

Unbelievable Doubts (and other skeptical discoveries)

by

Jonathan Faerber
B.A., University of Alberta, 2013

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of
MASTER OF ARTS

in the Department of Philosophy

© Jonathan Faerber, 2017
University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part, by photocopy or other means, without the permission of the author.

Supervisory Committee

Unbelievable Doubts (and other skeptical discoveries)

by

Jonathan Faerber
B.A., University of Alberta, 2013

Supervisory Committee

Dr. Scott Woodcock, Department of Philosophy
Supervisor

Dr. Colin Macleod, Department of Philosophy
Departmental Member

Abstract

Supervisory Committee

Dr. Scott Woodcock, Department of Philosophy
Supervisor

Dr. Colin Macleod, Department of Philosophy
Departmental Member

Moral skeptics sometimes argue that science is at odds with morality. These arguments sometimes privilege scientific explanations of moral belief at the expense of objective moral knowledge. More specifically, since morality is (arguably) a biological adaptation involving our species' cognitive states, Richard Joyce and Sharon Street doubt the justification and objective truth of moral belief, respectively. This thesis defends objective normative facts from this empirical problem. Reasons for moral skepticism are not compatible with arguments against objective normativity. Put simply, without objective normativity, skeptics have no ultimate reason to doubt anything in particular, moral or otherwise. So, on pain of incoherence, moral skeptics should doubt the truth, rather than the objective normativity, of moral belief.

Table of Contents

SUPERVISORY COMMITTEE.....	II
ABSTRACT	III
TABLE OF CONTENTS.....	IV
LIST OF TABLES.....	V
ACKNOWLEDGEMENTS.....	VI
PREFACE	VII
1: KNOWLEDGE AND VARIETIES OF DISBELIEF	1
1.0: INTRODUCTION	1
1.1: MAKING SENSE OF MORALITY.....	4
1.2: OBJECTIVE MORAL FACTS AND THE EMPIRICAL PROBLEM.....	8
1.3: MORAL IRREALISM AND THE EMPIRICAL PROBLEM	14
1.4: MAKING SENSE OF MORAL SKEPTICISM(S)	20
2: OBJECTIVE MORAL FACTS AND ERROR THEORETIC VALUE.....	25
2.0: INTRODUCTION	25
2.1 OBJECTIVE MORALITY AND ERRORS OF BELIEF	26
2.2: WORLDLY FACTS AND FICTIONAL WORLDS	31
2.3 DOXASTIC ERROR AND EPISTEMIC DEFEAT	36
2.4: OBJECTIVITY AND NORMATIVITY	38
3: SUBJECTIVE MORAL FACTS AND VALUE IRREALISM.....	45
3.0: INTRODUCTION	45
3.1: SUBJECTIVE MORALITY AND TRUE BELIEF	46
3.2: IRREALISM AND MORAL KNOWLEDGE.....	49
3.3: IN DEFENSE OF JOYCE	59
3.4: MORAL IRREALISM AND NORMATIVE TRUTH.....	62
4: SKEPTICISM AND OBJECTIVE EPISTEMIC FACTS	64
4.0: INTRODUCTION	64
4.1: EPISTEMIC IRREALISM AND NORMATIVE TRUTH.....	65
4.2: EPISTEMIC FACTS AND THE META-NORMATIVE EMPIRICAL PROBLEM.....	68
4.3: OBJECTIVE EPISTEMIC FACTS.....	75
4.4: META-NORMATIVE EXPLANATIONS AND MORAL FACTS.....	85
5: SKEPTICISM AND OBJECTIVE NORMATIVE FACTS	90
5.0: INTRODUCTION	90
5.1: NORMATIVE EXPLANATIONS AND OBJECTIVE MORAL FACTS.....	91
5.2: CONCLUSION	100
BIBLIOGRAPHY	103

List of Tables

TABLE 1: Comparison of Meta-Ethical Theories.....7

TABLE 2: Comparison of Meta-Epistemic Theories 65

Acknowledgements

Upon completing this thesis, I recall the naïve moral of my project's primitive ancestor, some years ago, in my initiative MA pro-seminar. I remember thinking then that although we often fail to believe and do what's right, *we should know better*. Many people have engaged with many of my ideas on this over many months, and many details have changed since I first argued for the rationality of moral belief. As a result, I better appreciate the task of believing and doing what's right in this world. I now know better.

For this, I owe this university's philosophy department sincere gratitude. Thank you Dr. Scott Woodcock for your many interventions and tireless patience; it is but for your efforts that others tolerate this thesis, if at all, spared of some protracted and trying prose. Dr. Colin Macleod, thank you for tolerating this thesis as well, and for providing the right advice, at the right moments. Alastair, Emma and Buddy, thanks for the encouragement and inspiration. I take some responsibility, but little credit, for what's best in what's left.

I acknowledge the broader campus community's support of this learning: many of you helped make my time here count. Heather and the C6 team, especially, thank you for my most productive months of this project. I will miss your work dearly, and do not doubt that many other students will too. Everyone I haven't yet named, I apologize in advance for forgetting how to thank you here, and for what you're about to read (for which, indeed, I thank you nonetheless!). Most of all, a most considerable thank you to Mari, Tobias, Julia, Josiah and Simon. Your unbelievable companionship has made this an unforgettable experience. May we celebrate many more together!

Preface

This thesis is about moral belief. It is a thesis about what we do when we talk about morality, why we do it, and if we succeed in doing what we want to do as moral creatures. I also look at if some theories of moral truth and belief measure up to what the science can tell us about the topic. For example, in this thesis, I try to show that we *at least* take ourselves to tell each other about the world when we say that murder is wrong, or that Hitler is evil, or that caring for children is good. In this thesis, I am interested in whether what we say when we believe or talk about morality is ever true, and if so, whether the true things that we do say and believe are common because of their truth.

One possible answer to this question is that we don't ever say or believe anything true when we judge (for example) that murder is wrong. Our moral beliefs are never literally true, because there are no facts that would make them true. So our moral beliefs are always mistaken. I argue against this error theoretic approach in my thesis.

Another kind of answer is that we can believe things that are true in morality, because there are moral facts. In fact, when our beliefs are correct, this is because we are acquainted directly with their truth. There are two ways of making this answer work: one way is the realist way, and the other is the not-realist way. The realist insists that morality is real. Morality is real if (for example) it is true that it is wrong to torture kittens for fun *because* the world is set up in just such a way that kittens have certain rights, such that there are real prohibitions against torturing them, whether we recognize this or not!

Notice that when I say kittens have some moral standing, this is dissimilar from saying that kittens are cute. (Usually I don't care too much about whether you think that puppies are cuter than kittens. That's not really a dispute that has a right answer for everyone!) But morality is different. I am not *merely* sharing my opinion when I express a moral belief. I am not reporting it to you for the sake of giving you information *only* about me, or about how much I like kittens, or just about the rules that *I* want to follow. I am telling you that you must not torture kittens either, when I share this belief with you. This is the default moral theory we have inherited in our (Western) philosophical tradition, from Plato and Aristotle, to Kant, Mill and Sidgwick.

Irrealists reject all this. Their alternative is that our moral beliefs can be true, only because of attitudinal "facts" about us. (David Hume is often credited with inspiring this kind of approach to morality.) On this kind of view, the success of moral belief works differently from my belief that, say, it is raining outside. The rain doesn't make only my belief true, it makes the same belief true for anyone, even if they are very different from me. But if realism is not true, very different people with very different attitudes *can* be right in their moral beliefs, even if these are radically different from what *I* am right to believe in the moral case.

For example, what makes it wrong for me to kill children depends on normative facts about my practical identity as a (Canadian) citizen, as a parent, as a human, and so on. In other words, whether I am right to think I should care for children depends on whether I in fact accept what those roles require of me. As it happens, if I do ultimately value happiness, or if I value humanity, I would be making a mistake if I thought I should kill children. But the irrealist can't accuse me of making a mistake for ignoring or opting out of these roles and

rules altogether, granted I do this in a coherent way. Maybe some people don't really have a coherent identity of being a citizen or a human in the usual way—they don't care about the rules that govern this kind of identity. They don't value humanity. They might have some other kind of identity altogether. It's possible for these people to be correct to think that they should torture people, if irrealism is correct.

This view is the primary target of arguments in my thesis, especially in the last three chapters. But along the way, I also try to discuss the various ways that non-moral explanations of moral belief do (or don't) give us more reason to prefer one theory of moral facts over another. Think of this as a kind of contest. You and I are the judges of this contest: we get to decide whether the science, along with other considerations, give us more reason to decide in favour of a winner. If we think that one view wins, we have to say something about why that is. We have to give reasons that convince the neutral to think what we think. And that's exactly what I try to do.

I start by acknowledging that each theory of moral belief and truth has to reconcile the adaptive features of morality with its apparent factive features. Morality, for example, is adaptive if it turns out that our capacity to things like "torture is wrong", or that "caring for children is good", can encourage an adaptive kind of cooperative or moral behavior. For example, the behavior of caring for children in the past is part of the explanation of our past reproductive success, which is part of the explanation of our current characteristics, including our moral habits, including our moral beliefs. We care for children because we believe it's the right thing to do. *Of course, we think we believe it's right because we recognize the appropriate moral truth—we think, first-personally, that morality is factive.* But third-personally, science can't decide whether we're correct to think this. All that the

science shows (some say) is that it *was* adaptive for to care for children, and *that* explains our beliefs on the topic. In general, in nature, the way we see things doesn't have to be accurate or correct. It just has to make us make more of us.

Now the realist thought that the **fact** that it's wrong to kill children is the reason that we believe that it's wrong. But actually, the fact that killing children is not adaptive for us is enough to explain our belief. The scientific explanation does not require moral facts and it does not help explain moral facts—it only explains *beliefs* in such moral facts. So the realist has significant difficulty demonstrating the reality of particular moral facts. We can't use the methods of science to show us these facts, and since the scientific methods don't involve the facts we want to explain, our factive methods of explaining moral belief are provisionally off the table. It looks like this view is disqualified from this contest. It can't compete on these grounds alone.

So this thesis moves on, and looks for another view that does better. Let's say our moral beliefs never succeed in representing facts, because there aren't any. For the error theorist—for people like John Mackie and Richard Joyce—it still makes sense to talk about morality *as if* our beliefs are true according to the world, in the same way it makes sense to talk about the characters and plot of a story *as if* there are real people and events that make the story true, even though the story is a fiction.

This view gives us a unified explanation of moral belief. Morality isn't really part of the world, so there are no facts that require explanation. All that requires explanation is our belief, and error theory and science both make sense of *that*. There isn't anything the sciences can discover, it seems, that would show the theory to be false, unless there is a

scientific way to demonstrate the truth of a moral fact. For those of us who think that science can't find room for moral properties, error theory should be very attractive.

But one difficulty is that error theory doesn't do the same thing for us when it comes to normativity in general. (When I talk about normativity in this thesis, I am talking about the rules, or the truths about thinking and doing that we endorse or are committed to when we give reasons for belief and action, or when we evaluate the people, or things and events around us in terms of these reasons.) Murder is wrong (for example) because it prevents our happiness or disregards our right to life. That's a "normative fact" that error theory denies. But here's another that error theory is often okay with: Sam knows that it's raining. Usually when we talk like this, we're saying that he in fact thinks it's raining, and that it *is* raining, and most importantly, Sam's thought is somehow good, or correct (usually because Sam has reason to think that it's raining). But I think the error theorist has to decide whether this evaluation about what Sam knows or what he should believe is normatively factive, and they have to make sure that this decision is consistent with the denial of objective facts about what Sam does well, or what Sam should do. The problem is that error theory is often internally inconsistent or ambivalent on this point. The point here is that for those of us who believe that knowledge is real, or for those of us that are sure that there *are* really better and worse reasons in general, error theory might seem inconsistent, or schizophrenic, in accepting these reasons in general but denying them in particular cases.

Now, the irrealist possibility I considered throughout this thesis is designed to avoid exactly this kind of problem. Sharon Street, for example, agrees with the error theorist that the objective moral facts I just described are not real. In fact, for Street, there is no objective fact about what counts as knowledge either. But then Street thinks we are not in error at all

when we believe murder is wrong. This is because her view is committed to what I (want to) call *alternative* moral facts. This is an incredibly frustrating view. Actually, I think it is literally incredible. According to Street, our normative beliefs are often true, even though there is nothing in the world that fixes this truth for everyone. This *subjectivity* allows Street to make plenty of room for normativity, for plenty of norms, including facts about what Sam knows, or what Sam should believe, as well as facts about what Sam should do. So the irrealist manages to capture the internal structure of reasons really well: they manage to say exactly the same thing about reasons for belief and reasons for action. Again, usually, the error theorist doesn't do this. So there's a curious kind of advantage to these alternative facts. They might look illegitimate, but they are certainly attractive.

