. The exact labels are, I think, inessential. What is important is that if or when *more information* becomes available this presumptive support *might* be defeated — and at a minimum we should be open to assessing *whether* it has been defeated. In that case, we will need to assess whether the support we thought our belief enjoyed still holds.

6.2.2 Constraining Arguments

The second feature of Moorean claims is that they constrain philosophical arguments. This is standardly characterized in the following way: because of Moorean claims there are necessarily limits to where philosophical arguments can lead. In particular, arguments will not be able to lead to the *denial* or *rejection* of a Moorean claim.

The constraining influence of Moorean claims is widely held as their defining feature. Thomas Kelly says that a Moorean claim is a claim that cannot be undermined by a philosophical argument (Kelly 2005: 180). David Lewis says that a Moorean claim "is one of those things that we know better than we know the premises of any philosophical argument to the contrary" (Lewis 1999: 418). Moore says that he is justified in believing **Hands**

[12] The notion of a 'presumption in favour of taking appearances at face value' is also similar to what Michael Huemer (2001) calls *phenomenal conservatism*. This view says that if it seems to you that p, then you are prima facie justified in believing p. For an overview of phenomenal conservatism, see Moretti (2015).

because it seems to me more certain that I do know [it] rather than any sceptical alternative (Moore 1959: 226).

These are all different locutions, but they are compatible. In saying that a Moorean claim cannot be undermined by philosophical argument Kelly takes himself to be identifying a genuine constraint. Whatever else philosophical arguments can achieve, they cannot undermine Moorean claims.[13] On the other hand, when Lewis and Moore say that one knows a claim *better than* an alternative, or that one is *more certain* of it as compared to something else, they are each making an expressly comparative claim.[14] This comparative claim is importantly different than what Kelly says, since Kelly does not identify a contrast class of propositions.[15]

Nevertheless, the comparative claims that Lewis and Moore each present seem to entail the constraining approach to inquiry that Kelly suggests. We can see this by thinking the

[13] There are two ways of understanding this constraint. If C is a Moorean claim, one might say that C is not undermined by philosophical argument simply because *there is currently not* a philosophical argument for not-C that is rationally compelling. Or, instead, one might say that C is not undermined by philosophical argument because *there cannot be* a philosophical argument for not-C that is rationally compelling. Call the former view *Open Mooreanism* because it allows (in principle) for C to be undermined — that is, it is *open* to assessing whether a given philosophical argument undermines C. While there is no philosophical argument that is currently able to do so, perhaps a *future* philosophical argument could successfully do so. And call the latter view *Closed Mooreanism* because it denies this, suggesting instead that the case is now and forever *closed*; philosophical argument is in principle incapable of undermining Moorean claims. Regardless of whether one accepts *Open* or *Closed* Mooreanism, both are clear constraints on philosophical inquiry.

[14] Michael Huemer (2005) also argues for a comparative claim. He says that of two equally valid arguments, "the better is the one whose premises are more initially plausible (Huemer 2005: 116). In order to reject a Moorean claim one will therefore "have to produce premises more plausible than any moral judgment — more plausible than 'Murder is wrong', more plausible than 'Pain is worse than pleasure', and so on' " (Ibid., 117).

[15] While it is natural to read Kelly as suggesting that a Moorean claim *is not* undermined by a sceptical argument (Open Mooreanism), one can also read him as suggesting that a Moorean claim *cannot be* undermined *full stop* (Closed Mooreanism).

following: if one knows C better than not-C, then one might be justified in dismissing an argument for not-C. For example, Jeremy Fantl says that "some of us don't have to feel so guilty about flatly dismissing [an] opponents arguments ... [because] flat dismissal of arguments against [controversial] propositions is often legitimate" (Fantl 2013: 36). Indeed, this is why Mooreans often say that philosophical inquiry must be tethered to basic beliefs. It is why Lewis suggests that we should build our philosophical theories around certain fixed points, and why Kelly reads him as suggesting that "we should make adjustments elsewhere as needed in order to hold on to [particular] fundamental commitment[s]" (Kelly 2005: 1).[16]

What should we make of the suggestion that this issue is fundamentally comparative, that we need to assess whether one claim is more compelling or intuitive than another? Its obvious that we can have varying degrees of certainty about our beliefs. We might want to accept that ones certainty or confidence in a given proposition provides *defeasible* justification for believing that proposition. However, it seems less obvious that ones certainty or confidence in a specifically Moorean claim will always enjoy (or always *continue* to enjoy) justification on those grounds.

So those are the two features of Moorean arguments: they are basic beliefs, and they constrain arguments.

6.3 Moorean Arguments

How do Moorean claims relate to Moorean arguments? In particular, what role does a Moorean claim play within a Moorean argument like MAAMET? To see the role, we will need to understand how a Moorean claim contributes to the conclusion of a Moorean argument. The easy answer is that Moorean claims are simply a premise within Moorean arguments.

[16] The idea of 'holding onto a fixed commitment' is compatible with both Open and Closed Mooreanism. That is, if there are no current not-C arguments that are rationally compelling, then that will serve as grounds for holding onto a particular fundamental commitment. Ditto if one anticipates that any not-C argument will fail to be rationally compelling.

But this easy answer obscures the special epistemic status that Moorean claims are thought to have. For that reasons, we will have to explore the issue in more detail.

6.3.1 Justification and Moorean Arguments

In the previous section we saw why a Moorean claim can be justified: because it is a basic belief, and because basic beliefs can have justification owing to presumptive support. But that justification stands apart from the Moorean argument itself. Instead, it lurks in the background. To bring it out we will need to introduce a second argument. Recall the initial schema:

If RC then not-MC

But MC

∴ Not-RC

In the previous section we saw why we should think that one can have a justified belief that MC. This is because MC can be a basic belief. What we now want to know is how this information helps us support MAAMET. Recall:

If the moral error theory is true then it's false that some things are morally wrong

But some things are morally wrong

∴ The moral error theory is false

Here is one way to see the Moorean claim and the Moorean argument relating. We can accept something like the following: if one is justified in believing a Moorean claim, and one can tell that believing the Moorean claim entails that one should believe that some revisionary claim is false, then one is justified in believing that the revisionary claim is false.

Notice how this relies on the *justification* one has for believing the Moorean claim rather than relying on the mere content of the claim itself. Notice too that the justification one has

for believing the Moorean claim appears to *transmit* to the justification one has for denying the revisionary claim.[17] This explanation is an instance of a familiar epistemic principle:

Justification Closure (JC) if S is justified in believing P, and S can tell that P entails Q, then *ceteris paribus* S is justified in believing Q

Strictly speaking this is **Justification Closure** construed as 'closure under *known* entailment'. If instead **JC** was simply closure under entailment we would not have to identify that S can tell that P entails Q. It would suffice that P entails Q, irrespective of S's identification of this entailment. But we should prefer closure under *known* entailment in order to account for ignorance, on the part of S, that P entails Q, or worse, a rejection of the entailment. Consider: if we accepted just closure under entailment, then if S knows P, and if P entails Q, then S could explicitly reject that P entails Q but still be characterized as knowing Q. Closure under known entailment avoids this problem, and is therefore preferable.[18]

I think that **JC** can help us make sense of how Moorean claims contribute to Moorean arguments. Note how **JC** seems to be lurking behind most Moorean arguments. Moore thinks that he is justified in believing **Hands**. He also thinks that this belief is justified, and he seems to (implicitly) think that if this justified belief entails that the belief in the sceptical conclusion is false, then he is justified in believing that the sceptical conclusion is false. In doing so, Moore takes himself to be justified in denying that external world sceptics revisionary claim.

Now suppose that **Hands** is indeed (defeasibly) justified. If **Hands** is justified then what **JC** does is explain how that justification can transmit to another claim, namely the

[17] For the previous explanation of justification transmission, see 4.2.

[18] Other possibilities here are justification under *actual* and *justifiably* believed entailment. For more, see Luper (2016).

denial of the revisionary claim. And the way it does so is through an argument. Consider the following:

(i) If S is justified in believing MC, and S knows that MC entails not-RC, then S is justified in believing not-RC

(ii) S is justified in believing MC

∴ S is justified in believing not-RC

If this is right, then we are in a position to see how the Moorean claim (MC) helps to support MAAMET:

(i*) If S is justified in believing that some things are morally wrong, and S knows that this entails that moral error theory is false, then S is justified in believing that moral error theory is false

(ii*) S is justified in believing that some things are morally wrong

∴ Therefore, S is justified in believing that moral error theory is false

When MAAMET is thought of as supported by this argument the result is a better argument. It is better because it makes explicit how justification is thought to transmit to the conclusion of the argument. Without **JC** in the background, MAAMET can look too much like a flat-footed denial of moral error theory. But Mooreans should try to avoid this appearance. Towards this end, drawing out **JC** and making it explicit is helpful.

So far we have separated Moorean claims (like 'but some things are morally wrong') from Moorean arguments (like MAAMET). The reason that we did this was so that we can identify the important features of each component, and the relationship between the two. One important feature of Moorean claims is that they are plausibly construed as basic beliefs. These beliefs enjoy justification because we have a presumption in favour of taking appearances at face value. This in turn makes it plausible to interpret a Moorean claim as a folk judgment. We also saw that one important feature of Moorean arguments is that they use Moorean claims as their star premise. Moorean arguments have another argument lurking

in the background, namely one that uses justification closure to transmit the justification for the Moorean claim to the denial of a revisionary claim. With all of these pieces in place, we are now in a position to see why MAAMET violates **Inference Debunking**.

6.4 MAAMET and Inference Debunking

At one stage, the Moorean claim 'but some things are morally wrong' can be plausibly viewed as a folk judgment that enjoys presumptive support. It is a judgment that laypersons will make that will help them navigate their environments and which is justified by the presumption in favour of taking appearances at great value. I take it that there is no great mystery as to whether folk judgments can in principle be justified, and that usually the meatier question is whether a particular folk judgment is *true*. But let's focus on justification. In our case, the Moorean claim enjoys justification because of what I have called *presumptive* reasoning. Because we have a presumption in favour of taking appearances at face value, and because it certainly appears as if some things are morally wrong, I think we can say that the claim 'but some things are morally wrong' can (at least initially) be justified. The justification for this folk judgment stems from our presumption in favour of taking appearances at face value.