Unfortunately, irrealism does not have quite the same attraction as error theory when it comes to preserving our common sense understanding of morality. It actually seems quite unnatural, since it's not the view I think we actually believe, or the view that it would be adaptive to believe (at least pre-theoretically). When we make a moral argument, or when we believe it's wrong to kill children, we really do behave and act as if the moral truth was out there in the world, independent of us (often not at all unlike the way we act and behave according to beliefs we acquire from perception). In fact, Richard Joyce argues that it's just this objectivity that made morality adaptive. We're not supposed to think that caring for children is optional. That's not the point of morality *at all*. If Joyce is right, moral belief won't work this way, if it's a poor strategy for reproductive success. We're supposed to think we're making the wrong choice when we choose to hurt children. Of course, I love my children, that's one reason *I* care for them. But that's not always why others care for them. And I don't think that I can go ahead and hurt your children more, just because I don't like

them as much. I think the non-moral explanation predicts exactly that: it predicts that the best cognitive strategy in our species is that each of us really believe that hurting each others' children is wrong whatever each of our own attitudes. We may not like children—they may sometimes annoy or inconvenience us. But we still cooperate as a society and as a species in caring for each other because we have a system of moral norms that appear to be genuinely objective. Error theory and realism seem to at least capture more of the moral psychology and intuitions that the disciplines in the social sciences describe. On the other hand, if irrealism is true, the way we think about morality, and even our phenomenology of moral experience, is confusing and at odds with what's *natural* to believe. So irrealism might involve re-thinking morality, or changing who we are. That's a high price to pay.

There are many more trade-offs between the candidates in this kind of debate. My strategy in this thesis is to point out problems with these different approaches to morality. I survey more substantial problems in what follows. There are good objections to all these views. And I try to respect our default intuitions and concepts in addressing these problems. This strategy isn't a strong proof for my own position, but my solution is meant to capture *at least* what's attractive about each of the options I reject. That's what my solution should do. It should take what works and what seems right and leave what does not. Of course, you are welcome to poke holes in my own view too, in order to eliminate it from the contest. That's only fair.

But let me sketch my position first. My position is that normativity is objective. It's real. There are real facts about what I should believe and how I should act, and I could be wrong about this truth, especially in morality. If it is true that murder is wrong, this would be true even if no one believed it. So my "skepticism" accepts the reality of moral facts, but merely

questions whether we get these right. I do **not** rule out reasons altogether. I want the right kind of doubt, in the right degree. This is because we are sometimes more right when we doubt our natural beliefs than when we blindly believe what we want. And that's right because it is correct to care about truth. And if it is correct that true beliefs are better than false beliefs, then there are evaluative facts, because that's one of them. That's exactly why we think that truth itself solves the evaluative problem about what to believe.

If this is true, we can make sense of moral skepticism. The view I am describing is right if there are real reasons to think there is a right view. If there aren't, there's no point to our contest. There's no reason to decide on a theory, if each of the theories denies that we have access to the reasons that require this contest in the first place. In this thesis, I hope to show how to win the contest between meta-ethical approaches, often by shifting the narrow focus from the moral to the normative, and challenging the doubters to win the fight on *that* ground. And you can't win that contest by doubting too much. You can't win a contest if you don't play by the rules. You have to doubt something, but you can't doubt everything at once. And giving up on objective normativity is a kind of radical doubt that is impermissible in this contest.

That's why I argue that there are better or worse beliefs: there are correct beliefs, and there are incorrect beliefs. And so there must be right answers about what we should do and believe, even if we aren't certain about what these are. If anything, science does not show that morality false, or unreal, but it helps illuminate a certain kind of distance between our beliefs and their truth. I don't think that should surprise us. Morality, for all we know, is something we're wrong about a lot. (And sometimes I am not sure we are getting any better at it either.) I am happy to concede the skeptic this point.

In fact, I encourage this intellectual humility. It is a good thing. We need (more of) it. I am sympathetic to *this* skepticism because it is the best way of *not losing*: it is closest to the realist view that is internally consistent (like irrealism) but also thoroughly objective and conservative (like error theory). But this preface, and indeed the thesis, is just a broad kind of comparative exercise. At the end of the day, I don't really think that the scientific stuff, with all due respect, decides a theory of moral truth. These are two very different domains of inquiry. Empirical information is certainly useful. At the very least, examining the way we actually work encourages a more rigorous study of morality. It allows us to better compare the trade-offs and consequences of different theories in moral philosophy.

A consequence of moral realism for example, is that we lose some moral certainty. So we must be realistic. But maybe the main reasons to believe moral realism are independent of and sometimes immune to these skeptical debates. So (although I call myself a skeptic) I think the realist can take some encouragement from surveying the ways different varieties of moral doubt are unsatisfying and incomplete, if not altogether incorrect or unbelievable.

1

Knowledge and varieties of disbelief

1.0: Introduction

Knowledge involves, at minimum, beliefs, the facts these beliefs concern, and some evaluative standards or measure of doxastic success with respect to some facts. For example, I know that I exist just in case (1) I think that I do; (2) this thought is true; and (3) I have good (enough) reason to believe (2). In this instance, my awareness justifies my belief or reliably “tests” its truth, while the more general fact of my existence causes or explains this awareness, and therefore my belief, and my knowledge of my existence.

Of course, knowledge is often more difficult to come by than this. I certainly think that torture is wrong. I also think I have good reasons to believe this—say, because I think pain is bad, among other things. But there are sometimes good reasons to doubt beliefs like this. For instance, my reasons for my beliefs on torture may be misguided. Maybe politicians or the media have misrepresented facts in their rhetoric about torture, and maybe such propaganda has influenced my beliefs. That sort of influence is a reason to reconsider my beliefs about torture. This thesis is about such reasons for disbelief.

Skepticism is a position of considered disbelief with respect to some propositions. More precisely, skeptics argue against at least one of (1) the facts some beliefs concern or (2) the

evaluative success of such beliefs with respect to these “facts”. For instance, say Sophie used to believe in Santa Claus. But then, Sophie figured out where the presents *really* come from. So after understanding how Christmas could work just as well without Santa, Sophie became a Santa-skeptic, so to speak.

The moral skeptic thinks that we should doubt our moral beliefs because our understanding or explanation of these beliefs cannot show that they are true. The skeptic argues that rather than knowing what’s right, we inhabit a shared, albeit illusory moral experience in much the same way some can carry on with their beliefs in Santa, or witches, or homeopathy, and the like. In these cases, we believe what we do because it was natural or seemed like a good idea at the time—someone told us something was true, or we just grew up believing it, sometimes without considering why we believe what we do. This is the way most believing works. And much of the time, these ways of believing work well, since they are reliable ways of learning many true things.

But then there are the cases where beliefs like this do not work out at all. Skeptical cases of the sort I consider challenge the merits of our beliefs in two ways. The first challenge explains away our reasons for some of our default beliefs, such that we cannot be certain of their truth. A child, for example, may struggle to explain the presents under the tree once she begins to reconcile the Santa myth with the basic physical characteristics of, say, reindeer, or aeronautical travel, and so on. We sometimes say—quite rightly—that someone who carries on with their childish beliefs despite these considerations *should know better*.

On the other hand, the second sort of skepticism provides conclusive reasons against the truth of our default beliefs—they show that these beliefs are false.¹ This is more like a child finding out who wraps up her Christmas presents—as soon as she knows that her parents are behind it, she also knows that it cannot be Santa. The two ways can coincide, but the differences are important. Briefly, the first skeptical challenge is a problem about justification, the second about truth. Although these problems are closely related, and although both can work to rule out whatever knowledge we thought we had, they should not be confused. Or so I will argue.

In arguing this, I agree with many moral skeptics that some moral beliefs are not true or justified. Basically, it is difficult to show that many of our moral beliefs are objectively true. However, in what follows, I show that this moral doubt conflates the two distinct objections to knowledge outlined above. First, I might mean that the usual way of evaluating truth does not apply to moral belief successfully. For instance, if there is nothing in the world to make moral beliefs *objectively* true, it follows that our moral beliefs aren't true as such. Second, the standard analysis might itself be correct, even if our beliefs are often false. This thesis takes this second stance. I am not sure, in short, that the world makes all our moral beliefs true. I am quite sure, in fact, that many of our moral beliefs are not true, in that they are mistakes about objective features of our world. So while I conclude that our moral knowledge is often incomplete, I do not doubt that moral facts are objective.

In this first chapter, I define knowledge and skepticism, and introduce some varieties of the latter (§1.0). I compare some analyses of moral sentences, and focus on two veridical

¹ Or, to be precise, we cannot (quite literally) believe it any longer, once we think “it” is not true. I return to this result in Chapter 4, where I argue that the concept of belief *necessarily* involves truth.

accounts of moral belief (§1.1). I then introduce a general empirical problem for these explanations of moral belief. Recent examples of this empirical problem in meta-ethics use biological explanations of morality to illustrate the comparative inadequacy of naïve explanations of moral belief. I introduce Richard Joyce's and Sharon Street's discussion of this problem (§1.2, 1.3), before concluding the chapter (§1.4) with a brief comparison of these two arguments and their analogous cases from §1.0.

1.1: Making Sense of Morality

Much of this thesis will compare two arguments against objective moral knowledge. The relative success of these two arguments depends on the character of moral truth. On an irrealist account, moral truth is not objective, so objective moral knowledge is not possible. On a realist account, objective moral knowledge is not possible if our beliefs are false, or if they are not justified. In both cases, knowing a moral fact involves more than just believing it. Again, I know that torturing *S* is wrong just in case (1) I think this; (2) this thought is true; and (3) I have good (enough) reason to believe (2). For example, I might think it is true that torture is wrong in a particular case because I think pain is bad, or that people have rights, and so on. These are reasons to think that there's a fact about what's right in this case. The difference between the two views is the sense in which (2) is true. For the cognitive irrealist, the truth of <it is morally wrong to torture *S*> is not objective: it is not true independent of our particular practical attitudes. The realist claims instead that the truth of this sentence is objective, in that it is true independent of our values.

Consider Ronald Dworkin's statement of the view:

...when I say that the badness of abortion is objective, or that abortion is objectively or really bad [...] I would be emphasizing that, in my view, the deliberate destruction of human life at any stage is impermissible for reasons that in no way depend on my or anyone else's reactions or tastes. The claim that abortion is objectively wrong seems equivalent, that is, in ordinary discourse, to another of the further claims I made: that abortion would still be wrong even if no one thought it was. That, read most naturally, is just another way of emphasizing the content of the original moral claim, of emphasizing, once again, that I mean that abortion is just plain wrong, not wrong only because people think it is. (98)

This is the view, in Russ-Shafer Landau's words, that: "we don't, singly or together, have the final say about what is right and wrong [...] even the ultimate moral commitments of individuals and societies can be mistaken" (11). I think many of us already accept much of this approach to moral belief. We make evaluative judgments, for example, when we talk about torture: you might state that we should not torture other people, and I might believe that we sometimes should. And you might object to my belief because you think I am missing some *moral* facts about the cases we consider. You might say that pain is *bad*. You might add that no one reasonably agrees to torture as a general rule, and that torture is *wrong* because it does not acknowledge our fundamental human rights. And so on.

I make several assumptions when I say that moral sentences represent some moral facts more or less correctly. I assume, first, that at least some moral sentences express beliefs. In saying this, I mean that these sentences are sometimes true, or sometimes false. Ultimately, I also think that the truth of a particular moral belief is objective: when I learn about the torture of Omar Kahdr, say, and I judge that he was wronged, I think this belief is true for anyone who thinks it, at any time they think it. The belief is true because we should not torture Kahdr, and that fact about his case is the truth, period.

I characterize what moral realism has in common with other alternatives in the table below. The two most relevant alternatives to realist theories in this thesis are moral error theory and moral irrealism. The use of my term “irrealism” is relatively precise: sometimes, the other alternatives to realism below are also sometimes called irrealist or “anti-realist”, because they are *not* realist theories either. But I think the irrealism I describe is an explicit rejection of realism. For example, the cognitive irrealism of Sharon Street has less in common with realism than error theory. For Street, the truth of a moral sentence depends on our attitudes. It is possible, say, that the torture of Kahdr is wrong for some people, or some societies, at some times, and not others. The truth of beliefs like this depend on what we happen to value. Were we to value pain, for instance, our belief that <the torture of Kahdr is wrong> may be false.

In realist theories, however, the truth of our moral beliefs does not depend on what we value. For example, if it is right to care for our children, this would be true whether or not anyone cares about children, or should care about children based on their own values and opinions. Of course, I may happen to love my children, and maybe this is why I care for them. Maybe I will stop caring for them if I stop loving them. But caring for these children is either the right or wrong thing to do, and our own attitudes do not make caring for these children right or wrong. If it is true that I should care for my children, this is true even if I do not value them. In the table below, realism and error theory both have this result: possible changes in what we value do not change the truth of particular moral beliefs. In my thesis, I will argue that this is the right way to understand an evaluative belief. As a consequence, however, there is a real worry that these beliefs are often in error.

TABLE 1: Comparison of Meta-Ethical Theories

	Do some moral sentences express beliefs?	Do some moral sentences express true beliefs?	Do true moral sentences only and necessarily follow from our particular practical attitudes? ²
Cognitive irrealism	✓	✓	✓
Cognitive realism	✓	✓	✗
Cognitive error theory	✓	✗	✗
Non-cognitivism	✗	-	-

The theory in the third row above simply states that the truth of a moral belief is objective, although these moral beliefs are *never* true. Finally, a non-cognitive analysis of sentences like “torture is wrong” rejects the above three positions altogether, since on this view such sentences endorse a norm, express an attitude, or convey a command, rather than state a belief. This last alternative shows that the other positions, for all their differences, are similar because they are *cognitive*: moral sentences, according to these views, represent the world (more or less successfully).³ On all three positions, some moral sentences are beliefs. And on the first two, these beliefs are sometimes true.

² More exactly, if there are true sentences, are they only and ultimately entailed by our evaluative attitudes? As Louise Hanson puts it, realism is the view that: “Moral statements have truth-conditions; at least some moral statements are true; and these truth-values are (relevantly) *independent of the attitudes of (even idealised) agents.*” (2016, emphasis mine)

³ In taking moral judgments to express beliefs that are true or false, this thesis assumes moral cognitivism. This is because there is already a cognitivist consensus (or bias) among the contemporary moral skeptics I consider here. That said, according to this skepticism, some of this consensus is just a consequence of the

The way these beliefs are true will matter in what follows. In fact, this thesis will often focus on the difference outlined between the first two views in the third column above. I am primarily interested in the *objectivity* of cognitive realism, and will label the collection of views in this category “moral objectivism”.⁴ “Moral subjectivism”, or theories that take the truth of our moral beliefs to depend on *us*, have an apparent advantage over objective (realist) views, since the truth of such irrealist beliefs are immune to an empirical problem in meta-ethics.⁵ On the other hand, it is very difficult to show that we have good enough reason to think our moral beliefs are true if such truth is objective. If this is true, we may have to give up objective moral knowledge on the realist, or objectivist account.

1.2: Objective Moral Facts and the Empirical Problem

We think we are obligated to care for our children—we think it is good, or required, or compulsory, and perhaps praise-worthy, that we feed, dress, shelter, and mentor our own offspring. At the very least, most of us agree that to do otherwise is impermissible: for

cognitivist bias in our species, and these discussions are often explicitly conscious of and critical of this bias. For sustained arguments against non-cognitivist interpretations of moral discourse or “folk morality”, as it were, see Joyce 2001 (9-16). For a taxonomy of the skeptical positions in the literature, including the skepticism about “truth-aptness” I will dismiss hereafter, refer to Sinnott-Armstrong’s 2006 (pp. 13).

⁴ Throughout this thesis, I prefer the term moral objectivism to moral realism (despite the latter’s pervasive—even, perhaps, homonymous—use in the literature I discuss (Copp 2007: 12, Cuneo 20, Enoch 2011: 2, Joyce “Anti-Realism”). This terminology may help prevent misunderstanding. For example, moral irrealism is *not* immoral or less moral than moral realism. For Street, rejecting realism is not about rejecting morality at all, but more properly about denying that moral reality is *objective*. Rather than rejecting value outright, Street’s constructivist project is positive, cognitive, and ambitious: she hopes to regain and defend moral knowledge—she calls herself an anti-realist, not a nihilist. My disagreement with her project is over the subjectivism it involves, rather than her more easily mis-read rejection of “real” morality. In other words, Street believes in morality, she just disagrees with a certain meta-ethical view on the nature of moral truth.

⁵ Very general problems of a similar sort date back to at least John Mackie’s and Gilbert Harman’s discussion of moral properties (or lack thereof) in natural explanations of the world. This attitude, in turn, has a long history involving G.E. Moore, and his objection to reductionist explanations (of the moral to the non-moral). See Darwall, Gibbard and Railton for a discussion of this history (1997).

example, it is wrong to kill our own children. The fact that this is wrong, according to whatever way our preferred moral theory spells out this fact, is all we usually require to explain why we believe that we should not kill or intentionally and unnecessarily harm children. But this explanation of why we care for children already assumes that morality involves beliefs and that these beliefs are more or less connected to facts. In particular, it takes for granted that the moral facts are part of the explanation of what we believe.

This is what we might hope for with all of our beliefs. But if it turns out that we believe most of these things because that's just how we survived *whether or not these beliefs happen to be true*, the truth of whatever we happen to believe is far from obvious. In fact, our naïve account of moral belief is itself far from obvious if it is right that morality is just something we are born into—something our species can't help *but* believe.⁶ According to a “naturalistic” multi-disciplinary research program, there are many non-moral facts—the various characteristics of our species in the past, for instance—that explain our present characteristics and functions, including our moral practices at the level of cooperative groups.⁷ In particular, the behaviour of caring for our children is a characteristic our species has had for a long time. We believe we have moral reasons to contribute to and continue the cooperative care of our children. And the non-moral or natural explanation of

⁶ There are, to begin with, obvious (short-term) influences on fitness via our doxastic etiologies: that of social or cultural customs on the one hand and more self-interested biases on the other. Here, I intend to generate a more simple, immediate epistemic puzzle about our doxastic etiologies that both previews and postpones the more detailed and important discussion of long-term biological etiologies below.

⁷ Robert Nozick articulated an early version of this empirical problem, writing: [scientific] “explanations seem to undercut the objective status of ethics by showing that ethical behaviour can be explained without in any way introducing ethical truths. If ethical behaviour is adaptive, if that behaviour increases inclusive fitness, and is genetically based and heritable, over the generations it will spread more widely. Ethically behaving individuals will leave more great-grandchildren or (given kin selection) great grand-nephews and nieces similarly disposed” (344).

the moral belief that <we should care for our children>, for example, shows that the truth of such beliefs is not involved in explaining the present characteristics of our caring for children. Put simply, our moral beliefs were adaptive even if they were not true.⁸

Richard Joyce argues that our beliefs in objective moral truth are adaptive *even if false*.⁹ More specifically, beyond just thinking it wrong to kill children, the sorts of beliefs that make certain helpful behaviours more or less *mandatory* are adaptive because they most effectively guard against “free-riding”—that is, the phenomenon of defecting from reciprocal exchanges. In biology, such beliefs, desires, intentions and other psychological attitudes can operate as *proximate mechanisms*: natural processes do not select for beliefs directly, but select for behaviours that are characteristic of these beliefs. For instance, believing that it is right to care for children characterizes successful parenting if beliefs of this sort accompany the continued trait of successful parenting in successive generations. Conversely, individuals believing instead that caring for each other and their children is merely optional will find it a lot harder to 1) live cooperatively in large groups and 2)

⁸ It only seems fair to credit Alvin Plantinga with an original articulation of this argument. According to Plantinga, “blind” natural selection makes it extremely unlikely that we have any knowledge at all, since there is no reason, in principle, to think that the biological processes that shape our beliefs guarantee their truth. Plantinga pre-supposes a “super-natural metaphysics” to solve this puzzle (237). Here, I do not pretend to decide the question of whether truth-tracking is an evolutionary adaptation; I merely want to point out that, in at least some cases (mysticism, morality, etc.), our species can get by with at least some mistaken beliefs.

⁹ Richard Joyce and Philip Kitcher provide the most complete, and perhaps, original presentation of this problem, and so provide a useful overview of the scientific details it involves. However, since 2006, when Philip Kitcher, Michael Ruse, Sharon Street and Richard Joyce individually published influential analyses against objective moral knowledge, further discussion of these details have flourished. And unlike some of their more generic, historical precedents, these discussions post-2006 do more than argue that moral truth does not explain moral belief—they argue that our moral beliefs are mistaken *whatever* the status of moral facts. And once we realize that our moral beliefs are more often false than not, we cannot justify any one of them. At any rate, we can’t know anything about morality, if none of our moral beliefs are true.

produce children that exhibit the altruistic traits that best promote their genealogical future in large groups.

To see this, compare the two beliefs <I should not kill my children> and <I should not kill a deer>. If Joyce is right, the first belief is more adaptive, biologically, than the latter: killing deer for food is far more likely to keep our species around than killing our children. In this way, moral facts make no difference to a natural explanation of why we think it is significantly worse to kill a human—especially a member of our own family—than it is to kill an animal, especially for food. In this way, the *belief* that we should not kill our children is responsible, among many things, for why we are still around. But the *fact* that it is wrong to kill children, even if it is a fact, is not involved in a natural explanation of why we think killing children is wrong. Put simply, moral beliefs are adaptive because of facts about biological fitness, rather than objective moral facts.

Indeed, a large variety of scientific explanations for our moral beliefs *are* available and these variously predict that moral beliefs are adaptive in groups throughout human history.¹⁰ More specifically, moral beliefs are involved in a long and complex history of promoting helpful behaviour among hominid groups. The helpful behaviour, in turn, has its

¹⁰ That said, in what follows, I am not interested in whether Joyce's and Street's empirical premises are correct; I am primarily interested in evaluating the conclusions they draw from their data; these may not be correct, after all, even granted all the relevant biological facts they assume. This is why the last section set the stage, by way of examples, for these wide-ranging assumptions. Of course, if some of the descriptive details assumed in the antecedent of the arguments examined here turn out to be false, the skeptical challenge may well fail on these independent empirical grounds. However, even if the skeptical challenge fails in the case of moral belief for particular *empirical* reasons, its relevance can extend to other domains with other details involving analogous doxastic processes (see, for instance Copp's 2008 (190) and Joyce's 2016 (143)). It is worth examining, therefore, both whether the challenge works as advertised (even granted these empirical details), and whether what makes it work, if at all, operates independently of these details.

probable basis in a combination of kin selection, mutualism, reciprocal exchanges, and group selection (2006: 40). According to Joyce, the consequent practice of behaving according to prescriptive beliefs confers adaptive advantages on groups by regulating helpful or reciprocal behavior. So natural processes selected for our doxastic characteristics in much the same way that it has selected for physical characteristics, so that we have ended up with the particular kinds of moral habits (e.g. caring for our children, because we believe it is *right*) in the same way we ended up with particular kinds of bodily functions.¹¹

According to Joyce, these available scientific explanations of our moral beliefs variously show that those beliefs are not justified once we become aware of the relevant biological etiology or explanation of those beliefs.¹² His argument from 2006 runs as follows: 1) The proximate mechanism of moral judgments confer adaptive advantages on creatures like us 2) biological processes select for moral creatures like us—that is, moral facts are not involved in the natural explanation of our moral judgments, and, finally 3) moral judgments are unjustified, (earlier: moral beliefs are “largely untrue” (2001)). So, because we have no scientific account of moral facts themselves (at least, none that makes sense of their

¹¹ See Street 2006 (143) for this comparison.

¹² Admittedly, this statement glosses over the complexity of a multi-disciplinary research program involving mechanisms beyond biological natural selection—Joyce acknowledges, for instance, that these natural explanations “involve data from numerous scientific disciplines: experimental economics, neuroscience, anthropology, primatology, and various fields of psychology” (2008: 3). Philip Kitcher agrees: “To work out a convincing story about how current moral practices might have emerged would require a vast amount of information from diverse fields: evolutionary theory, primatology, psychology, anthropology, and history (and maybe more)” (2006: 183).

“practical authority” (2006: 191)), Joyce argues that our beliefs in moral facts are not justified.

This is to say nothing of moral facts themselves. And the moral non-naturalist—the theorist who thinks that science *never* accounts for moral facts—will not find much comfort in this concession. After all, it is primarily beliefs in *non-natural* moral facts, argues Joyce, that are not justified (2006: 210). And these beliefs are not justified, according to Joyce, because they cannot track the *possible* non-natural moral facts. For instance, were it a non-natural fact that killing children were permissible, our belief would not be sensitive to this fact. Biological processes would still select for the sorts of beliefs that keep children alive, and since the (possibly) true belief that killing children is permissible will not prevent deaths as well as the belief that killing children is wrong, we would believe that killing children is wrong no matter what the fact of the matter is.

So said, the current scientific explanation does not directly involve moral *truths* at any stage of the genealogical process. What makes reciprocal traits predominant in hominid groups over time is not some fact that, for example, humans have a right to life, but the *belief* that we usually shouldn’t hurt the humans we live with (because we *think* they have a right to life, for instance). Nonetheless, Joyce has recently recognized that “the possibility that moral facts might find a place in the evolutionary genealogy was always acknowledged” (2016: 155), and has revised his 2001 claims that “moral discourse consists largely of untrue assertions” (125)—that is, “[moral beliefs] are, for the large part, untrue” (186)—to “all moral judgments [currently] lack justification” (2016: 146). Joyce’s conclusion, as such, makes some concessions to the moral naturalist—it’s not that science

has no place for moral facts, in principle. It's just that these facts are not part of our current scientific theory of moral belief.¹³

So far, we have been agnostic about what these facts look like. Although Joyce does not decide what moral facts are in particular, he believes that such facts, natural or otherwise, are at least a real *possibility*. But this possibility is not good enough: more exactly, we do not have a good enough justification or source of moral beliefs to support beliefs “one way or the other” (2006: 211). This puts Joyce in the position of Sophie the Santa-skeptic when she cannot fully explain the Santa story, but doubts it anyway (or our position when we dismiss whatever some suspect news source tells us about torture). Put simply, Joyce knows just enough science to doubt that we have the correct moral story. At the very least, the moral stories we've told ourselves ever since antiquity are, he thinks, unjustified.