However, folk judgments often give way to theoretical judgments. Where the former are implicit or tacit judgments that we make which help us navigate our environments, the latter (while often still navigational tools) are marked by their explicitness. One reason that folk judgments give way to theoretical ones is because the latter are judgments that can be analysed and scrutinized, which means that they are open to negotiation and perhaps fine-tuning. This allows us to decide which of our judgments we wish to endorse, assess whether any are in tension with others, and so on. Our tacit judgments (being as they are 'mere' navigational tools) are not suited to these tasks.

So folk judgments can give way to theoretical judgments. And, importantly, presumptive support, by its very nature, is open to being defeated. As we saw above, Moorean claims

play a key role within Moorean arguments. I explained that Moorean claims are basic beliefs that enjoy presumptive support because we are licensed to take appearances at face value. But what happens when we are presented with arguments specifically *against* those same appearances? If it seems that some things are morally wrong, and if we are presumptively justified in believing this on that basis, then we need to take seriously attempts at overturning the initial appearances. I suggest that when our presumptively justified beliefs are challenged in this way, we are forced into a position where we must move on from relying on presumptive support and towards relying on dialogical support. When there is a philosophical argument against one of our presumptively justified beliefs, in engaging with that argument we are accepting certain terms of debate. One such term is that we will provide a philosophical defense of our belief. And this, as I understand it, is precisely what MAAMET is: it is a philosophical argument against another (revisionary) philosophical argument (namely AMET), an argument for why we can continue to believe the Moorean claim.

But a problem for metaethical Mooreanism arises here. As we saw above, Moorean claims play a starring role with Moorean arguments, and Moorean arguments purport to transmit the justification of the Moorean claim to the denial of the revisionary claim. The problem, however, is that the epistemic status of the Moorean claim undergoes a subtle shift during this process. If one happens to believe that 'some things are morally wrong', then this is likely a folk judgment. But if one uses that judgment within a philosophical argument — in particular, within a metaethical argument *against* a revisionary metaethical position — then one cannot be said to be relying on a 'mere' folk judgment any longer. Rather, the judgment here requires a different status. It is a judgment that is explicit, which (as we said) is how we mark the distinction between folk and theoretical judgments.

However, while a *folk* judgment might enjoy justification because we have a presumption in favour of taking appearances at face value, a *theoretical* judgment cannot be justified in the same way. Thinking so is a violation of **Inference Debunking**.

I am not suggesting that a claim that enjoys presumptive support is once-and-for-all defeated simply because there is a counterargument against it. The mere presence of AMET does not rule out that the Moorean claim can be justified. Rather, I am suggesting that AMET eliminates the possibility that the Moorean claim can remain *presumptively* justified. Recall that presumptive support is what we have in the absence of further information. Thus, since AMET is not only further information, but also an argument specifically against initial appearances, the Moorean claim loses its presumptive support. However, it can gain other support, namely *dialogical* support — but only if Mooreans explains why and how.

If this is right then we can see how a problem arises for metaethical Mooreanism. As we saw above, Moorean claims play a starring role within Moorean arguments, and Moorean arguments purport to transmit the justification of the Moorean claim to the denial of the revisionary claim. The problem, however, is that the required epistemic status of the Moorean claim undergoes a subtle but critical shift during this process. If one initially believes that 'some things are morally wrong', then this is presumably licensed by the presumption in favour of taking appearances at face value. But if one uses that judgment within a metaethical argument *against* a revisionary metaethical view, then one cannot be said to be relying on a 'mere' presumption any longer. The justification we gain from initial appearances will no longer do, since AMET targets those initial appearances. Rather, the Moorean claim here requires a much different type of epistemic support. It requires dialogical support

Because the contents of the folk and the theoretical judgments are identical, the the shift in required epistemic status is easily missed. So here is perhaps one way to bring out the distinction between the two types of epistemic support. Let **FJ** be a folk judgment that is justified presumptively, and **TJ** be a theoretical judgment that is justified dialogically:

FJ 'some things are morally wrong'

TJ 'some things *are* morally wrong'

It is plausible that **FJ** is a folk judgment, that it is something we implicitly believe, that is justified because of our presumption in favour of taking appearances at face value and so on. But those things are not true of **TJ**. It is not plausible that **TJ** is a folk judgment, because it is a *response* to a philosophical argument. It is not the case that **TJ** is something we *implicitly* believe, rather it is something we explicitly state. And most importantly, **TJ** cannot be justified simply because we have a presumption in favour of taking appearances at face value — it must be licensed by something *besides* that presumption.

We can certainly accept that **FJ** is at one point justified because it is a folk judgment. But what the Moorean owes us now is an explanation for why **TJ** (which seems to have the same content) is also justified, and moreover justified as a *response* to a sceptical argument. If **TJ** is justified because we infer it from **FJ** then **Inference Debunking** rears its head. But the Moorean needs this inference to go through, otherwise the key Moorean claim will lack justification entirely, and there will be no justification lurking behind the Moorean argument. To see why, lets look more closely at the dialectical between the Moorean and the revisionary opponent.

6.4.1 Dialogical Reasoning

When faced with a revisionary or sceptical thesis, like the moral error theory, some Mooreans will respond in the following way: if we believe the Moorean claim ('some things are morally wrong'), and if we have no reason to accept a competing revisionary thesis, then we are justified in believing the Moorean claim. There is a subtle difference here as compared to the presumptive reasoning in favour of taking appearances at face value, but I think the difference is important.

On the one hand, to say that we should take appearances at face value is just to say that we have a default *presumption* in favour of thinking something is true if it strikes as us true. Presumptive reasoning about appearances can provide defeasible justification regarding our beliefs about those appearances. On the other hand, when we say that we are justified in

believing a Moorean claim *if we have no reason to accept a revisionary claim*, we are making a *dialogical* point. In this latter case, we are highlighting that there is a *response* that we think we are licensed to give when confronted with a revisionary or sceptical argument.

It is important to recognize that there is a dialogical *process* involved in this latter case. If there is justification in the former case it is achieved presumptively. But this justification can be defeated, specifically by an argument that explains away those appearances. When this happens, if justification is going to be regained it will have to be because a claim has dialogical support. This latter form of support is one that can only be gained when faced with arguments or when pressed to give responses

We can adopt the following characterization for this dialogical support: if it appears to S that p, and S has no reason to suspect that any revisionary or sceptical alternative to p is true, then S is justified in believing p (White 2006: 527).[19] It is important here that by 'any' alternative we at least restrict this to any argument that has been made. Working from this characterization, let us separate out the following claims:

Appearance It appears to S that p

No Reason S has no reason to suspect that any revisionary alternative to p is true

Justification* If **Appearance** and **No Reason** then S is justified in believing p

The inclusion of **No Reason** makes the argument different than the presumptive reasoning that we looked at above. What does it mean to say that S has *no reason* to suspect that any revisionary claim is true? Plausibly, that there is no revisionary claim that is rationally

[19] White (2006) calls this *Dogmatism*. But, as said, in order to avoid pejorative connotations I eschew the term in favour of *Mooreanism*. And see Huemer (2007) for a more recent defense of phenomenal conservatism, in particular one that explicitly mentions the presence of defeaters.

compelling.[20] Whether that is the case or not depends on the quality of the arguments involved.

If we have presumptive support for a claim then we can easily move from **Appearance** to **Justification**. In the absence of further information, we are justified in believing that things are just as they appear to be. But when we include **No Reason**, if we are to get **Justification*** it must be because we are specifically entertaining or evaluating other possibilities. What needs to be stressed here is the difference between **Justification** and **Justification***, and how it is **No Reason** (or the presence of defeaters) that makes the latter different. The former relies on presumptive support, which is something we have in the absence of further information, and the latter on dialogical support, which is something we gain only after the back and forth of arguments. And the reason the difference between **Justification** and **Justification*** is important is because it shows how MAAMET violates **Inference Debunking**.

The Moorean claim 'but some things are morally wrong' (or **FJ**) was justified presumptively. Because it appears to us that some things are morally wrong we are thereby justified in believing that some things are morally wrong. But the similar claim 'but something things *are* morally wrong' (or **TJ**) cannot be justified presumptively. Rather, it must be justified dialogically. This is again because **TJ** is being given as a *response* to a sceptical argument (namely AMET). So the Moorean must offer an explanation for why **TJ** is justified here, because they cannot say that **TJ** is justified because **FJ** is justified. To do so would be a violation of **Inference Debunking**. It would ignore the fact that **TJ** is being given as a

[20] Plausibly, **No Reason** means that the revisionary or sceptical alternative to p is not rationally compelling. This would entail that S has no reason to suspect that the alternative to p is true. Note how this understanding of **No Reason** is compatible with both Open and Closed Mooreanism. Recall that Open Mooreanism says that there is no known philosophical argument for not-C that is rationally compelling. And Closed Mooreanism says that there will be no philosophical argument for not-C that is rationally compelling.

response to a sceptical argument. So **TJ** is going to have to be justified in some other way, and not because we have a presumption in favour of taking appearances at face value.

Though this seems like a challenge that Mooreans might be able to meet, there is a serious problem awaiting them. Recall that one of the key features of Moorean claims is that they are basic beliefs, and that such beliefs gain justification because (all else being equal) we can believe that things are just as they appear to us to be. But if the Moorean cannot rely on that particular Moorean claim and its attendant justification when confronted with AMET then what can they use in their defense? If they double-down and say that **TJ** is justified because **FJ** is justified then the move is subject to **Inference Debunking**. So they will have to provide some other reason to think that **TJ** is justified.