1.3: Moral Irrealism and the Empirical Problem

Street argues instead that moral knowledge is possible *according to a constructivist theory of moral facts*. On this view, our moral beliefs count as knowledge if they cohere with each other in reflective equilibrium—roughly, our moral beliefs are true and justified if they follow from certain facts about our practical attitudes (2006: 154). This is Street's view. I

¹³ Notice, however, that even if science does not have a place for moral facts, this does not mean that they cannot be true. For example, the right moral facts might be roughly analogous to mathematical facts in that they are derived from normative concepts (analogous to mathematical axioms) rather than discovered in the empirical world. The philosophers who hold this view are non-naturalists, and their sort of “facts” (indeed, it is at least as debatable in morality as in math whether abstract definitions count as facts) are the primary target of Joyce's argument. So I am granting the moral skeptic this target—a robust theory of cognitive non-naturalist realism—in characterizing their argument(s). For the most part, however, I want to remain neutral in the debate between naturalists and non-naturalists, because I think both have adequate responses to the moral skeptic. For a characterization of the sort of view I assume above, see David Enoch (2011). For a characterization of Joyce's and Street's anti-naturalist bias, see Ramon Das (2016).

will not evaluate Street's theory of moral knowledge here (see Chapter 4). What's important for our immediate purposes is Street's argument against objective moral knowledge, and whether her irrealism solves the empirical problem.

Street discusses the empirical problem in the form of dilemma. If morality is objective, either our beliefs are sensitive to objective moral facts or they are not. If they are not sensitive to these facts—if there is no relation between, for instance, non-natural objective moral facts and our moral beliefs—then our beliefs do not correspond to moral facts. On the other hand, if our beliefs are sensitive to the facts—if there is some causal relation, for instance, between natural objective moral facts and moral beliefs—then these moral facts are part of the natural world. But surely, the scientific explanation does not involve or discover moral facts, and the moral explanation is superfluous. So objective moral facts are incredible.

The first horn of Street's dilemma discusses the view that moral facts and moral beliefs—that is, the adaptive doxastic product of natural selection—are not related in any scientific way. In other words, on this view, the facts themselves have no direct influence on what we happen to believe. Consider Street's characterization:

On this view, allowing our evaluative judgments to be shaped by evolutionary influences is analogous to setting out for Bermuda and letting the course of your boat be determined by the wind and tides: just as the push of the wind and tides on your boat has nothing to do with where you want to go, so the historical push of natural selection on the content of our evaluative judgements has nothing to do with evaluative truth. Of course, every now and then, the wind and tides might happen to deposit someone's boat on the shores of Bermuda. Similarly, every now and then, Darwinian pressures might have happened to push us toward accepting an evaluative judgment that accords with one of the realist's independent evaluative truths. But this would be purely a matter of chance, since by hypothesis there is

no relation between the forces at work and the “destination” in question, namely evaluative truth. (2006: 122)

In other words, we have evolved “towards affirming whatever evaluative content tends to promote reproductive success” (122) and so, “many or most of our evaluative judgments are off track” (122).

The alternative is a natural explanation of moral facts and their relation to moral belief. On this second horn, the objectivist might argue that moral facts are reducible to some naturalistic facts and are therefore part of the scientific explanation of whatever beliefs turned out to be adaptive. Consider Street’s characterization of this approach. First off, she acknowledges, with Joyce, that “[the tracking account] is the more plausible route for the realist to take” since “we think that a lot of our evaluative judgments are true” (125). For instance, “we may understand evolutionary causes as having tracked the truth: we may understand the relation in question to be a tracking relation” (2006: 125). This *tracking account* “offers a specific hypothesis as to how the course of natural selection proceeded and what explains the presence of these judgments is explained by the fact that these judgments are true, and that the capacity to discern such truths proved advantageous for the purposes of survival and reproduction” (2006: 126).

But Street argues that this hypothesis doesn’t survive scientific scrutiny. Among other things, Street argues that her adaptive-link account is *better* than the moral naturalist’s alternative. More specifically, “it is more parsimonious, it is much clearer, and it sheds much more light on the explanandum in question, namely that human beings tend to make some evaluative judgments rather than others” (2006: 129). In particular, the advantages of such an “adaptive link account”, according to Street, are just that “it illuminates a

striking, previously hidden unity behind many of our most basic evaluative judgements, namely, that they forge links between circumstance and response that would have been likely to promote reproductive success in the environments of our ancestors” (2006: 123). These “judgements about reasons – and the more primitive, ‘proto’ forms of valuing that we observe in many other animals – may be viewed, from the external standpoint of evolutionary biology, as another such mechanism” (2006: 127). Street concludes with the epistemic upshot of these evaluative observations:

The tracking account has no comparable explanatory power. Its appeals to the truth and falsity of our judgments in question sheds no light on why we observe the specific content that we do in human evaluative judgements in the end, it merely reiterates the point that we do believe or disbelieve these things. When we couple this final point with the points about the parsimony and clarity of the adaptive link account as compared to the tracking account, it is clear which explanation we should prefer. The tracking account is untenable. (2006: 134)

So said, Street thinks that the tracking account is *unbelievable* according to the normative constraints on our best science, according to our most plausible theory of knowledge.

According to the relevant epistemic account, we *should* prefer the theoretical virtues of parsimony, plausibility, clarity, coherence, explanatory power and so on that are characteristic of our *best* scientific or philosophical explanations.

All of this may not be of much concern, were Street agnostic on particular evaluative facts. But of course, her argument (especially with respect to the second horn of the dilemma against the naturalist) repeatedly invokes a number of significant and explicit normative considerations—that is, apparent evaluative *truths* about human knowledge—*without argument*. For instance, Street claims that irrealism “acknowledges the point that we are

self-conscious and reflective creatures [...] *better* than alternative views, by asking what reflective creatures like ourselves *should* conclude when we become conscious of what Kant would call this ‘bidding from the outside’ affecting our judgements” (123, emphasis my own). She continues:

The very fact of our reflectiveness implies that something *must* happen—that something *must* change—when we become conscious of any foreign influence (such as these Darwinian forces) on our evaluative judgments. What that change *should* be is exactly what I am exploring in this paper. (123, emphasis my own)

In ordinary circumstances, there isn’t much wrong with an appeal to normative or prescriptive considerations in the course of an argument. But in the course of an argument against a particular category of objective normative considerations—that is, in the course of an argument *against* objective morality—normative prescriptions of the sort the author ultimately denies (in general!) cannot be taken for granted. As it turns out, Street denies that there are no objective facts about (for example) what counts as knowledge, so her appeal to epistemic principles without careful discussion of their truth isn’t warranted in the obvious, objective sort of way her readers might expect, especially according to Street herself!

In the first case, in response to the first horn of her dilemma for moral objectivists, Street points out that their explanation is “implausible”, in that it does not make sense of moral knowledge. But the absence of moral knowledge does not entail the absence of moral facts. If I do not know that <torture is wrong>, this might be because (1) I *do not* think this **or** (2) this belief is false, **or** (3) I *do not* have good (enough) reason to believe (2). For example, if the belief happens to be true, but I do not have reason to believe it, then I do not know that

torture is wrong. Or, I may have good reasons for the above belief, but I still do not know that torture is wrong if it is false that torture is wrong. Basically, Street concludes from the fact that we do not know whether torture is wrong that it is not objectively true that torture is wrong. But objectivists do not think that they do not know moral facts because there are no such objective facts. If anything, the objectivist would only admit that we do not have good enough reason to believe that torture is wrong.

For example, since Joyce does not conclusively decide questions of moral truth, and focuses on the justification of moral beliefs instead, he cannot know which of our moral beliefs are objectively true. But this is precisely what one would expect from a committed *skeptic* about moral knowledge. On the other hand, a positive account of moral *knowledge* has to provide a particular account of the relation between the first two conditions above that the skeptic might distance himself from. As such, Street's argument against moral *objectivity*—from empirical data alone—may stack the cards against objectivists from the outset!

For now, however, I'll grant Street these assumptions, and describe her argument as follows: there are no objective moral facts because we cannot have knowledge of them. More exactly, we cannot have objective knowledge of moral facts that are not connected at all to the processes that select for adaptive moral habits (including beliefs), and natural moral facts do not explain the processes that do select for adaptive moral habits either (including beliefs). This shows that we do not have moral knowledge if morality is

objective. But Street thinks it is false that we do not have moral knowledge. So Street concludes that moral facts are not objective.¹⁴

1.4: Making Sense of Moral Skepticism(s)

In §1.0, I compared two kind of skepticism. The first argued that some belief lacked justification, while the second ruled out its truth. This is the difference between Sophie noticing that she does not have good reason to believe in Santa, and simply discovering, instead, that there is no Santa. In what follows, I will argue that doubts about morality loosely resemble the first, rather than the second case. In comparing the differences between these two kinds of skepticism, I will claim that the latter is committed to more conclusions than the former. That is, while one kind of skepticism showed that we do not know something, the other goes beyond this doubt to show that something is not real. So, really, the second sort of skepticism is more properly a rejection of some fact(s), than of knowledge.

Recall Sophie changing her mind about Santa Claus. Sophie knew next to *nothing* about how various Christmas events *in fact* work despite all her beliefs about Santa Claus. This sort of pervasive error will become more and more significant in the next chapter. For Sophie's problem is the moral skeptics' problem: despite believing everything everyone tells her

¹⁴ On the face of it, this begs the question against a moral objectivist, and this is precisely what I want to highlight. That is, although an objectivist argues that moral knowledge is possible iff the relevant moral facts are objective and accessible to us, Street seems to assume the conclusion of her argument—that there are no objective moral facts—in her denial of objective moral knowledge as well. In other words, it trivially follows that we cannot have moral knowledge if we assume that there are no moral facts. Without this assumption, the absence or improbability of objective moral knowledge does not entail the irreality of objective moral facts. This question-begging accusation will take up much of Chapter 3 (§3.2).

about Santa—despite learning about the elves, and the reindeer, and the north pole, and the rest—once it turns out that the real world does not contain (for instance) elves and flying reindeer, Sophie is left with the rather serious question of what all these Santa beliefs amount to, after all.

Now, as common sense would have it, these beliefs won't amount to much at all. For when Sophie comes up with an explanation of how Christmas could work without Santa, she has all the reasons she needs to doubt whether Santa is *real*. But Sophie need not have decided what *exactly* to believe instead. After all, there will still be the presents, the red-suited man in the mall, the missing milk and cookies, and a whole lot of other Santa-phenomena left to explain. And there may be a number of alternatives yet to consider. For instance, she might think that there are multiple gift distributors, or very few good children, or very few houses with children, and so on. She might just decide that the Santa in the mall is the real Santa, and the Santa at the north pole is fictional. The point is that, before she has enough evidence against the totality of her beliefs, she may simply admit that very many of her beliefs are false without deciding which ones these are.

I think that the empirical problem in meta-ethics encourages this second sort of doubt. On the one hand, although we cannot help but have moral beliefs, we may have a hard time knowing which are correct (if any). On the other hand, if we find reason to reject moral facts out of hand, we have little reason for any particular moral beliefs. Note that while the first variety of skepticism gives us reason to doubt particular beliefs, the second makes this agnosticism unnecessary or misguided from the start. After all, absent an objective analysis of belief, it hardly makes sense to look for ways to justify our beliefs. Despite our initial

convictions, an anti-objectivist approach already assumes that the world is simply not the sort of world that makes our (default) moral beliefs objectively true or false at all.

In other words, it is in only the first of the two cases above that the search for a further objective explanation—for the moral facts themselves—still make sense. And it is just my contention in what follows that this is exactly the case in which we find ourselves as persons considering moral claims. Despite powerful reasons for doubting a large variety of moral claims, we cannot help but think that some of them might be true.

Furthermore, we often cannot help but think that if they are not true, our mistake can be explained in thoroughly objective terms: we make mistakes about the wrongness of torture, for example, if torture is *in fact* wrong, and we somehow missed this fact. This objectivity is just what makes skepticism such a pressing problem: as long as it is possible that questions of right and wrong make sense—as long as moral questions have right and wrong answers in some absolute sense—it is often possible that we are not right about many of our beliefs at all, although we think we are right about each of them.

Or so the skeptic claims. And I argue, in broad agreement with this claim, that particular worries about justification are difficult to solve in moral epistemology. This does not mean that we ought to conclude from skeptical arguments *of this sort* that morality is not real or that there aren't reasons for moral beliefs. On the contrary, I will argue that objective normative facts generate epistemic puzzles about particular normative facts. These puzzles explain why the metaphysics of morality—the right interpretation of moral facts—remains an open and important question (FitzPatrick 898). This divides me from moral skeptics of the metaphysical variety outlined in the previous section (§1.3). Simply put, I do not think

that we have conclusive reason to doubt the *objectivity* of moral truth. And I certainly do not think an irrealist skepticism along these lines follows from our uncertainty about such truth. Rather, I will ultimately argue that if normative claims are not objective all the way down, then we do not have very good reason to doubt moral beliefs.

So said, most of my thesis takes up and responds to Street's first horn of the dilemma for moral objectivists, and compares this argument to Joyce's moral skepticism. My comparison focuses on what sort of skepticism a moral objectivist—especially a non-naturalist—ought to consider, and why this commitment is required by value objectivism more generally. I will argue that as long as objective moral knowledge is at stake, there is still some point in projects that work out which objective moral facts, if any, justify beliefs.

In particular, in Chapter 2, I am going to use Joyce's skepticism as an example of an empirical problem that the moral objectivist (particularly of the sort that does not assume a relation between moral facts and scientific explanation) must take seriously. In Chapter 3, I am going to discuss whether the first horn of Street's dilemma really presents an empirical problem of the sort Chapter 2 defended. So far, I will not have argued that our moral beliefs are justified or that they are true. I will have merely acknowledged that they are difficult to justify, precisely because we have no way of knowing which of our many beliefs are true. Chapter 3 will point out that although it is the case that our moral beliefs are not warranted even if true, they are not *therefore* false. So Street's conclusion against objective moral beliefs begs the question against objectivism: that is, the non-moral explanation does not show that our natural moral beliefs, *because objective*, are mistaken.

In Chapters 4 and 5, I will argue that objective evaluative facts make the most sense of our doxastic experiences. In Chapter 4, I am going to examine and then argue against Street's value subjectivism at length. Street thinks objective values are incompatible with natural explanations, but Street assumes irrealism in arguing against epistemic objectivity. Street's solution is therefore vulnerable to a dilemma: if epistemic irrealism is true, there is no ultimate reason to believe it, but if true beliefs are not better than false beliefs there is also no reason to believe irrealism, because it is false. So ultimately, since Street's solution doubts ultimate reasons for belief, her solution is incoherent. In Chapter 5, I'll return to the work I started off with, and argue that the justification of moral beliefs, rather than their metaphysics, is primarily at stake in these contemporary moral skepticisms. Skepticism makes sense, that is, for someone who believes in objective normative facts to begin with. A corollary of this argument is that skepticism of these facts is ultimately redundant for those who do not believe that any normative facts really exist independently of us to begin with. That said, I've already sketched the essential arguments of this thesis in this first chapter. The rest of the thesis will situate these arguments in the literature and spell them out with more precision. I am afraid some of these details are not very easy to write or read about, so I have done my best to refer to this chapter throughout to remind the reader of some key points in what follows. My intent so far, at any rate, has been to introduce the analogies and tools that these arguments will employ.

2

Objective moral facts and error theoretic value

2.0: Introduction

The next two chapters compare Joyce's analysis of the empirical problem in meta-ethics (Chapter 2) with the first horn of Street's dilemma for the moral realist (Chapter 3). In §2.1, I clarify the connection between Joyce's (and John Mackie's) moral skepticism and their commitment to moral and normative objectivity. I then introduce and employ a modal mechanism—what I've called the "False World Hypothesis"—to pull apart the *possibility* that there are no moral facts from the error theorists' eventual denial of moral facts (§2.2). I will illustrate this modal mechanism with an analogy before revisiting the biological details that make such "false worlds" a relevant epistemic defeater in the moral case (§2.3). Finally, however, I argue that Joyce's confidence in his evaluative conclusion against moral belief is not obviously compatible with the natural explanation of evaluative beliefs in general, absent a principled distinction between moral and epistemic facts. Without such a distinction, it may be difficult for Joyce to establish cases of evaluative knowledge in epistemic facts as well as moral facts.

2.1 Objective Morality and Errors of Belief

At the start of this thesis, I briefly outlined an objective moral theory. According to such a theory, we describe the world more or less successfully when we make moral pronouncements and assessments. At least, we use these sentences to represent moral phenomena. Such sentences are true or false just in case they represent some moral phenomena more or less correctly. And it is because we think such phenomena are objective, or true independent of us, that we expect them to influence our perceptions and beliefs accordingly.

This analysis of morality is common. Much *default* moral believing *seems* semantically and cognitively objective, and many of us want some objective purchase on morality. When we come across hoodlums burning a cat, at first glance, we see something wrong, and we mean to report what we see when we say that burning cats is immoral.¹⁵ But this is a naïve approach to morality. It simply assumes that morality involves beliefs and that these beliefs

¹⁵ This example is from Gilbert Harman's influential *Ethics and Observation* (from *The Nature of Morality*). Harman points out that there is hardly any good reason to think that moral facts themselves help us explain everyday occurrences, even though we cannot help but think they do (7). Morality, on his view, does very little to explain even the most horrendous or important events. Consider the gang of hoodlums burning a cat. Most of us see something wrong here. But it seems that we would "see" this "wrongness" even if moral facts did not help explain what is actually happening. If we look closely enough, we can break down the event into psychological and, ultimately, physical facts: the cat is burning, it is in pain, and we cat lovers are repulsed and mortified and so forth (perhaps, ultimately, for non-moral reasons). But are there really any moral fact(s) that help us explain why we feel the way we do, or why we "see" something wrong? Is not the exhaustive descriptive detail of our psychology, including the physical circumstances of the event, more than enough? This chapter outlines the error theorists' response to these questions.

are more or less connected to facts, and that we usually believe whatever's true *because of this truth*.

The previous chapter suggested instead that the truth of whatever we happen to believe does not directly explain why we believe what we do. Rather, science provides a powerful explanation of why such *apparent* objective truths are an essential feature of our moral beliefs. To begin with, these explanations make sense of much of our moral experience. After all, moral beliefs *are* powerful and natural, they are influenced by our environment and upbringing, and they are remarkably diverse. While the standard analysis attempts to explain these moral features on its own, there is a lot that science explains without it.¹⁶ Science continually expands our historical horizons: it explains why we are social creatures, and why our behavior, and the beliefs involved in our psychology, are adaptive. More importantly, science does not need to decide difficult moral questions or prove the truth of our beliefs. We often believe the things we do, in short, because of what others believed in the past. Those without the right adaptive habits died off long ago.

Although natural explanations of morality do not directly involve moral facts, they represent our attitudes towards such facts in the standard, objective way that introduced this section. Even if there are no such facts, the scientific account still shows that moral belief is intelligible as fictional discourse about semantically and metaphysically objective properties. If this is right, morality is the stuff of mistaken beliefs—a way of thinking and

¹⁶ The advantage of scientific explanations, so stated, is itself normative: scientific explanations of moral belief exhibit some *evaluative* advantage over the (traditional) philosophical explanations. Note that this seems inconsistent with the objection that science cannot make sense of evaluative facts of just this sort. I will argue, at any rate, that science cannot rule out evaluative facts without presupposing them (in Chapter 4).

talking about the world in terms of objective values. So said, science can explain the objective features of moral discourse and psychology even if our moral beliefs are false.¹⁷

The parallels to Sophie's case from Chapter 1 are instructive. When Sophie talks about Santa, she does not *mean* to pick out a fictional character as such—she is making claims instead about the way the world really is. And so, again, we have an example where someone is committed to some objective facts, and is wrong all along about the world whether she knows it or not. Sophie was wrong to interpret the Santa discourse as ostensibly real, but breaks free from the myth of Santa, in the end, when she comes to describe the (Santa-less) world correctly.

Similarly, Richard Joyce argues that our moral discourse about apparent properties is ultimately fictional because there are no such properties in the real world. In *The Myth of Morality*, Joyce's interpretation of moral discourse preserves the realist commitment to objectivity, but also goes hand in hand with the worry that we systematically get *all* of our moral beliefs wrong (2001: 9). In this way, Joyce agrees with many of us that our moral beliefs are almost always about what is objectively true or false, but does not think that they are mostly objectively true rather than objectively false. As will become clear in my discussion of his second book, Joyce is just not convinced that our moral beliefs are

¹⁷ Some disagree. Robert Nozick, for example, argues that the facts objective morality describes must be explained by their truth rather than science, since "the capacity to recognize ethical truths unveils surprising structures, convolutions, refinements, modulations and asymmetries. Reflective people [...] recognize these features and give them some weight, features having little apparent connection with inclusive fitness or actual schedules of reinforcement and shaping. This is evidence of a realm with its own objective standing and characteristics" (347). Nozick concludes "it need no longer be true of our pattern of moral behaviour (and beliefs) that no one tried to bring it about; it need no longer be true that the correct explanation of why it exists and continues is a blind explanation. Since what arises via a blind process need not itself be blind, this can lead to new things that cannot themselves be explained by an invisible hand explanation (347).

justified. This is certainly moral skepticism. But it is also distinct from thinking that our moral beliefs are not objectively truth-apt or that they are truth-apt, but not objectively so.

Following John Mackie, Joyce argues that a moral error theory makes more sense of our moral experiences and statements than the non-cognitivist or irrealist interpretations of moral discourse outlined in §1.1 (2001: 9). Error theory denies that beliefs in objective moral facts are successful, but remains, for all that, an analysis of morality as an objective enterprise. Error theory shares an objective semantic analysis of morality with the realist, but rules out the *reality* of the objective properties this analysis represents. Of course, Mackie frequently dismisses objective values, but this way of talking is rather unfortunate: it isn't that objective values are the wrong commitments for our morality, it is just that morality is inherently objective, and that there are no particular objective values to be had.

The most familiar of Mackie's arguments to this end focuses on the *queer* features of these values—they are "facts" that *must* motivate us, but can't be "seen", or reliably "verified" if they are real (39). So by the realist's own admission, objective values do not share the standard properties we ordinarily expect of facts; as Mackie puts it, "they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe" (38). Mackie has a legitimate (albeit obscure) worry here. I do not know what to make of it, and I am not entirely sure that Mackie's worry is entirely independent of the arguments Joyce articulates against the explanatory relevance of purported moral facts. Perhaps the concern that motivates both Mackie and Joyce is just that "objective values [...] are not the kind of stuff likely to appear on the list of what our best science says or is likely to say exists" (Enoch 135). I will not decide this question, because I think what can be said

against the scientific argument for error theory applies also to Mackie's objection from queerness (in short, queerness also rules out other facts—more facts than the error theorist initially bargained for).

So error theory focuses on explaining our moral discourse—on demystifying morality—as it were, without taking on the intractable task of making moral facts, mythical as they are, any less mysterious. This account agrees that there are no objective values, but it argues that this follows from the nature of morality and moral discourse itself, rather than the way the world is like. For example, rather than describing an objective reality, some subjective theories argue that we create a moral reality for the purpose of evaluating actions, decisions and so on. But for the error theorist, as for the realist, this is not what we take ourselves to do in making moral judgments at all. So error theory is careful not to confuse the rejection of objective moral facts with the rejection of objective morality altogether.

For example, Mackie is careful to append his opening statement to his *Ethics: Inventing Right and Wrong* (that there are no objective (moral) values whatsoever), with the following qualification:

If all our statements were such subjective reports, it would follow that, at least so far as we are aware, there are no objective moral values. If we were aware of them, we would say something about them. In this sense this sort of subjectivism entails moral skepticism. But the converse entailment does not hold. The denial that there are objective values does not commit one to any particular view about what moral statements mean, and certainly not to the view that they are equivalent to subjective reports. (18)

Unfortunately, Mackie announces this skepticism without talking in much detail about his semantic or ontological commitments in meta-ethics. It is no wonder that error theorists

thereafter are called “anti-realists” together with non-cognitivists, irrealists, subjectivists, and so forth. But putting Joyce and Mackie in such company is misleading; if anything, they are in about as much agreement with straightforward cognitivists and moral objectivists as they are with anti-objectivists and non-cognitivists outlined in §1.1. The moral objectivists provide an accurate picture of morality as a human practice involving more or less successful assertions and beliefs, and although this picture itself is the right one, the particulars it involves are riddled with fictions and mistakes.

With this caveat, I will take these connected claims to be definitive of error theory, at least for the purposes of this thesis:

Ordinary moral judgements include a claim to objectivity, an assumption that there are objective values in just the sense in which I am concerned to deny. And I do not think it is going too far to say that this assumption has been incorporated in the basis, conventional, meanings of moral terms. Any analysis of the meanings of moral terms which omits this claim to objective, intrinsic, prescriptivity is to that extent incomplete; and this is true of any non-cognitive analysis, any naturalist one, and any combination of the two [...] the denial of objective values will have to be put forward not as the result of an analytic approach, but as an ‘error theory’, a theory that although most people in making moral judgments implicitly claim, among other things, to be pointing at something objectively prescriptive, these claims are all false. It is this that makes the name ‘moral skepticism’ appropriate. (Mackie 35)

2.2: Worldly Facts and Fictional Worlds

Mackie and Joyce both conclude that we believe moral claims in error. More exactly, our moral judgments are mistakes because they involve beliefs about the world—they “[point] at something objectively prescriptive”—even if the world does not contain the properties that would make these beliefs true (Mackie 35). Furthermore, this possibility is at least

consistent some common phenomena: we talk about moral properties all the time, and we refer to all sorts of moral facts when condemning our cat-burning hoodlums, even if we are at a loss to explain what in the world makes all these moral judgments true.