Let me be clear about what I take this to demonstrate. It does not show that the key Moorean claim ('but some things are morally wrong') cannot be justified. Rather, it shows that while a Moorean claim can be initially justified (i.e., presumptively) because we are licensed to take appearances at face value, when it is given in response to AMET that claim can no longer be justified in that same way. This is because presumptive support only works *in the absence of further information*. But further information has been given, in the form of AMET. AMET is specifically an argument *against* the key Moorean claim, and so an argument against initial appearances. So now, if the Moorean claim is going to regain its justificatory standing then it must not be presumptively but rather dialogically. The Moorean must provide some other argument in favour of the justification of the key Moorean claim, because the presumption in favour of taking appearances at face value no longer holds. But Mooreans have not provided these arguments.

In light of that, we should think that metaethical Mooreanism cannot adequately explain *that* we do in fact have successful moral judgments. The view needs to be able to explain why Moorean claims are justified. Ordinarily it can do so by noting that Moorean claims are basic beliefs, and that basic beliefs are justified via a presumption in favour of taking appearances at face value. But this avenue of explanation is not available for MAAMET.

Because MAAMET is a *response* to a sceptical argument, it cannot rely on the presumptively justified Moorean claim. To do so would be to flatly ignore that it is being given in *response* to a sceptical argument. So the Moorean owes us an explanation for why we can continue to think that the key Moorean claim is justified. But Mooreans have not offered this explanation. And it is doubtful that they can.

6.5 Conclusion

There are two conclusions to draw from the above discussion. First, like moral intuitionism and reflective equilibrium, metaethical Mooreanism is not in a position to explain *how* we have successful moral judgments. But, as I have said, metaethical Mooreanism is not best understood as attempting to offer a positive epistemic account in the first place. So this conclusion should come as no surprise.

However, the second conclusion is a bigger problem. I have shown that metaethical Mooreanism cannot adequately explain *that* we do in fact have successful moral judgments. The view needs to be able to explain why Moorean claims are justified. Ordinarily it can do so by noting that Moorean claims are basic beliefs, and that basic beliefs are justified via a presumption in favour of taking appearances at face value. But this avenue of explanation is not available for MAAMET. Because MAAMET is a *response* to a sceptical argument, it cannot merely rely on the presumptively justified Moorean claim. To do so would be to violate **Inference Debunking**. So the Moorean owes us an explanation for why we can continue to think that the key Moorean claim is justified. But Mooreans have not offered this account.

Metaethical Mooreanism set out to show that, though we may not be able to explain *how* we have successful moral judgments, we can nevertheless be sure *that* we have them. I have argue that metaethical Mooreanism fails in this task. It is thus no help to nonnaturalist cognitivist realism.

CHAPTER 7
Conclusion

With the previous chapters I have argued that we do not have a plausible moral epistemology for nonnaturalist cognitivist realism. And because we do not, we should accept epistemic moral scepticism. Here is a brief recap of how I arrived at this position.

I began by suggesting that nonnaturalist cognitivist realism is the view to beat in metaethics. This is the view that best explains our moral discourse and practice. From this we saw that there are two important desiderata for a plausible moral epistemology: *cognitivism* about moral judgments, and *nonnaturalism* about moral facts. I argued that we should not be satisfied if nonnaturalist cognitivism realism says simply that there *are* such moral facts. Rather, we must have a plausible account of how we can *know* those facts — that is, how we have *successful* moral beliefs about nonnatural moral facts. Lacking this we should reject nonnaturalist cognitivist realism and accept epistemic moral scepticism.

I then explained why a general epistemic challenge — the evolutionary debunking argument — does not threaten nonnaturalist cognitivist realism. Only views that violate what I called **Inference Debunking** are ruled out. From there I examined two positive epistemic accounts. I argued that while neither moral intuitionism nor reflective equilibrium were guilty of **Inference Debunking**, neither can plausibly explain how we have successful moral beliefs about nonnatural moral facts. Then I looked at another way of vindicating nonnaturalist cognitivist realism, one without a positive epistemic account. Metaethical Mooreanism is a view that says even if we cannot explain *how* we know the moral facts, we can nevertheless be sure that we *do* know them. I showed how Mooreanism fails even with respect to this more limited aim, because it violates **Inference Debunking**. And on

the basis of these failures we seem to run out of explanatory options. Therefore, we are compelled to adopt a sceptical position regarding moral knowledge.

7.1 Some Objections

It might go without saying that the foregoing should not be thought of as a *proof* of epistemic moral scepticism. It is of course a *defense* of it. And because it is a 'mere' defense there are of course ways to counter the sceptical conclusion. I think we can identify at least a few ways of responding here. I will briefly mention three, and then I will respond in detail to a fourth. To be clear, I think I have already defended against each of the first three ways, but they nevertheless remain points at which one might disagree.

First, you might reject one or both of the desiderata that I said best explain our moral discourse and practice. If so, it could follow from this that the failure to explain how we have successful moral *beliefs* about *nonnatural* moral facts is not problematic, precisely because we should not be attempting to characterize either moral judgments as beliefs or moral facts as nonnatural. While I think the available metaethical arguments favour both cognitivism and nonnaturalism, you might instead think that there are better arguments against those positions. Because Chapter 2 was simply an overview of the metaethical landscape, there are avenues available for rejecting epistemic moral scepticism. In particular, noncognitivists and naturalists will not have seen any new arguments against their positions, so whatever has led them to accept their views in the first place likely remains convincing.

Second, even if *cognitivism* and *nonnaturalism* do in fact best explain our moral discourse and practice so far as I've construed it, you might think that I should have taken a much broader view of morality. Here one might take issue with my focusing on *moral judgments* in particular. Indeed, some moral philosophers have argued that focusing solely

on moral judgments obscures other important features of our moral discourse and practice.[1] If this is right, then even if our moral *judgments* are best explained according to *cognitivism* and *nonnaturalism*, our moral discourse and practice might be much more varied than my focus on judgments alone allows.

And third, you might reject the claim that we need an account of *how* we have successful moral judgments at all. I said that nonnaturalist cognitivist realism can only vindicate our moral discourse and practice if it can explain how we know the moral facts. This introduced an epistemic notion regarding the success of our moral judgments. But you might instead think that nonnaturalist cognitivist realism stands or falls based only on ontological issues — that it only matters whether or not there *are* moral facts. As I said, I have set aside questions of ontology and was willing to at least assume that there are nonnatural moral facts. But if you take the ontology-only approach then you might then think that the important questions regarding nonnaturalist moral realism are thereby settled, and that epistemic questions are beside the point.

I think that all of these avenues are technically open (if difficult to travel) for those who insist on rejecting epistemic moral scepticism. I grant that the preceding chapters have not shown that these options (and perhaps others) are closed. That is, I have not offered knockdown arguments for why *only* nonnaturalist cognitivist realism is tenable, why the two desiderata *alone* must be focused on, or why the focus on knowledge is *itself* the key to vindicating a metaethical view. Rather, what I take myself to have done is provide good arguments in favour of each of those claims, and I do not think philosophical arguments can aspire to much more than that.

[1] See for example Alice Crary (2007). She argues that rather than focusing on moral judgments in particular, we should focus on so-called 'moral thinking' more generally (Crary 2007: 44). Crary distinguishes moral thinking from moral judgment by saying that the former includes both moral and non-moral concepts, whereas moral judgments are narrowly construed as involving only applications of moral concepts. However, see Kirchin (2008) for a critical response.

And because I have discussed each of those avenues there is no point in going over them again. Instead, by way of concluding the dissertation, I will raise and address yet another challenge — one that concerns the scope of the preceding arguments. Crucially, whereas the three challenges just mentioned are ways of rejecting key parts of the preceding chapters, the final challenge I will address is one that arises *even if* you think that all of the chapters are successful in their stated aims.

7.2 The Over-Generalization Worry

Suppose that you accept the case I have made for why we lack moral knowledge, so you now think that when we canvass the options for how we could know nonnatural moral facts we come up empty-handed. You might then wonder about the extension or scope of that argument. Does the sceptical argument regarding moral knowledge affect *only* moral knowledge, or does it also extend to our knowledge of some non-moral matters? The threat here is that the preceding sceptical arguments are too broad, taking down much more than moral knowledge, and that this in turn threatens the plausibility of the initial case for epistemic moral scepticism. We can call this the *over-generalization* worry.

As I argued, if moral facts are nonnatural properties then we lack a plausible epistemology to explain how we can know anything about them. This leads to epistemic moral scepticism, the view that we do not have moral knowledge. But what other properties are nonnatural? According to some philosophers, certain epistemic (or *meta*epistemic) properties are nonnatural.[2] In that case, do we lack knowledge about them as well? Chris Heathwood presents the problem in the following way:

> Just as there are properties like that of an acts being morally wrong or a state of affairs being intrinsically good, there are the properties of having a reason to believe something or of a beliefs being epistemically justified or reasonable.

[2] Cuneo (2007), Enoch (2011), Parfit (2011), Wedgwood (2007), etc.

> About these epistemic properties, the same questions arise. Can they be reduced to properties from some other domain? (Heathwood 2018: 47).

If some epistemic properties — like something being reason to believe p, or the fact that S's belief that p is epistemically justified — are not identical to any natural properties, then those epistemic properties might be nonnatural. And if epistemic properties are nonnatural then, just as it was for nonnatural *moral* properties, it seems we do not not have knowledge of them either.

So the over-generalization problem is this: a defense of epistemic moral scepticism works equally well as a defense of *normative* or *metanormative* scepticism. According to these latter forms of scepticism, we do not have knowledge (or perhaps justified belief) regarding any normative domain. It is presupposed here that the moral and the epistemic domains are united insofar as they are *normative*: they concern what we ought, should, or have reason to do, believe, etc.[3] And if we characterize normativity in nonnatural terms, then any scepticism about the normativity of morality *ipso facto* carries over to the normativity of non-moral matters as well. That is how a defense of epistemic moral scepticism might work equally well as a defense of normative or metanormative scepticism. In turn, if the defense of epistemic moral scepticism is really a defense of normative scepticism, then that might threaten the plausibility of epistemic moral scepticism itself.