It might be instructive, in understanding the moral error theoretic account of such phenomena, to consider the theory of phlogiston, according to Joyce:

The chemists Stahl, Priestly, et. al. [...] were able to ostensibly pick out paradigm examples of phlogiston: they could point to any flame and say “There is the phlogiston escaping!” and for all that they were failing to state truths, for there wasn’t any phlogiston. [...] users of the term (considered collectively) thought and said certain things about phlogiston such as “It is that stuff stored in bodies,” “It is that stuff that is released during combustion” and “Soot is made up almost entirely of it,” and these concomitant statements are false. (2006, 2-3)

In the above, Joyce provides us with an example of some discourse concerning some properties that turned out to be fictions—that is, a false description of reality. But before Lavoisier’s theory of combustion replaced all talk of phlogiston with talk of oxygen, chemists proposed the properties of phlogiston as *matters of fact* in exactly the same way Lavoisier inferred the existence of oxygen from very real phenomena. This is why we think Stahl and Priestly were wrong about phlogiston—their concept was useful and meaningful and superficially consistent with the way the world looked, but it was just the sort of thing that needed to *be there* in order to make all this talk of it *true*.

Truth like this does not depend on the way we think or talk about the world.¹⁸ For Stahl, for Priestly, and for Lavoisier, the truth of what they said ultimately and exclusively depended

¹⁸ As Shafer-Landau puts it: “For a claim to be objectively true, it must be true independently of what anyone, anywhere, happens to think of it. Mathematical truths are like this, as are truths in such areas as astronomy, microbiology, and chemistry. We don’t make them up. The true claims within these disciplines are not true because we think they are; we think they are because they really are. We may have invented the *vocabulary*,

only on what the world, *in fact*, is like. And since their beliefs concerned the real world and the way it is independent of what they thought of it, they turned out to be right or wrong with respect to some attitude-independent or objective truth of the matter. Certainly, Stahl and Priestly were convinced that *they* referred to an objective property of the world, and both would propose that theirs is a theory of the world itself. Of course, it is easy enough for us *now* to talk about phlogiston as a curious, fictional entity—saying things like “Priestly believed in phlogiston” or “Phlogiston doesn’t exist”—but were we to apply these words quite literally, saying something like “*a* is phlogiston” while pointing at a fire, we would be just as committed to some objective fact, and just as mistaken, as Stahl and Priestly were (Joyce 2001: 5).

Joyce argues that we are much like Stahl and Priestly in our moral habits: we state our (false) moral beliefs and judgments as *matters of fact*. Similarly, Mackie claims that the objective moral theories of Plato, Aristotle, Kant and Sidgwick are outdated and incorrect (30-31). These philosophers certainly *thought* they had provided some theory of what an objective moral fact would consist in, but these were false as descriptions of the world. At least, these philosophers were also like Stahl and Priestly, argues Joyce, in that their theories are now incompatible with an alternative explanation about the way the world *is*—our current science is much like Lavoisier’s, since it exposes the objects of our more traditional beliefs for the useful fictions that they were, or still are.

If contemporary moral philosophy is as mistaken as 18th century chemistry, the way we think about the world is different from the way the world is. Even if the world was more or

but we didn’t invent the *fact* that (say) the earth is in a solar system, that genetic information is carried by DNA, that one oxygen and two hydrogen molecules, will, at certain temperatures, bind to form a liquid” (12).

less the way it looked—even if our moral beliefs were true, we have yet to rule out the error theorists’ hypothetical amoral world. For example, Mackie argues, in agreement with R.M. Hare, that a world with objective values might *look* quite like a world without them. Quoting Hare’s claim that between “one world into whose fabric values are objectively built; and [...] another in which those values have been annihilated” there is no “difference between the states of affairs in these two worlds”, Mackie writes “it does not follow, and it is not true, that there is no difference whatever between these two worlds. In the one there is something that backs up and validates some of the subjective concern which people have for things, in the other there is not” (21-22). Our task, then, is to show that the way the world looks to us is the way it is. And if a very different world could look quite the same to us, this task is significant. Ours could just as easily be the world with no moral facts at all, even if we were convinced we had discovered the moral facts after all.

For example, take some fact Φ . Mackie and Joyce propose a skeptical problem for beliefs in some fact Φ just in case it is possible that our psychology in a world with Φ would not change in a world without Φ . Let the world with Φ be the True World (Wt) and call the world without Φ the False World (Wf). Say that Rene, a skeptic, lives in Wt . And say Wt contains a candle, while Wf simulates a phenomenological replica of the candle, absent an actual candle. According to a standard skeptical argument (Descartes 115), if Rene can imagine Wf , and if he has no conclusive reasons to rule out his being in Wf , he may well doubt whether he really is in Wt . Of course, Rene’s being in Wt is determinate: it either is or isn’t the case that the candle is an illusion. But despite the objective truth of the matter, Rene doesn’t know that the candle is real unless the relevant justification conditions—for instance, discovering certainties through introspection—rule out his being in Wf . Some

(very strict) epistemic accounts state that if Rene's belief is true all along, he only knows rather than merely believes this conclusion, *once alternative (non-truthful) explanations of that belief are ruled out.*

Now, I don't want to argue that the False World Hypothesis counts against belief in *any* circumstances. For instance, the above circumstance posits an illusion that many of us would object to: it is not a very serious possibility, or an intuitively realistic possibility for many of us, anyway. But it is a possibility nonetheless. And some skeptics will argue that these sorts of possibilities always make knowledge about anything impossible. I am not concerned to argue with that sort of skeptic here. The circumstance with respect to morality, according to Joyce, is a lot more like our beliefs in phlogiston, or even beliefs in elves and flying reindeer, than it is like everyday perceptual beliefs. This is because many of us find objective moral facts hard to believe in quite the same way that, at one point or another, one or more of us have come to think that our world is, in fact, without phlogiston, or without elves and flying reindeer. As long as these possibilities exist, we could be just as convinced of our moral theories as 18th century chemists were of theirs, and likewise all the worse for our conviction.

At some point, in other words, we changed our minds about whether our world contains some fact Φ . We thought, in the case of phlogiston, that Φ (phlogiston) was all around us. And now we are quite sure that we were in the False World (the world without phlogiston) all along. So say the facts Φ that Joyce and Mackie are concerned to refute are the particular objective facts that would make at least some of our moral beliefs true. Mackie and Joyce argue that Φ is absent from this world—at the very least, they argue that there is very little

evidence that we live in *Wt* rather than *Wf*. Further, Joyce wants to make the False World Hypothesis about morality just as plausible as the False World Hypothesis about phlogiston. In summary, Mackie and Joyce think that a False World Hypothesis about morality counts against moral belief in our particular circumstances.

2.3 Doxastic Error and Epistemic Defeat

A False World Hypothesis operates as an epistemic defeater for our beliefs in some fact(s) Φ . For instance, Lavoisier's theory of combustion made the False World Hypothesis about phlogiston relevant, and his alternative, more accurate conception of the world defeated our beliefs in phlogiston. Likewise, once Sophie discovers that a world without Santa might very well look just like the world with (stealthy) flying reindeers and the rest, the False World Hypothesis also operates as an epistemic defeater for beliefs in Santa. Our epistemic defeaters for beliefs in astrology, or homeopathy, might well work like this as well.

Joyce uses another analogy to illustrate such cases:

Suppose that there were a pill that makes you believe that Napoleon won Waterloo, and another one that makes you believe that he lost. Suppose also that there were an antidote that can be taken for either pill. Now imagine that you are proceeding through life happily believing that Napoleon lost Waterloo (as, indeed, you are), and then you discover that at some point in your past someone slipped you a "Napoleon lost Waterloo" belief pill. It is not a matter of your learning of the existence (of) such pills and having no way of knowing whether you have ever taken one; rather, we are imagining that you somehow discover beyond any shred of doubt that your belief is the product of such a pill. (179, 2006)

In this case, Joyce continues, the knowledge that someone slipped you a belief pill should defeat your belief <Napoleon lost Waterloo>. Notice that your skepticism has nothing much

to do with whether Napoleon really won or lost Waterloo. After all, there could be the two worlds: the world where Napoleon won (Ww), and the world where Napoleon lost (Wl). And you could be in the losing world, believing that Napoleon lost, and still not know that Napoleon lost. To be exact, you would be in Wl , thinking that Napoleon lost, but also thinking that Ww is likely to be actual, especially if you discover you've been arbitrarily influenced (by, in this case, a belief pill) to think Napoleon lost. So you had better take the antidote, get the belief out of your system, and look for the truth.

Nonetheless, none of this shows that the pill itself is *evidence* that Napoleon won Waterloo nor evidence that Napoleon lost Waterloo; it is, rather, merely the *absence* of sufficient evidence one way or another. Now, suppose that you found yourself in the world where Napoleon won Waterloo, and also had good evidence that Napoleon won Waterloo. In this case, you may know that Napoleon won Waterloo. At the very least, the evidence would certainly rule out our knowing that Napoleon *lost* Waterloo. So epistemic defeaters make it difficult to justify beliefs in some facts Φ *only* once we consider whether your etiology is not sensitive to a world without facts Φ . In particular, Joyce illustrates an example in which your belief in some fact Φ is not justified, because the fact does nothing to explain your belief. On the other hand, the actual status of Φ —the truth of your belief—does not really matter to our epistemic evaluations, for the etiology of your belief defeats true beliefs just because the truth does not explain or cause the true beliefs.

In this way, there's a clear difference between not being able to tell Ww and Wl apart from our own perspective, and simply *knowing* that we are in Ww . In both cases, we cannot know that we are in Wl . And in both cases, we have reasons to doubt that some facts Φ

obtain. But the reasons for our doubt are different between the two cases. In the former, our discovery of some suspect influence on our belief <Napoleon lost Waterloo> gives us reason to suspend beliefs about Napoleon altogether, and in the latter, the evidence that Napoleon won Waterloo gave us even *stronger* reason to disbelieve *in particular* that Napoleon had lost Waterloo. But it is only the latter case that one of these worlds was ruled out. So Joyce's analogy is not concerned with whether these beliefs are ultimately correct, and focuses instead on how the pill works. And it has everything to do with what, *in fact*, counts as knowledge.

2.4: Objectivity and Normativity

I had stated in §2.1 that moral facts do not exist, according to Joyce. And I also admitted in section §2.3 that Joyce and Mackie think some objective moral facts Φ are not "real". But I have cautioned throughout this chapter *against* reading an outright refutation of moral objectivity into Joyce's moral skepticism. Indeed, since writing *The Myth of Morality*, Joyce's robust ontological claims have become considerably more modest meta-ethical doubts. The analogy above illustrates this more considered stance: we are not Lavoisier with theory of combustion in hand, we are drugged believers without any real access to the facts we believe in. This is why Joyce has recently and explicitly cautioned *against* arguing that these facts do not exist (2013: 354). For him, the possibility of our having been "drugged" with respect to all of our beliefs in morality makes the reality of some objective moral facts Φ hopelessly intractable (2013: 355).

To see this, let's substitute moral facts for the relevant facts about Napoleon in our analogy. Call the world with some objective moral facts Φ the True World (Wt), and the world

without moral facts Φ the False World (*Wf*). And call Joyce's hypothesis that there are no objective moral facts the False World Hypothesis. In the False World, our belief <we should not kill children> isn't paired with any fact that it is wrong to kill children. For instance, even though a False World contains the natural properties that the True World contains, it seems that our beliefs about morality would also be the same, even absent the particular objective moral facts that make our moral beliefs true in the True World. This is because there are no False Worlds where we have true beliefs in an objective fact that it is right to kill our children. Worlds where we believe this will be worlds without us. This is just to say that our biological world, rather than the moral facts, explains why we believe we should care for our children.

In this way, Joyce's problem is really just that two competing hypotheses explain the same set of beliefs, where " \rightarrow " symbolizes some causal relation:

Hypothesis A: non-moral genealogy \rightarrow moral beliefs (False World)

Hypothesis B: (objective) moral facts \rightarrow moral beliefs (True World)

Of course, there is the third hypothesis that the moral facts explain the non-moral genealogy that in turn explains the moral beliefs. Call this the True World*.

Hypothesis C: (objective) moral facts \rightarrow non-moral genealogy \rightarrow moral beliefs (True World*)

In the most difficult variation of this problem, there are perhaps an infinite number of possible worlds with an infinite combination of possibly true moral facts (Street 2006: 122).¹⁹ But the possibilities of our moral beliefs varying with the facts across all worlds is

¹⁹ See also Justin Clarke-Doane (2014).

constrained by our biological history. If this seems right, the task of guaranteeing that the True World is actual is virtually insurmountable. Worse still, in an actual True World, there is no obvious moral explanation of our capacity to make moral judgments and think in moral terms. So, we cannot be sure we are in the True World, even if it is actual.

The conclusion of the False World Hypothesis in either case is simply that *S*' belief is unjustified. In one case, *S* believes that <Napoleon lost Waterloo>, in the other *S* believes that <I should never kill my children>. Call these beliefs *p*. In the first case, the pill provides a relevant explanation for this belief, in the latter, it is our biological history. Call this explanation *E*. This yields us the following causal premise:

Causal premise: *S*'s belief that *p* is explained by *E*.

And now, we add in the False World Hypothesis, which yields an epistemic premise:

Epistemic premise: *E* is insensitive to *p*.

Finally, we have Joyce's (skeptical) conclusion:

∴ *S*'s belief that *p* is not justified.

Or, in other words, *S* does not know that *p*, and has no reason (yet) to believe *p*.²⁰

In this way, Joyce proposes normative considerations in favour of good beliefs that override the default "natural" habits of believing something that's false (2001: 171). Here's a quick example: it's quite natural, at least before some scientific education, to think that the world is flat. It just *looks* that way to us. But of course, we have plenty of reason to believe that the world is spherical instead. So in these and other cases, considerations having to do with truth can override the kinds of adaptive or naïve beliefs we started off

²⁰ See Eric Wielenberg (2010) and Guy Kahane (2010) for similar presentations of Joyce's argument.

with in favour of *better* beliefs, however unnatural they might seem at first blush.²¹ We might say we humans typically believe what we do, all things considered, by *normative* necessity.

I take it that these are cases in which beliefs are ultimately responsive to facts about what to believe, as well as apparent non-normative facts, even if belief is (generally) adaptive. In the above argument, although both the causal and the epistemic premise are descriptive, Joyce has no trouble jumping to a prescriptive conclusion from these premises. In other words, Joyce gives us facts about our biological history, and together with an (implicit) modal analysis, takes these facts to convince us to change our minds about morality. But this is an inexact picture. What's really going on is that we have a normative fact, in the end, about *S*. That is, in short: "*S* does not know that *p*", where *p* might stand in for the semantic or metaphysical content of our moral beliefs.

Notice, however, that even if this statement is objectively true, its truth does not seem to matter to explanations of our belief that it is true. We might believe that "*S* does not know that *p*" even if were false. In other words, as Chapter 4 will argue, the possibility of a False World with respect to the truth of these statements also create problems for our knowledge of such statements. You might recall that in §1.1, I mentioned the sort of evaluative standards that make it the case that our belief in some fact, together with this fact, is knowledge of that fact. These standards resemble the sort of evaluative standards

²¹ Joyce does acknowledge, however, that these "better" beliefs are far from automatic—for instance, we have a habit of perceiving and treating people who look like or live near or with us differently than others, sometimes despite "knowing better" (see Street 2006). Nevertheless, even Joyce is confident that "with the appropriate education, role models, and cultural climate, all traces of the trait [our xenophobic dispositions] can be abolished" (2001: 171).

that make it the case that some action is more or less right. And these standards are, of course, just the evaluative sort of objective moral facts Joyce has been discussing and doubting all along.

This is why I am now pointing out that Joyce makes an *objective* epistemic claim when he concludes that we do not have knowledge of some objective moral facts (2013: 353). After all, we think that some sorts of beliefs are more or less right, and we mean this also in an evaluative sense: we mean to say that some beliefs are *better*, or more warranted, than others. But of course, our habits and practices of believing is itself also biologically adaptive. There are real world consequences between, say, believing that facts do or don't matter. This is just what a biological explanation will tell us, irrespective of which of our particular beliefs are true.

Certainly, this is worry is a hard pill to swallow, but it's no less relevant for all that. It's not clear that evaluative self-referential beliefs about what to believe *aren't* adaptive. We would think truth matters, after all, whether or not there were such an objective epistemic fact at all. We might wonder, then, whether the False World Hypothesis is relevant with respect to some objective epistemic facts Φ . I will not decide this question now. I will argue later, in Chapter 5, that Joyce is right to conclude that there are objective epistemic facts of the sort above. Briefly, this is because these are the sort of facts that hold in all worlds if they hold in any one world. And we have every reason to believe that they hold in this world, if we have any reason to believe anything at all.

But for similar reasons, I am not so sure that Joyce is right to disbelieve the evaluative facts that would make objective morality possible. This does not mean that we know *which* of

our moral beliefs are true. I am just pointing out that they *can* be true, if there really are evaluative facts of the sort Joyce is ultimately committed to. But even on this more modest conclusion, it is not clear that Joyce can know which of our normative beliefs are true, even in epistemology, if he does not know which of our normative beliefs are true in the case of morality. Our normative beliefs are in tension with our meta-epistemic beliefs, and Joyce does not explain why our beliefs in evaluative facts of the latter sort are warranted, while the former aren't in the case of morality. But that will be a topic for another section (§5.1).

I have argued in this chapter that Joyce doubts whether we can ever know much about whether some objective moral facts obtain. That is, he argues that we cannot know what our *moral* world really looks like, ontologically, since the possibilities are overdetermined by our biological reality. But I want to remind the reader that this conclusion is distinct from arguing that moral facts are not objective. At any rate, Joyce and Mackie don't think subjective or non-cognitivist moralities fully capture the formal, objectively *normative* characteristics of morality and epistemology. On the contrary, I suspect that both would concede that normativity is an objective domain, and that there are some facts (primarily objective epistemic facts) that are actual.²²

²² Again, Joyce himself argues explicitly for this reading in 2013, 2014, and 2015. In 2014, Joyce writes: "I don't think many meta-ethical theories deny the existence of normative truths. Not even the typical error theorist denies them. Certainly, it is possible to be an error theorist about normativity across the board, but I don't know anyone who endorses the view [...] people like Mackie (1977) and myself (2001)—think that moral normativity is a defective concept, but happily allow that other normative frameworks are trouble free" (844). In 2015, Joyce puts the same point as follows: "one could be a moral error theorist *by implication* [...] because one endorses an error theory about all normative phenomena—typically the moral error theorist thinks that there is something *especially* problematic about morality, and does not harbour the same doubts about normativity in general" ("Moral Anti-Realism", Joyce's emphasis).

This analysis supplies us with some useful tools for the next chapter. Briefly, the next chapter contrasts Street's approach to the empirical problem with Joyce's. Street disagrees with Joyce, and attacks the "standard analysis" of realism itself. According to Street, moral facts are not objective. For Street, the False World Hypothesis illustrates the impossibility of objective moral knowledge, and it also rules out objective moral truth in all possible worlds. This is, I think, the metaphysical conclusion Joyce is careful to avoid. It is analogous to deciding whether Napoleon won Waterloo without justification or evidence for beliefs one way or another (see §2.3). For on Street's subjective solution to the empirical problem, we ought to choose the truth of our moral beliefs over the commitment to objective truth. We do not have moral knowledge if morality is objective, so it is a mistake to believe in moral objectivity. So, according to Street, our moral beliefs are sometimes true, *because* they are not objective.

3

Subjective moral facts and value irrealism

3.0: Introduction

In this chapter, I focus on Sharon Street's irrealist alternative to moral skepticism. I argue that skeptical arguments do not decide the nature of moral truth. Street's conclusion that morality is not objective is inclusive of, but ultimately distinct from, Joyce's argument that we do not have moral knowledge if morality is objective. The conclusion that moral knowledge is possible requires a positive argument for moral truth, objective or otherwise, rather than doubts about their justification on one such account.

In §3.1, I briefly rehearse Street's case for an irrealist conclusion, I then detail her dilemma for the realist in §3.2. As I understand the dilemma, the objective approach of the realist either results in doubting moral beliefs, or revising our current scientific beliefs. Street argues that the realist should reject objective morality instead, since it is better to believe irrealism. But the conclusion that morality is not objective does not follow from the skeptical argument so stated. Simply put, it is not clear that we *should* believe that such truth is not objective. In §3.3, I compare Street's argument with Joyce's from §2.4. In §3.4, I conclude this comparative task: although moral knowledge is not possible if morality is objective, Street's dilemma for the realist cannot show that morality is not objective.

3.1: Subjective Morality and True Belief

Earlier in this thesis (§1.2, §1.3), I understood both Richard Joyce and Sharon Street to agree that science explains the human capacity, if not the content, of believing all things moral. According to this explanation, these beliefs were indirectly adaptive: for instance, parents who believe that it is wrong to kill their children will, *ceteris paribus*, have more surviving children than those that believe otherwise. But when I tried to make sense of moral skepticisms at the end of Chapter 1, I suggested that Sharon Street defends a much stronger conclusion than Joyce's. In particular, although Street and Joyce agree that objective moral facts are irrelevant to the scientific explanations of moral belief, I concluded that only Street argued against the reality of such facts in general.

Street argues quite explicitly for this metaphysical conclusion. For Street, moral facts cannot be objective. If moral facts are objective, they therefore have nothing to do with the biological processes that select for moral beliefs. And if such facts are not part of scientific explanations, we do not know anything at all about moral truth. But if moral facts are part of the scientific explanation, then our current science is radically incorrect. On other hand, an irrealist theory of moral facts does not require this revision of our science. As well, such theories make moral knowledge possible. On this irrealist view, our scientific and moral explanations are both correct. Since Street thinks that we are right about many of our particular scientific and moral beliefs after all, she concludes that moral facts are not objective.

More exactly, Street's solution *conclusively* abandons an objective moral theory in favor of a subjective moral theory constructed from the beliefs we already have. For Street, we ought

only to have further beliefs if and when they cohere, constructively, with our basic evaluative attitudes (2006: 153). This is a version of value irrealism, since a moral “truth” that depends on our attitudes and beliefs is not ultimately the objective truth of the cognitive realism or error theory examined earlier. In Street’s words:

Anti-realism explains the overlap [between evaluative truths and our contingent evaluative judgments] not with any scientific hypothesis such as the tracking account, but rather with the meta-ethical hypothesis that value is something that arises as a function of the evaluative attitudes of valuing creatures -- attitudes the content of which happened to be shaped by natural selection. The breaking of our bones is bad, on other words, and we’re well aware of this. But the explanation is not that it is true independently of our attitudes that the breaking of our bones is bad and we were selected to be able to notice this; the explanation is rather that we were selected to take the breaking of our bones to be bad, and this evaluative judgment withstands scrutiny from the standpoint of our other evaluative judgements (to speak, for example, in the voice of the constructive antirealist). (2006: 154)

On the other hand, Street argues that the question(s) of whether our evaluative beliefs are *ultimately* true or false independent of us and our evaluative attitudes is a unique difficulty for moral objectivists. That is, the irrealist insists that the epistemological burden of proof falls squarely on the moral objectivists’ shoulders, given their metaphysical commitments to “real” truth:

Ultimately, the fact that there are any good scientific explanations of our evaluative judgments is a problem for the realist about value. It is a problem because realism must either view the causes described by these explanations as distorting, choosing the path that leads to normative skepticism or the claim of an incredible coincidence, or else it must enter into the game of scientific explanation, claiming that the truths it posits actually play a role in the explanation in question. The problem with this latter option, in turn, is that they don’t. The best causal accounts of our evaluative judgments, whether Darwinian or otherwise, make no reference to the realist’s independent evaluative truths. (2006: 155)

This metaphysical conclusion stands in stark contrast with Joyce's more modest moral agnosticism. Whereas Joyce argued for a purely epistemological variety of skepticism, Street is establishing what I will call (following Ramon Das) a *metaphysical* skepticism, for want of a better term (Das 422). This metaphysical skepticism is riddled with dialectical difficulties that Joyce avoids. First, Street tries to argue that the skeptical argument can work just as well without objective moral truth. But rather than defeating objectivists with their own resources, the analogous *subjective* skeptical conclusion follows trivially from its premises: the epistemic premise shows that the scientific explanation is insensitive to something that, *by hypothesis*, simply does not exist! (Recall, in contrast, that Joyce does not conclude that our moral beliefs are not justified because morality is not in fact real). From this, the objectivist is supposed to admit metaphysical defeat for thoroughly normative reasons—it is “bad” science to look for something that the explanation doesn't need. We should stop hoping for objective moral facts for the same reasons that hunting witches or looking for phlogiston in this day and age will waste valuable time and epistemological resources. “Better” that scientists explain morality without worrying about such objective moral truth.