There are I think two ways to express the over-generalization worry — one in terms of 'biting bullets' and the other in terms of 'self-defeat'. First, if accepting epistemic moral

[3] Broome writes that "in one sense 'normative' simply means to do with norms, rules or correctness. Any source of requirements is normative in this sense. For example, Catholicism is. Catholicism requires you to abstain from meat on Fridays. This is a rule, and it is incorrect according to Catholicism to eat meat on Fridays. So Catholicism is normative in this sense. But I do not use 'normative' in that sense. In my sense, it means to do with ought or reasons. Given a rule or a requirement, we can ask whether you ought to follow it, or whether you have a reason to do so. Is there any reason for you to abstain from meat on Fridays, for instance?" (Broome 2007: 162-163).

scepticism also requires us to accept normative or metanormative scepticism then this might be too big a bullet to bite. It might be difficult or unpalatable to accept that we do not have moral knowledge, and so we do not know whether torturing people merely for fun is wrong. But it can be even more problematic to accept that we do not have knowledge of what would or would not be reasonable to believe, which beliefs are or are not epistemically justified, and so on. Adjusting our view of the world in light of moral scepticism would be dramatic, but not as dramatic as adjusting it in light of normative or metanormative scepticism. It is in that sense that normative scepticism seems to be too big a bullet to bite.

But I do not think that characterizing the over-generalization worry in this way constitutes a serious philosophical challenge. Which bullets we are comfortable biting should of course be secondary to which arguments are correct and for which we have good reasons to accept. If the arguments presented for why we lack moral knowledge are sound then we should accept them. And if those same sound arguments then compel us to accept normative scepticism then we should accept that too. Whether or not these bullets are too big to bite is not something that affects whether an argument is cogent and its conclusion true.

However, here is the second way of expressing the over-generalization worry: if my defense of epistemic moral scepticism also works equally well as a defense of normative or metanormative scepticism, then this might in turn *undermine* epistemic moral scepticism itself. The worry here is that over-generalization leads to self-defeat. If my defense of epistemic moral scepticism works equally well as a defense of normative or metanormative scepticism, and if these latter forms of scepticism entail that there is nothing we ought or should believe, then that *also* entails that we do not know that my defense of epistemic moral scepticism *itself* is true. Therefore, epistemic moral scepticism is self-defeating: it entails a position that says we do not know whether epistemic moral scepticism is true.

To make this clearer, if we accept normative or metanormative scepticism then we cannot say that we have any knowledge — because doing so would mean we are in a position to say that a given epistemic fact obtains. But if we do not have any knowledge of whether an

epistemic fact obtains, then we do not know that the arguments in favour of epistemic moral scepticism are true. In this way, my defense of epistemic moral scepticism works equally well as a defense of a broader form of scepticism - namely, one that entails that we do not know that my defense of epistemic moral scepticism is true. Therefore, we should not accept epistemic moral scepticism, because it is self-defeating.

Unlike the first worry about biting bullets, the second worry about self-defeat seems to me a serious philosophical challenge. It is serious because it suggests that the arguments in the preceding chapters undermine themselves. If epistemic moral scepticism entails normative or metanormative scepticism then, we do not know whether my defense of epistemic moral scepticism is cogent. As a result, "no one would make a rational mistake in rejecting [my claims] and no one would be epistemically praiseworthy in accepting [them]" (Cuneo 2007: 117). The entire dissertation would seem to be an excercise in futility.

Needless to say, I *do* think that one is making a mistake in rejecting my defense of epistemic moral scepticism, and that one *would be* epistemically praiseworthy in accepting it. I think my defense of epistemic moral scepticism *is* cogent, and I think the conclusions I have reached *are* true. So in this final section of the dissertation I will defend against the over-generalization worry. In particular, I will reject that epistemic moral scepticism entails normative scepticism. And on that basis we do not have to worry about epistemic moral scepticism (and my preceding defense of it) being self-defeating.

7.3 Rejecting Normative Scepticism

Here is my case against normative scepticism in brief. Moral properties, as we saw in Chapter 2, invite a particular characterization: they are best understood as *nonnatural* facts. Epistemic properties, on the other hand, invite a different characterization — which I will explain below. For this reason, one of the key desiderata that pushed us towards epistemic moral scepticism drops out of the picture entirely. If they do not invite the same problematic ontology as moral properties, then knowing about epistemic properties is going

to be comparatively simple. Recall that the arguments in the preceding chapters pick on our inability to know *nonnatural* moral properties. But if epistemic properties do not invite the same ontology then the preceding arguments do not obviously generalize.

As we have seen, nonnaturalist cognitivst realism best explains our moral discourse and practice. But what best explains our epistemic discourse and practice? I think it is best understood in terms of *epistemic constitutivism*. While moral properties are nonnatural, epistemic properties are *constitutive* — and these do not invite a nonnatural ontology.

In broad strokes, 'constitutivism' about some phenomenon is the view that there are features related to that phenomenon that commit us to certain norms. In the epistemic case, constitutivism is the view that there are features of belief that guarantee that we are committed to certain epistemic norms (like the fact 'that one ought not to believe against their evidence'). A constitutive argument is one that purports to show that certain epistemic norms get a grip on us simply because we have beliefs.

Epistemic constitutivism is the view that I think best suits our epistemic discourse and practice. When we talk about epistemic properties, we are not talking about nonnatural properties. Rather we are talking about the norms that arise simply in virtue of being agents who are capable of having beliefs at all.

We can perhaps get a better grip on epistemic constitutivism by looking at a parallel case: moral constitutivism. When it comes to morality, constitutivism is the view that owing to the fundamental features of action itself we are thereby committed to certain moral norms. The aspiration here is that "we can find out just which standards are the right ones by deriving them from what is constitutive of action" (Enoch 2006: 32). Moral constitutivists say that we can understand that we are necessarily committed to accepting that (for example) 'torturing people merely for fun is wrong' simply by better understanding

what it means to be the sort of beings who act.[4] Furthermore, moral constitutivists say that moral properties would not invite a nonnatural ontology. The properties are not 'out there' to be discovered, but simply arise because of the sort of agents that we are. This in turn makes it comparatively easy for us to know about them. Since there is no gap to bridge between our judgments and the nonnatural properties, we do not have to wrestle with questions about how we could possibly know such properties.[5]

So according to moral constitutivism, certain features of (e.g) action guarantee that we are committed to certain moral norms obtaining. In parallel, according to epistemic constitutivism, what exactly are we committed to just in virtue of having beliefs? It will be helpful here to distinguish between two things: first, the constitutive features *of* beliefs; and second, the norms that purportedly get a grip on us *in virtue of* beliefs having those constitutive features. The label 'epistemic constitutivism' I think should be reserved for the latter.

Here is one plausible constitutive feature of beliefs: they aim at truth. This is a widely accepted claim, but also a much discussed metaphor.[6] In what sense do beliefs have an *aim*? This of course does not mean that our beliefs *are* true, or even that they are *usually* true. Rather, it means the the following: that beliefs are truth-apt (i.e. capable of being true or false), and that if one recognizes that a belief is *false* then they will abandon that belief. Now, if beliefs in fact have this minimal aim, it would seem to follow that we are already committed to something: namely, that if we believe p then we also believe that p

[4] For prominent defenses of moral constitutivism, see Ferrero (2009), Korsgaard (1996; 2009), Velleman (2000), etc.

[5] Of course, I do not think that moral constitutivism best explains our moral discourse and practice. For a persuasive argument against moral constitutivism, see Enoch (2006; 2011).

[6] For acceptance of the claim, see Setiya (2008), Wedgwood (2002), Williams (1973), etc. For a more wide-ranging discussion of the idea of beliefs aiming at truth, see Chan (2013).

is true. We cannot have a belief that we ourselves take to be false.[7] Indeed, it is part of what allows us to distinguish those mental states which are beliefs from those mental states which are not. A mental state that was transparently false to us could not in fact be a belief — it might instead be something like a hope or a desire. And if this is right, then it is a constitutive feature of beliefs that they aim at truth.

How then do we get from beliefs aiming at truth to something about epistemic norms? The goal of a constitutive argument is to show that certain epistemic norms necessarily get a grip on us owing to the constitutive features of belief. So we need to see whether there are epistemic norms that believers are committed to, precisely because of the constitutive feature of belief just identified (or else some other constitutive feature of belief). Here is how we can see this.

It is typical of our ordinary epistemic discourse and practice that "we presume that [epistemic] facts should weigh with us" (Nolfi, ms). By parallel, we also seem to think that moral facts (like 'torturing people merely for fun is wrong') should weigh with us too. An epistemic (or *meta*epistemic) fact concerns the *status* of a belief: these are facts that tell us what counts in favour or counts against a belief. Again by parallel, a moral fact concerns the status of an *action*: these are facts that tell us what counts in favour or against an action. The constitutivist story across both domains is one that explains how those facts have corresponding norms that get a grip on us. Recall that the nonnaturalist says that such facts are 'out there' to be discovered. This then opens up an explanatory challenge: why should we be gripped by their corresponding norms? Nonnaturalists will of course have to answer that challenge (and I have argued in the previous chapters that they fail to do so). But the constitutivist says that such norms get a grip on us precisely because of the constitutive features of the relevant phenomena (like action, belief, etc.). In the specifically

[7] Indeed, this would be Moore's Paradox: i.e. 'p, but I do not believe that p', or 'p, but I believe not-p'). For a recent discussion, see Green & Williams (2007).

epistemic case, the claim is that just in virtue of having beliefs (which are mental states that necessarily aim at truth) we are going to think that the epistemic norms (which are about the status of our beliefs) are in play. Importantly, this does not depend upon having to be convinced to accept a particular epistemic norms, but rather it appears to follow from our having beliefs *at all*. The task, then, is to explain why we should think that epistemic norms necessarily grip us.

Suppose that I believe that I am a wise person. Moreover, suppose that the reason I believe this is because I read so in a fortune cookie, and on that basis alone I form my belief. We can say that my belief has a characterizable epistemic status: it was formed on the basis of *fortune-cookie-say-so*. Probably you will say that I should not believe things on this basis alone. Probably you will say that *fortune-cookie-say-so* should not contribute to my beliefs *at all*. But note what you would be doing here: you would not necessarily be criticizing the content of my belief, rather you would be criticizing the epistemic *status* of my belief. The more natural way to say this is that you think my belief is *unjustified* — not false, just formed or held on some improper basis.

What is important to recognize here is that your criticism of my belief-formation process presupposes that I *sensitive* in some way to the epistemic status of my beliefs. Or it at least shows that you think that I *should* be sensitive. That is what it means for epistemic facts to 'weigh with us', for epistemic norms to get a 'grip' on us. It means that we are necessarily sensitive to whether we form, maintain, and revise our beliefs in the right sort of way. It is to be sensitive to whether or not our beliefs are justified, among other things.[8] And

[8] Compare this thought to both Sellars (1956) and Brandom (1994), each of whom characterize our normative practices as involving the 'giving and asking for reasons'. In particular, Sellars says that "[i]n characterizing an episode or state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says" (Sellars 1956: sec. 36).

this suggests that we all think that there are epistemic facts that 'weigh with us', epistemic norms that get a grip on us, that inform us about what counts in favour or against beliefs.

One objection here might be that different people clearly have different standards of justification when it comes to belief formation, maintenance, and revision. Some people think that forming beliefs via fortune-cookie-say-so is perfectly legitimate. Other people disagree. Some people think that tabloids are a reliable source of news, that psychics can in fact predicate the future, or that horoscopes accurately explain our day-to-day lives. And those same people form beliefs on those bases. How then could it be that we are all sensitive to epistemic status?

But we can quickly dispense with this objection. Mere differences in *standards* does not undermine epistemic constitutivism. That some people accept *this* standard as opposed to *that* one does not spell trouble for the key constitutivist claim. The case for epistemic constitutivism is plausible if simply in virtue of having beliefs we are all necessarily sensitive to the epistemic status of our beliefs *at all*. The epistemic constitutivist can accept that there is disagreement about which standard one should accept. They will then point out that, irrespective of such disagreement, all of us necessarily accept that there is *some* standard that we ought to accept.

And I think that this is clearly case: we are all sensitive to the epistemic status of our beliefs, even if we disagree over what the correct standard is. If your crazy uncle chooses tabloids as his source of news, this does not show that he does not care about the epistemic status of his beliefs. Rather, it shows that he accepts a different standard than you do. Both of you both accept *some* standard, and it is that which seems to speak in favour of epistemic constitutivism. We are all gripped by epistemic facts about proper belief formation, maintenance, revision, and so on. Moreover, this is much more plausible in the epistemic case than it is in the moral case, as the following comparison can show.

A moral constitutivist would say that just in virtue of being an agent who acts we are necessarily sensitive to moral norms about our conduct. But here is one way we can

reject this claim. While we might want to say that people *should* care about their actions conforming to certain moral norms, it is not *constitutive* of action that people do *in fact* care at all.[9] There are two types of persons we can point to here: *amoralists*, who are not sensitive to the moral status of their actions, and *moral abolitionists*, who even if they are sensitive to the moral status of their actions nevertheless think there are philosophical reasons to work towards disassociating from that sensitivity. The existence of amoralists has been defended by Brink (1986), Svavarsdóttir (1999), and Zangwill (2009), while moral abolitionism as a prescriptive view has been defended by Garner (1994; 2007) and Hinckfuss (1987). Whatever we think about such persons, the mere existence of amoralists and moral abolitionists seems to me sufficient for undermining the claim that caring about moral facts *necessarily* stems from being agents who act. And we can offer a quick (but far from decisive) rejection of moral constitutivism on that basis. After all, amoralists and moral abolitionists still *act*, but they are not sensitive to the moral status of their actions.

On the other hand, there does not seem to be a purely epistemic equivalent of the amoralist or moral abolitionist. Where are those persons who do not care at all about *how* their beliefs are formed? Where are those persons who are actively attempting to *disassociate* their beliefs from any epistemic norms about good belief formation, maintenance, and revision? Theorizing about such persons, if they in fact exist, appears conspicuously absent in the epistemological literature.[10] One thing we might say is that there seem to

[9] This shares something in common with the so-called *shmagency* objection to constitutivism. See Enoch (2006) (2011). But as we will see, my argument here is not against constitutivism *as such*, but rather against normative scepticism. In making this case I rely on a *disanalogy* between the moral and epistemic cases, whereas the schmagency objection cuts acorss both domains.

[10] The closest we get to an epistemic equivalent of an amoralist are what we can call *epistemic deadbeats*. Russell & Doris (2008) describe an agent who doesn't care about their own practical interests, and so (per interest-relative invariantism) is able to easily secure knowledge. We find a similar character in Schaffer (2006), where someone can work towards

be persons who are sensitive to *bad* or *incorrect norms* of belief formation — as said, those who treat tabloids as reliable sources of news, those who seek out advice from psychics, and so on. But this of course wouldn't show *insensitivity* to epistemic facts, so much as a disagreement about how one can get in touch with those facts. Plausibly, the tabloid reader thinks *you're* insensitive to the epistemic facts, which according to them say that tabloids *are* reliable sources of news.

So it seems to me that we can push back against moral constuitutivism by showing that there are people for whom purported moral facts have no weight. These people still act, but mere action alone does not suffice for the moral facts getting a grip on them. But we cannot push back against epistemic constitutivism in the same way. Indeed, if we think about the purely epistemic parallel to amoralism or moral abolitionism, then we can see that epistemic constitutivism is in fact a highly plausible view. If there are no such persons who do not feel the weight of epistemic facts, then it suggests that epistemic facts weigh with us all. And the reason they seem to weigh with us all is simply in virtue of our having beliefs in the first place. For this comparative argument, note also that it is irrelevant whether amoralists or moral abolitionists are actually common. It is only relevant whether there are *any* such persons, or whether they are psychologically possible. Because the moral constitutivist says that by the very fact that we are acting or reasoning agents we are necessarily sensititive to moral facts, any exception seems to disprove that claim. Of course, the strict standard applies to epistemic constitutivism. But for the latter case, there do not seem to be any such persons.

---

lowering their practical stakes to attain knowledge. But neither of these characters is insensitive to the epistemic status of their beliefs. Rather, it is precisely *because* they are sensitive to the epistemic status of their beliefs that they have an *interest* in being *insensitive* to their own practical interests.

While this is far from a rich defense of epistemic constitutivism, it is enough to motivate my suggestion that the moral and epistemic domains appear to invite much different ontologies. What I think this shows is that the epistemic domain does not invite a nonnatural ontology. We do not see the same dialectic across both morality and epistemology. In the moral case, we begin with particular ambitions set by our moral discourse and practice — that our moral judgments are beliefs about nonnatural properties. But in the epistemic case, the ambitions are different — our epistemic judgments do not seem to be about nonnatural properties.

And if epistemic facts are not best understood as nonnatural facts, then the arguments from the preceeding chapters are not going to generalize. Because they do not invite the same problematic ontology as moral properties, knowing about epistemic properties is comparatively simple. Recall that the arguments in the preceding chapters pick on our inability to know *nonnatural* moral properties. But if epistemic properties do not invite the same ontology then the preceding arguments do not obviously generalize.

I think the sketch here is sufficient for explaining why I do not think we should be bothered by the over-generalization worry. The arguments I have given for why we lack moral knowledge do not necessarily compel us to accept normative scepticism. And because they do not over-generalize, we do not have to worry about moral scepticism being self-defeating. For that reason, I think that my defense of epistemic moral scepticism can be limited to the moral domain.

7.4 Wrapping Up

In the introduction to the dissertation I noted two characters: Amy, who was an atheist, and Gary, who was a grass-counter. These two characters allowed us to understand the difference between the ontological and epistemic routes to scepticism. I have not argued that moral facts do not exist — so I have not pursued the ontological route to scepticism. Rather, I have argued that *even if* there are moral facts we do not have a plausible account

of how we know them — this is the epistemological route to scepticism. Because we have no plausible moral epistemology for nonnaturalist cognitivist realism I think we should accept epistemic moral scepticism (and, per the above, *only* moral scepticism).

But where should we go from here? I think a next step could be to think more about our moral discourse and practice itself. Supposing that you accept the sceptical thrust of the dissertation, you might think that what I have ultimately defended is not epistemic moral scepticism but actually moral nihilism. For, if we cannot know anything about the moral facts, then we might not have any reason to continue thinking that they exist. In a similar way, if we continually come up empty-handed in our attempts to know anything about God, then perhaps we should stop thinking that God exists. While I had set aside questions of ontology, the next step might be to see if my epistemic conclusion extends to or helps support an ontological one. Ultimately, epistemic moral scepticism might just be a holding station on the way to moral nihilism.

REFERENCES

[1] Anscombe, G.E.M. 1957. *Intention*. Harvard University Press.

[2] Armstrong, David. 1978. *Universals and Scientific Realism*. Cambridge University Press.

[3] Armstrong, David. 1999. "A Naturalist Program: Epistemology and Ontology". *Proceedings and Addresses of the American Philosophical Association*. 73 (2): 77 — 89.

[4] Au, T.K. (1994). "Developing an intuitive understanding of substance kinds". *Cognitive Psychology*. 27: 71 — 111.

[5] Audi, Robert. 1999. "Self-Evidence. *Philosophical Perspectives*. 13: 205 — 228.

[6] Audi, Robert. 2001. "A Kantian Intuitionism. *Mind*. 110 (439): 601 — 635.

[7] Audi, Robert. 2004. *The Good in the Right: A Theory of Intuition and Intrinsic Value*. Princeton University Press.

[8] Audi, Robert. 2008. "Intuition, Inference, and Rational Disagreement in Ethics. *Ethical Theory and Moral Practice*. 11 (5): 475 — 492.

[9] Austin, J.L. 1962. *How to Do Things with Words*. Carendon Press.

[10] Bagnoli, Carla. 2013. *Constructivism in Ethics*. Cambridge University Press.

[11] Bagnoli, Carla. 2017. "Constructivism in Metaethics. *The Stanford Encyclopedia of Philosophy*. Editor Edward N. Zalta.

[12] Band, Megan, Douglas L. Medin & Scott Atran (2007) "Cultural mosaics and mental models of nature". *PNAS*. 104(35): 13868 — 13874

[13] Bealer, George. 1998. "Intuition and the Autonomy of Philosophy. *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. Editors Michael DePaul & William Ramsey. Rowman & Littefield. 201 — 240.

[14] Bedke, Matthew. 2009. "Intuitive non-naturalism meets cosmic coincidence. *Pacific Philosophical Quarterly*. 90 (2): 188 — 209.

[15] Bedke, Matthew. 2014. "No Coincidence?. *Oxford Studies in Metaethics*. 9: 102 — 125.

[16] Benacerraf, Paul. 1973. "Mathematical Truth. *Journal of Philosophy*. 70 (19): 661 679.

[17] Bengson, John. 2015. "The Intellectual Given. *Mind*. 124 (495): 707 — 760.

[18] Berker, Selim, 2014. "Does Evolutionary Psychology Show That Normativity Is Mind-Dependent?. *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*. Editors Justin DArms & Daniel Jacobson. Oxford University Press. 215 — 252.

[19] Berker, Selim. 2011. "Guptas Gambit. *Philosophical Studies*. 152 (1): 17 — 39.

[20] Berlin, B., Breedlove, D.E. and Raven, P.H. (1973). "General Principles of Classification and Nomenclature in Folk Biology". *American Anthropologist* 75. 214 — 242.

[21] Bertamini M, Spooner A, & Hecht H. (2004)."The representation of naive knowledge about physics". *Multidisciplinary Approaches to Visual Representations and Interpretations*. Editor G. Malcolm. Elsevier.

[22] Blackburn, Simon. 1984. *Spreading the Word: Groundings in the Philosophy of Language*. Clarendon Press.

[23] Blackburn, Simon. 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford University Press.

[24] Bloom, Paul. 2013. *Just Babies: The Origins of Good and Evil*. Crown Publishers.

[25] Bogardus, Tomas. 2016. "Only All Naturalists Should Worry About Only One Evolutionary Debunking Argument. *Ethics*. 126 (3): 636 — 661.

[26] Bonjour, Laurence. 1985. *The Structure of Empirical Knowledge*. Harvard University Press.

[27] Bourget, David & David J. Chalmers. 2014. "What Do Philosophers Believe? *Philosophical Studies*. 170 (3): 465 — 500.

[28] Braddon-Mitchell, David. 2004. "Folk Theories of the Third Kind. *Ratio*. 17 (3): 277 — 293.

[29] Brandom, Robert. 1994. *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.

[30] Brandom, Robert. 2000. *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press.

[31] Brink, David. 1986. "Externalist Moral Realism. *Southern Journal of Philosophy*. 24 (S1): 23 — 41.

[32] Broome, John. 2007. "Is Rationality Normative? *Disputatio*. 2 (23): 161 — 178.

[33] Brosnan, Kevin. 2011. "Do the Evolutionary Origins of Our Moral Beliefs Undermine Moral Knowledge?. *Biology and Philosophy*. 26 (1): 51 — 64.

[34] Burgess, John P. & Gideon Rosen. 1997. *A Subject with No Object: Strategies for Nominalistic Interpretation of Mathematics*. Oxford University Press.

[35] Burgess, John. 2007. "Against Ethics. *Ethical Theory and Moral Practice*. 10 (5): 427 — 439.

[36] Camp, Elisabeth. 2017. "Metaethical Expressivism. *The Routledge Handbook of Metaethics*. Editors Tristram McPherson & David Plunkett. Routledge. 87 — 101.

[37] Carey, Susan and Mary Spelke (1996). "Science and Core Knowledge". *Philosophy of Science*. 63(4): 515 — 533.

[38] Carey, Susan (1985). *Conceptual Change in Childhood*. MIT Press.

[39] Carey, Susan. 1988. "Conceptual Differences Between Children and Adults. *Mind and Language*. 3 (3): 167 — 181.

[40] Carnap, Rudolf. 1937. *Philosophy and Logical Syntax*. K. Paul, Trench, Trubner & Co.

[41] Cath, Yuri. 2016. "Reflective Equilibrium. *The Oxford Handbook of Philosophical Methodology*. Editors H. Cappelen, T. Gendler & J. Hawthorne. Oxford University Press. 213 — 230.

[42] Chan, Timothy. 2013. *The Aim of Belief*. Oxford University Press.

[43] Chrisman, Matthew. 2008. "Ought to Believe. *Journal of Philosophy*. 105 (7): 346 — 370.

[44] Chrisman, Matthew. 2010. "From Epistemic Expressivism to Epistemic Inferentialism. *Social Epistemology*. Editors Adrian Haddock, Alan Millar & Duncan Pritchard. Oxford University Press.

[45] Chrisman, Matthew. 2016. *The Meaning of 'Ought: Beyond Descriptivism and Expressivism in Metaethics*. Oxford University Press.

[46] Chudnoff, Elijah. 2011. "What Intuitions Are Like". *Philosophy and Phenomenological Research*. 82 (3): 625 — 654.

[47] Churchland, Patricia. 2011. *Braintrust: What Neuroscience Tells Us About Morality*. Princeton University Press.

[48] Clarke-Doane, Justin. 2012. "Moral and Mathematics: The Evolutionary Challenge. *Ethics*. 122 (2): 313 — 340.

[49] Clarke-Doane, Justin. 2015. "Justification and Explanation in Mathematics and Morality". *Oxford Studies in Metaethics*. 10. 80 — 103.

[50] Clarke-Doane, Justin. 2017. "What is the Benacerraf Problem?" *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity*. Editor Fabrice Pataut. 17 — 43.

[51] Colyvan, Mark. 2001. *The Indispensability of Mathematics*. Oxford University Press.

[52] Copp, David. 2008. "Darwinian Skepticism About Moral Realism". *Philosophical Issues*. 18 (1): 186 — 206.

[53] Crary, Alice. 2007. *Beyond Moral Judgment*. Harvard University Press.

[54] Crisp, Roger. 2006. *Reasons and the Good*. Clarendon Press.

[55] Cuneo, Terence. 2002. "Reconciling Realism with Humeanism". *Australasian Journal of Philosophy*. 80 (4): 465 — 486.

[56] Cuneo, Terence. 2006. "Saying What We Mean: An Argument Against Expressivism". *Oxford Studies in Metaethics*. 1. 35 — 71.

[57] Cuneo, Terence. 2007. *The Normative Web: An Argument for Moral Realism*. Oxford University Press.

[58] Das, Ramon. 2016. "Evolutionary Debunking of Morality: Epistemological or Metaphysical?. *Philosophical Studies*. 173 (2): 417 — 435.

[59] De Cruz, Helen. 2007. "The Role of Intuitive Ontologies in Scientific Understanding the Case of Human Evolution". *Biology and Philosophy*. 22 (3): 351 — 368.

[60] de Maagt, Sem. 2017. "Reflective Equilibrium and Moral Objectivity". *Inquiry: An Interdisciplinary Journal of Philosophy*. 60 (5): 443 — 465.

[61] de Waal, Frans. 1996. *Good Natured: The Origins of Right and Wrong in Humans and Other Animals*. Harvard University Press.

[62] de Waal, Frans. 2009. *Primate and Philosophers: How Morality Evolved*. Princeton University Press.

[63] DePaul, Michael. 2009. "Pyrrhonian Moral Skepticism and the Problem of the Criterion". *Philosophical Issues*. 19 (1): 38 — 56.

[64] DeRose, Keith. 1999. *Skepticism: A Contemporary Reader*. Oxford University Press.

[65] Descartes, Rene. 1641/1993. *Meditations on First Philosophy*. Translator Donald A. Cress. Hackett Publishers.

[66] Dworkin, Ronald. 1996. "Objectivity and truth: Youd Better Believe It". *Philosophy and Public Affairs*. 25 (2): 87 — 139.

[67] Ellis, B.D. 1990. *Truth and Objectivity*. Blackwell.

[68] Enoch, David. 2007. "An Outline of an Argument for Robust Metanormative Realism". *Oxford Studies in Metaethics*. 2. 21 — 50.

[69] Enoch, David. 2009. "Can There Be a Global, Interesting, Coherent Constructivism About Practical Reason?". *Philosophical Explorations*. 12 (3): 319 — 339.

[70] Enoch, David. 2010. "The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It. *Philosophical Studies*. 148 (3): 413 — 438.

[71] Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.

[72] Enoch, David. 2017. "Non-Naturalistic Realism in Metaethics". *The Routledge Handbook of Metaethics*. Editors Tristram McPherson & David Plunkett. Routledge. 29 — 42.

[73] Fantl, Jeremy. 2013. "A Defense of Dogmatism". *Oxford Studies in Epistemology*. 4. 34 — 57.

[74] Ferrero, Luca. 2009. "Constitutivism and the Inescapability of Agency". *Oxford Studies in Metaethics*. 4. 303 — 303.

[75] Field, Hartry. 1989. *Realism, Mathematics & Modality*. Blackwell.

[76] Fine, Kit. 2002. *The Limits of Abstraction*. Oxford University Press.

[77] FitzPatrick, William. 2008. "Robust Ethical Realism, Non-Naturalism, and Normativity". *Oxford Studies in Metaethics*. 3. 159 — 205.

[78] FitzPatrick, William. 2015. "Debunking Evolutionary Debunking of Ethical Realism". *Philosophical Studies*. 172 (4): 883 — 904.

[79] Fraser, Benjamin. 2014. "Evolutionary Debunking Arguments and the Reliability of Moral Cognition. *Philosophical Studies*. 168 (2): 457 — 473.

[80] Garner, Richard. 1994. *Beyond Morality*. Temple University Press.

[81] Garner, Richard. 2007. "Abolishing Morality". *Ethical Theory and Moral Practice*. 10 (5): 499 — 513.

[82] Gazzaniga, Michael S. 2005. "Facts, Fictions and the Future of Neuroethics". *Neuroethics: Defining the Issues in Theory, Practice, and Policy*. Editor Judy Illes. Oxford University Press. 141 — 148.

[83] Gibbard, Allan. 1990. *Wise Choice, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press.

[84] Gibbard, Allan. 2003. *Thinking How to Live*. Harvard University Press.

[85] Gibbard, Allan. 2012. "Five Girls on a Rock." *Review of Books*. 34 (11): 22-24.

[86] Godel, Kurt. 1947. "What is Cantors Continuum Problem?". *Journal of Symbolic Logic*. Editors Solomon Feferman, John Dawson & Stephen Kleene. Oxford University Press. 176 — 187.

[87] Goldman, Alvin. 2007. "Philosophical intuitions: Their target, their source, and their epistemic status. *Grazer Philosophische Studien*. 74 (1): 1 — 26.

[88] Goldman, Alvin. 2009. "Internalism, Externalism, and the Architecture of Justification. *Journal of Philosophy*. 106 (6). 309 — 338.

[89] Goodman, Nelson. 1955. *Fact, Fiction, and Forecast*. Harvard University Press.

[90] Goodwin, Geoffrey P. & John M. Darley. 2008. "The psychology of meta-ethics: Exploring objectivism. *Cognition*. 106 (3): 1339 — 1366.

[91] Green, Mitchell S. & John N. Williams. 2007. *Moores Paradox: New Essays on Belief, Rationality, and the First Person*. Oxford University Press.

[92] Greene, Joshua. 2013. *Moral Tribes: Emotion, Reason, and the Gap Between Us and Them*. Penguin Press.

[93] Gupta, Anil. 2006. *Empiricism and Experience*. Oxford University Press.

[94] Hale, Bob. 2001. *The Reasons Proper Study: Essays Towards a Neo-Fregean Philosophy of Mathematics*. Oxford University Press.

[95] Hampton, Jean. 1998. *The Authority of Reason*. Cambridge University Press.

[96] Hanson, Louise. 2017. "The Real Problem with Evolutionary Debunking Arguments". *Philosophical Quarterly*. 67 (268): 508 — 533.

[97] Hare, R.M. 1952. *The Language of Morals*. Oxford Carendon Press.

[98] Hauser, Marc. 2006. *Moral Minds: How Nature Designed Our Universal Sense of Right and Wrong*. Harper Collins.

[99] Heathwood, Chris. 2018. "Epistemic Reductionism and the Moral-Epistemic Disparity. *Metaepistemology: Realism & Antirealism*. Editors Chrisos Kyriacou & Robin McKenna. Palgrave Macmillan. 45 — 70.

[100] Heney, Diane. 2016. *Toward a Pragmatist Metaethics*. Routledge.

[101] Hinckfuss, Ian. 1987. *The Moral Society: Its Structure and Effects*. Australian National University.

[102] Huemer, Michael. 2005. *Ethical Intuitionism*. Palgrave Macmillan.

[103] Huemer, Michael. 2017. "There Is No Pure Empirical Reasoning". *Philosophy and Phenomenological Research*. 95 (3): 592 — 613.

[104] Isserow, Jessica. forthcoming. "Evolutionary Hypotheses and Moral Skepticism". *Erkenntnis*.

[105] Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press.

[106] Jackson, Frank. 2000. "Psychological explanation and implicit theory". *Philosophical Explorations*. 3 (1); 83 — 95.

[107] Joyce, Richard & Simon Kirchin. 2010. *A World Without Values*. Springer.

[108] Joyce, Richard. 2001. *The Myth of Morality*. Cambridge University Press.

[109] Joyce, Richard. 2005. *The Evolution of Morality*. MIT Press.

[110] Joyce, Richard. 2006. "Metaethics and the Empirical Sciences". *Philosophical Explorations*. 9 (1): 133 — 148.

[111] Kahane, Guy. 2011. "Evolutionary Debunking Arguments". *Nous*. 45 (1): 103 — 125.

[112] Keil, Peter. 2010. "The Feasibility of Folk Science". *Cognitive Science*. 34 (5): 826 — 862.

[113] Keller, Simon. 2017. "Love and the Moral Error Theory: Is Love a Mistake?". *Philosophy and Phenomenological Research*. 95 (3): 709 — 721.

[114] Kelly, Thomas. 2005. "Moorean Facts and Belief Revision, or Can the Skeptic Win?". *Philosophical Perspectives*. 19 (1): 179 — 209.

[115] Kirchin, Simon. 2008. "Review of Alice Crary, Beyond Moral Judgment". *Notre Dame Philosophical Reviews*. (1).

[116] Kirchin, Simon. 2015. "Self-evidence, Theory and Anti-theory". *Intuition, Theory, and Anti-Theory in Ethics*. Editor Sophie Grace Chappell. Oxford University Press. 167 — 185.

[117] Kirchin, Simon. 2017. *Thick Evaluation*. Oxford University Press.

[118] Kitcher, Phillip. 2011. *The Ethical Project*. Harvard University Press.

[119] Korcz, Keith Allen. 2000. "The causal-doxastic theory of the basing relation". *Canadian Journal of Philosophy*. 30 (4): 525 — 550.

[120] Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge University Press.

[121] Korsgaard, Christine. 2003. "Realism and Constructivism in Twentieth-Century Moral Philosophy". *Journal of Philosophical Research*. 28 (Supplement): 99 — 122.

[122] Korsgaard, Christine. 2008. *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford University Press.

[123] Korsgaard, Christine. 2009. *Self-Constitution: Agency, Identify, and Integrity*. Oxford University Press.

[124] Kvanvig, Jonathan. 2003. "Propositionalism and the Perspectival Character of Justication". *American Philosophical Quarterly*. 40 (1): 3 — 18.

[125] Leary, Stephanie. 2017. "Non-Naturalism and Supervenience". *Oxford Studies in Metaethics*. 12. 76 — 105.

[126] Lehrer, Keith. 1971. "Why Not Scepticism?". *Philosophical Forum*. 2 (3): 283.

[127] Leibowitz, Uri D. & Neil Sinclair. 2016. *Explanation in Ethics and Mathematics: Debunking and Dispensability*. Oxford University Press.

[128] Lekan, Todd. 2003. *Making Morality: Pragmatist Reconstruction in Ethical Theory*. Vanderbilt University Press.

[129] Lekan, Todd. 2006. "Pragmatist Metaethics: Moral Theory as a Deliberative Practice". *Southern Journal of Philosophy*. 44 (2): 253 — 271.

[130] Lewis, David. 1986. *On the Plurality of Worlds*. Wiley-Blackwell.

[131] Lewis, David. 1999. "Elusive Knowledge". *Paper in Metaphysics and Epistemology*. Cambridge University Press. 418 — 445.

[132] Lucas, J.R. 1971. "Ethical Intuitionism II". *Philosophy*. 46 (175): 1 — 11.

[133] Lutz, Matt. forthcoming. "What Makes Evolution a Defeater?". *Erkenntnis*. 1 — 22.

[134] Lycan, William. 2001. "Moore Against the New Skeptics". *Philosophical Studies*. 103 (1): 35 — 53.

[135] Machuca, Diego E. 2018. *Moral Skepticism: New Essays*. Routledge.

[136] Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. Penguin Books.

[137] Mackie, J.L. 1980. *Humes Moral Theory*. Routledge and Kegan Paul.

[138] McGrath, Sarah. 2004. "Moral Knowledge by Perception". *Philosophical Perspectives*. 18 (1): 209 — 228.

[139] McPherson, Tristram & David Plunkett. 2017. *The Routledge Handbook of Metaethics*. Routledge

[140] McPherson, Tristram. 2009. "Moorean Arguments and Moral Revisionism". *Journal of Ethics and Social Philosophy*. 2. 1 — 24.

[141] McPherson, Tristram. 2011. "Against Quietist Normative Realism". *Philosophical Studies*. 154 (2): 223 — 240.

[142] Mogensen, Andreas. 2015. "Evolutionary debunking arguments and the proximate/ultimate distinction". *Analysis*. 75 (2): 196 — 203.

[143] Mogensen, Andreas. 2016. "Do evolutionary debunking arguments rest on a mistake about evolutionary explanations?". *Philosophical Studies*. 173 (7): 1799 — 1817.

[144] Moon, Andrew. 2017. "Debunking Morality: Lessons from the EAAN Literature". *Pacific Philosophical Quarterly*. 98 (S1): 208 — 226.

[145] Moore, G.E. 1903. *Principia Ethica*. Dover Publications.

[146] Moore, G.E. 1912. *Ethics*. Oxford University Press.

[147] Moore, G.E. 1959. *Philosophical Papers*. Macmillan.

[148] Nagel, Jennifer. 2007. "Ethical Intuitions". *Philosophy Compass*. 2 (6): 792 — 819.

[149] Nagel, Thomas. 1997. *The Last Word*. Oxford University Press.

[150] Nagel, Thomas. 2012. *Mind and Cosmos*. Oxford University Press.

[151] ONeill, Onora. 1989. *Constructions of Reason: Explorations of Kants Practical Philosophy*. Cambridge University Press.

[152] Olson, Jonas. 2009. "Reasons and the New Non-Naturalism". *Spheres of Reason*. Editor Simon Robertson. Oxford University Press.

[153] Olson, Jonas. 2011. "In defense of moral error theory". *New Waves in Metaethics*. Editor Michael Brady. Palgrave-Macmillan.

[154] Olson, Jonas. 2014. *Moral Error Theory: History, Critique*. Oxford University Press.

[155] Parfit, Derek. 2006. "Normativity". *Oxford Studies in Metaethics*. 1. 325 — 380.

[156] Parfit, Derek. 2011. *On What Matters: Volumes 1 & 2*. Oxford University Press.

[157] Parfit, Derek. 2017. *On What Matters: Volume 3*. Oxford University Press.

[158] Plantinga, Alvin. 2000. *Warranted Christian Belief*. Oxford University Press.

[159] Plantinga, Alvin. 2011. *Where the Conflict Really Lies: Science, Religion, & Naturalism*. Oxford University Press.

[160] Pollock, John L. 1986. *Contemporary Theories of Knowledge*. Hutchinson.

[161] Price, Huw. 2011. *Naturalism Without Mirrors*. Oxford University Press.

[162] Prichard, H.A. 1912. "Does moral philosophy rest on a mistake?". *Mind*. 21 (81): 21 — 37.

[163] Pryor, Jim. 2000. "The Skeptic and the Dogmatist". *Nous*. 34 (4): 517 — 549.

[164] Pryor, Jim. 2004. "Whats wrong with Moores argument?". *Philosophical Issues*. 14 (1): 349 — 378.

[165] Quine, Willard V. 1948. "On What There Is". *Review of Metaphysics*. 2 (1): 21 — 38.

[166] Railton, Peter. 2014. "The affective dog and its rational tale: intuition and attunement". *Ethics*. 124 (4): 813 — 859.

[167] Rawls, John. 1971. *A Theory of Justice*. Harvard University Press.

[168] Ross, W.D. 1930. *The Right and the Good*. Clarendon Press.

[169] Ross, W.D. 1939. *Foundations of Ethics*. Oxford University Press.

[170] Ruse, Michael & Edward O. Wilson. 1986. "Moral Philosophy as Applied Science". *Philosophy*. 61 (236). 173 — 192.

[171] Russell, Gillian & John Doris. 2008. "Knowledge by indifference". *Australasian Journal of Philosophy*. 86 (3): 429 — 437.

[172] Scanlon, T.M. 2014. *Being Realistic About Reasons*. Oxford University Press.

[173] Schafer, Karl. 2010. "Evolution and Normative Scepticism". *Australasian Journal of Philosophy*. 88 (3): 471 — 488.

[174] Schaffer, Jonathan. 2006. "The irrelevance of the subject: Against subject-sensitive invariantism". *Philosophical Studies*. 127 (1): 87 — 107.

[175] Schechter, Joshua. 2018. "Explanatory Challenges in Metaethics". *Routledge Handbook of Metaethics*. Editors Tristram McPherson & David Plunkett. Routledge. 443 — 459.

[176] Searle, John. 1984. "Intentionality and its Place in Nature". *Synthese*. 38 (October): 87 — 100.

[177] Searle, John. 1975. "The Logical Status of Fictional Discourse". *New Literary History*. 6 (2): 319 — 332.

[178] Sellars, Wilfred. 1956. "Empiricism and the Philosophy of Mind". *Minnesota Studies in the Philosophy of Science*. 1. 253 — 329.

[179] Sepielli, Andrew. 2015. "Moral Realism Without Moral Metaphysics". *Oxford Studies in Metaethics*. 11. Oxford University Press.

[180] Sepielli, Andrew. 2018. "Metaethics and Pragmatism". *Routledge Handbook of Metaethics*. Editors Tristram McPherson & David Plunkett. Routledge.

[181] Setiya, Kieran. 2008. "Believing at Will". *Midwest Studies in Philosophy*. 32 (1): 36 — 52.

[182] Shafer-Landau, Russ. 2001. "Knowing right from wrong". *Australasian Journal of Philosophy*. 79 (1): 62 — 80.

[183] Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford University Press.

[184] Shafer-Landau, Russ. 2012. "Evolutionary Debunking, Moral Realism, and Moral Knowledge". *Journal of Ethics and Social Philosophy*. 7 (1).

[185] Sidgwick, Henry. 1874. *The Methods of Ethics*. Thoemmes Press.

[186] Sinclair, Neil. 2018. "Belief pills and the possibility of moral epistemology". *Oxford Studies in Metaethics*. 13.

[187] Sinnott-Armstrong, Walter. 2006. *Moral Skepticisms*. Oxford University Press.

[188] Skarsaune, Knut Olav. 2011. "Darwin and moral realism: Survival of the iffiest". *Philosophical Studies*. 152 (2): 229 — 243.

[189] Smith, Michael. 1987. "The Humean theory of motivation". *Mind*. 96 (381): 36 — 61.

[190] Sober, Elliott. 1984. "Common cause explanation". *Philosophy of Science*. 51 (2): 212 — 241.

[191] Sosa, Ernest. 2007. "Experimental Philosophy and Philosophical Intuition". *Philosophical Studies*. 132 (1): 99 — 107.

[192] Sterelny, Kim & Ben Fraser. 2016. "Evolution and Moral Realism". *British Journal for the Philosophy of Science*. 68 (4): 981 — 1006.

[193] Stratton-Lake, Philip. 2002. *Ethical Intuitionism: Re-Evaluations*. Oxford University Press.

[194] Strawson, P.F. 1949. "Ethical Intuitionism". *Philosophy*. 24 (88): 23 — 33.

[195] Street, Sharon. 2006. "A Darwinian dilemma for realist theories of value". *Philosophical Studies*. 127 (1): 109 — 166.

[196] Street, Sharon. 2008. "Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About". *Philosophical Issues*. 18 (1): 207 — 228.

[197] Streumer, Bart. 2017. *Unbelievable Errors: An Error Theory About All Normative Judgments*. Oxford University Press.

[198] Stroud, Barry. 1979. "Inference, belief, and understanding". *Mind*. 88 (350): 179 — 196.

[199] Svavarsdottir, Sigrun. 2001. "Objective Values: Does Metaethics Rest on a Mistake?". *Objectivity in Law and Morals*. Editor Brian Leiter. Cambridge University Press. 144 — 193.

[200] Svavarsdottir, Sigrun. 1999. "Moral Cognitivism and Motivation". *Philosophical Review*. 108 (2): 161 — 219.

[201] Tersman, Folke. 2017. "Debunking and Disagreement". *Nous*. 51 (4): 754 — 774.

[202] Tersman, Folke. 2018. "Recent Work on Reflective Equilibrium and Method in Ethics". *Philosophy Compass*. 13 (6):

[203] Thomson, Judith Jarvis. 1996. *Moral Relativism and Moral Objectivity*. Blackwell.

[204] Tiefensee. 2016. "Inferentialist metaethics, bifurcations and ontological commitment". *Philosophical Studies*. 173 (9): 2437 — 2459.

[205] Tomasello, Michael. 2016. *A Natural History of Human Morality*. Harvard University Press.

[206] Toner, Christopher. 2011. "Evolution, naturalism, and the worthwhile: A critique of Richard Joyces evolutionary debunking of morality". *Metaphilosophy*. 42 (4): 520 — 546.

[207] Turri, John. 2010. "On the relationship between propositional and doxastic justiciation". *Philosophy and Phenomenological Research*. 80 (2): 312 — 326.

[208] Unger, Peter. 1974. "Two types of scepticism". *Philosophical Studies*. 25 (2): 77 96.

[209] Van Roojen, Mark. 2014. "Moral Cognitivism vs. Non-Cognitivism". *The Stanford Encyclopedia of Philosophy*. Editor Edward N. Zalta.

[210] Vavova, Katia. 2014. "Debunking Evolutionary Debunking". *Oxford Studies in Metaethics*. 9. 76 — 101.

[211] Vavova, Katia. 2015. "Evolutionary Debunking of Moral Realism". *Philosophy Compass*. 10 (2): 104 — 116.

[212] Vavova, Katia. 2018. "Irrelevant Influences". *Philosophy and Phenomenological Research*. 96 (1): 134 — 152.

[213] Velleman, David. 2000. *The Possibility of Practical Reason*. Oxford University Press.

[214] Vogt, Katja. 2016. "Ancient Skepticism". *Cambridge History of Moral Philosophy*. Editors Sacha Golob and Jes Timmermann. Cambridge University Press. 88 — 99.

[215] Vogt, Katja. 2018. "Ancient Skepticism". *The Stanford Encyclopedia of Philosophy*. Editor Edward N. Zalta.

[216] Wedgwood, Ralph. 2002. "The Aim of Belief". *Philosophical Perspectives*. 36 (S16): 267 — 297.

[217] Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford University Press.

[218] White, Roger. 2006. "Problems for Dogmatism". *Philosophical Studies*. 131 (3): 525 — 557.

[219] Wielenberg, Erik J. 2010. "On the evolutionary debunking of morality". *Ethics*. 120 (3): 441 — 464.

[220] Wielenberg, Erik J. 2014. *Robust Ethics: The Metaphysics and Epistemology of Godless Normative Realism*. Oxford University Press.

[221] Williams, Bernard. 1973. *Problems of the Self*. Cambridge University Press.

[222] Williamson, Timothy. 2004. "Philosophical intuitions and scepticism about judgement". *Dialectica*. 58 (1): 109 — 153.

[223] Wittgenstein, Ludwig. 1958. *The Blue and Brown Books*. Translator G.E.M. Anscombe Harper & Row.

[224] Wittgenstein, Ludwig. 1969. *On Certainty*. Translator G.E.M. Anscombe. Harper Torchbooks.

[225] Wodak, Daniel. 2017. "Why realists must reject normative quietism". *Philosophical Studies*. 174 (11): 2785 — 2817.

[226] Wright, Crispin. 1983. *Freges Conception of Numbers as Objects*. Aberdeen University Press.

[227] Zangwill, Nick. 2009. "Non-Cognitivism and Motivation". *New Essays on the Explanation of Action*. Editor Constantine Sandis. Palgrave-Macmillan. 416 — 424.

[228] Zimmerman, Aaron. 2010. *Moral Epistemology*. Routledge.