Street explicitly employs these evaluative epistemic *norms* to convince objectivists of her argument. Objectivists, valuing the approaches they do, take norms seriously, often as categorical imperatives. Paradoxically, however, despite taking a page out of the objectivists' book, Street simply isn't in agreement with *evaluative* objectivists: at the end of the day, it turns out that if there are epistemic norms, they are hypothetical imperatives constructed from our contingent evaluative attitudes (Street 2011). But her normative conclusion requires an even more radical revision of belief than Joyce's does—again, the

appropriate epistemic response to the scientific explanations is an extreme variety of moral disbelief. Indeed, it is because this conclusion is so antithetical to our naïve moral beliefs and values in general that Street constructs a viable, robust alternative to objectivism. Street argues that we should prefer this conclusion to moral disbelief. Rather than remain agnostic, the evaluative force of the skeptical objection *should* push us to embrace our beliefs and values more than ever as our own.

3.2: Irrealism and Moral Knowledge

In what follows, I will argue that objectivists need not accept Street's conclusion on either horn of her dilemma. On either horn, perhaps, at least some of our moral or scientific beliefs are not justified. But each horn entails, *at most*, that we cannot know whether or not our moral beliefs are true on an objective account. And this is distinct from concluding that our moral beliefs are false, or subjective. Instead, the conclusion that morality is not objective requires a *metaphysical* argument against the objectivity of both moral and epistemic realism.

Street argues that an objective moral theory must either give up the truth of moral explanations or the truth of our current natural explanations. The realist must choose between these two horns of the dilemma. On the first, Street argues, the realist or objectivist must admit that moral facts cannot cause true moral beliefs, and therefore are not likely to correspond to such beliefs. The alternative is establishing a causal account of the relation between moral facts and true, adaptive beliefs. On the causal account, beliefs overlap with their phenomenal objects in a relevantly truthful way, whereby a belief is

adaptive iff it is true (Street 2006: 126). On this view, the moral truth might play a causal role in ensuring that those with true moral beliefs survive, while the ones without, do not.

On either of these views, scientific theories explain our moral beliefs, since these beliefs may play a causal role in either preventing or promoting anyone's reproductive success (Street 2006: 128). In either case, some of our common-sense moral or scientific beliefs we currently hold are unlikely to be true—there is a distinct possibility, in fact, that at least some of them are false. We might formalize these two horns of Street's dilemma as two competing hypotheses:

True World^o: moral facts moral beliefs

True World^{*}: moral facts → moral beliefs (see §2.4)

Where “→” symbolizes causation and “...” represents the non-causal relation—or, more simply, the absence of a relation—between moral facts and moral beliefs.²³ To summarize §1.3, since any number of moral hypotheses is consistent with the scientific explanation of moral belief, establishing objective moral knowledge involves explaining which sort of “True World” is actual.

In this way, Street challenges objectivists to demonstrate that moral beliefs in either True World^o or True World^{*} are true. This dichotomy between whether our beliefs do or do not track objective truth represents a dilemma in which either option is plagued with epistemic

²³ This is the thesis corresponding to Street's “Bermuda analogy” from §1.3. According to this thesis: “every now and then, Darwinian pressures might have happened to push us toward accepting and evaluative judgment that accords with one of the realist's independent evaluative truths. But this would be purely a matter of chance, since by hypothesis there is no relation between the forces at work and the ‘destination’ in question, namely evaluative truth” (Street 2006: 121)

difficulties for moral objectivists. In True World^o, a non-causal theory of moral truth leaves an objectivist with the huge task of finding some true beliefs in the face of biological factors that cause adaptive beliefs, whether or not they are true (Street 2006: 122). In True World*, on the other hand, an objectivist's attempt to explain these adaptive beliefs in normative "truthful" terms either competes directly with the biological explanations of belief or, worse, revises and complicates such explanations (Street 2006: 129).²⁴ Simply put, objective moral explanation(s) in either world do not have any advantages in comparison with *just* a scientific explanation in the False World (Hypothesis A from §2.4). Such explanations, once sensitive to "moral truth" in a relevant scientific sense, are substantially less clear and parsimonious as a result.²⁵ If objectivists opt for an uneasy union of veridical and scientific explanations (True World*), their theory does not survive

²⁴ One problem for this sort of approach is the (admittedly, superficial) discrepancies between common-sense moral explanations and our physical explanations: we cite reasons, or desires, or the like, in the former, but not the latter case. When the hoodlums kill the cat, we cite their moral depravity as an explanation of the cat's death, while the physical account of how the cat *in fact* died (generally, gasoline burns, etc.) leaves this out. Street's attitude on this point echoes Dworkin's when he writes: "If someone thought that his moral beliefs were directly caused by moral particles, he would regard that as a vindication of those beliefs, as we regard the discovery that our ordinary observations are caused by what we claim to observe. [But] if he is asked for argument or explanation [of his beliefs] he will offer, if anything, very different propositions" (105).

²⁵ Of course, I do not want to dismiss these objective explanations out of hand, and Street usually tries not to prejudice the discussion in favour of non-moral scientific explanations either. As David Copp points out (much like Sturgeon had against Harman (237)), "if the moral facts are "reducible" to non-moral facts of the kind that are invoked in the Darwinian hypothesis, there is no loss of parsimony in supposing that there are moral facts" (2008: 190). Copp's "society-centred" naturalistic reduction of moral facts to non-moral relational facts holding between the properties of moral codes and a given society is one such reduction available to a moral objectivist. As Copp puts it: "a basic moral proposition is true only if a corresponding moral standard is included in or implied by the moral code the currency of which in the relevant society would enable the society better to serve its basic needs than would the currency of other sets of norms and better than would be the case if no sets of norms had currency in the society" (2008: 200). Copp claims that our doxastic processes under natural selection may have "quasi-tracked" just these sorts of relational facts, especially if there is a biological explanation of our tendency to live in cooperative groups (2008: 202).

scientific scrutiny. On the other hand, as Joyce also shows, the stance-independence of True World^o results in epistemological defeat.

This epistemological defeat, thinks Street, *ought* then to move the objectivist to metaphysical surrender. In this case, Street argues that objectivists *ought* to admit that moral subjectivism is the correct approach to the empirical problem; in particular, moral objectivists do not have enough evidence to justify their theories, and so they do not know that objectivism is true. So the objectivist can never have objective *evaluative* knowledge (Kahane 116). Again, since there is no obvious way of knowing *what* moral truth amounts to, there is no further reason to assume this truth. So Street assumes the negation of this truth: namely, *Wf* in which non-moral genealogy → moral beliefs → moral facts.

Street's initial point is just that so long as objectivists cannot rule out their being in the False World, their beliefs are not objectively true. So Street rejects Joyce's "True World" (of normative value) in quite the same way Joyce rejects our *knowledge* of the "True World" (of moral value). But I do not think Street can conclude that we are in a False World, even on the first horn of her dilemma. After all, Joyce did not conclude that we are, *in fact*, in a False World. Instead, Joyce argued that the scientific explanation of belief only defeats our justification of the belief without telling us whether the belief is (or is not!) *true*.

Recall the following argument:

(1) Causal premise: *E* explains *S'* belief that *p*.

(2) Epistemic premise: *E* is insensitive to *p*.

∴ *S* does not know that *p* (Human moral beliefs are not justified).

Joyce and Street *can* help themselves to (2) without begging questions so long as they are open to the possibility of an absolute, independent truth of the matter about whether *p* (say, whether <killing children is wrong>). But Joyce thinks his conclusion that “*S* does not know that *p*” follows as an evaluative fact from objective concepts of justification and knowledge, and the evaluative norms these involve (§2.4). So Joyce allows independent evaluative truths about morality on the most intelligible reading of (2) as well. For the most serious version of the challenge, we assume that we are in a True World, but cannot have access to the (moral) facts that would make this belief true. Even granting its possibility, we do not know that we are in the True World, after all.

The arguments Street and Joyce use, in other words, only have evaluative import for *believers antecedently committed to objective evaluative truths*. If arguments against objective moral knowledge do not even consider or address these truths, the irrealist simply restates their conclusion, without providing further reason for anyone else to believe it. Willam FitzPatrick makes this case when he writes:

the debunking models [e.g. of Street] would be compelling across the board if there were no moral truths to be discovered, leaving us nothing but evolutionary, psychological, sociological or other similar causal factors to appeal to in explaining beliefs. But it would be obviously question-begging in the course of an argument against realism just to deny from the start that there are [objective] moral truths of the sort realists posit. (2014: 897)

Put differently, FitzPatrick is more than willing to admit that biological creatures like us are more or less wrong about morality, precisely because the reasons we are wrong—the facts we are mistaken about, as it were—are objective (897). Somewhat ironically, it seems that part of FitzPatrick’s problem with assuming subjectivism in the course of an argument against moral objectivity is that it would allow *too many* of our moral beliefs to be correct.

That is, if the standards of moral evaluation are subjective, or if morality is not factive, or if morality is simply a statement of emotion, then the Darwinian dilemma and other consequences of the empirical problem seem *too trivial*.

However, Street writes as if we would evaluate these beliefs in precisely the same way in worlds without objective evaluative truth. In fact, Street argues that many of us think our moral beliefs are objectively veridical, even though they aren't. This implicit assumption that values *are not objective* is already built into her arguments against objective moral knowledge all along. That is, on Street's version of the 2nd premise above, scientific explanations are not only insensitive to particular objective moral truths; they are straightforwardly incompatible with the veridical explanations that posit such truth, by hypothesis.

This argument yields:

(1) Causal premise: *E* explains *S'* belief that *p*. (Street 2006: 115-121)

(2') Epistemic premise: if *p* is objective, *E* is insensitive to *p*.²⁶ (Street 2006: 122)

(3) Subjective premise: if *p* explains *S'* belief that *p*, then *p* is not objective.²⁷

²⁶ Notice that (2) merely says that our moral beliefs are adaptive whether they are true or false. But (2') says that our moral beliefs are adaptive whether or not they are objective: they would be adaptive if objectivism is correct, and adaptive *and*, in some cases, true if objectivism weren't. As far as I can tell (2) grants that these beliefs may be true and is neutral over what such truth consists in. But (2') implies that beliefs are not true or warranted if they are objective, and, if not objective, then at least sometimes true or warranted. Roughly, since our scientific theory would not be correct if morality was really objective after all, objectivism must be false (see Street 2006).

²⁷ See Street 2006 (esp. 154) and Street 2008. Roughly, "badness [of pain] depends on its being a sensation such that the creature who has it is unreflectively inclined to take it to be bad [...] badness hinges crucially on our unreflective evaluative attitudes towards the sensation which pain is (2006: 151).

(4) If (2') and not (3), then S does not know that p . (Street 2006: 130-31)

(5) If S knows that p , then p is not objective (see (3)).

(6) S knows that p . (Street 2006: 152, 53; Street 2016: 330)

$\therefore p$ is not objective. (Street 2006: 156; 2008: 15)

In summary, the argument states that if moral knowledge is possible, then moral truth is not objective. Conversely, if moral truth is objective, then moral knowledge is not possible. Street thinks that this result is implausible, so moral truth is not objective.²⁸

Consider the case of <killing children is wrong>. Street thinks that we do not know that <killing children is wrong> if we believe that this truth is objective. But if we know that <killing children is wrong>, then our belief that <killing children is wrong> corresponds to the fact that it is wrong to kill children. Since this cannot be the case if moral truth is not objective—*since it is implausible* that we cannot know whether or not we should kill children—the belief is subjectively true if it is correct.

Objectivists might find this argument compelling at first glance. It seems that they should agree with each premise. Objectivists would hope that (6) is true in many cases. Most

²⁸ Although Street does not state the argument this succinctly, the following pair of claims, more than 20 pages apart, do resemble the above summary: "...by hypothesis, there is no relation between the forces at work and the "destination" in question, namely evaluative truth. If we take this point and combine it with the first premise that our evaluative judgments have been tremendously shaped by Darwinian influence, then we are left with the implausible skeptical conclusion that our evaluative judgments are in all likelihood mostly off track, for our system of evaluative judgments is revealed to be utterly saturated and contaminated with illegitimate influence. [...] the antirealist opts for what I have said is the more plausible horn of the Darwinian Dilemma, arguing that *of course* there is some relation at work here—of course it is no coincidence that there is such a striking overlap between the content of evaluative truths and the content that natural selection would have tended to push us toward" (2006: 122; 153, emphasis mine).

would assess (6) on a case by case basis, and <killing children is wrong> seems as good a candidate as any for moral knowledge. If (6) is plausible, however, (1) is certainly surprising, but the first chapter argued that it is true as well. And the second chapter supported something much like (2') and (4). And if (2') is true, (3) (and (5), which is much the same) seem right as well. So Street thinks we must choose between (4) and (6): either morality is objective, and our moral beliefs are not warranted, or morality is not objective, and moral knowledge is possible. If we agree that some moral belief is true and warranted, Street thinks we must admit that its truth is not objective.

But notice that objective truth was never at stake in (2). What our original argument said is that our scientific explanation cannot decide the truth of p —it does not track this fact. And since it does not track this fact, we do not know that p . But this does not mean that the belief is not true. It may very well be a fact that it is wrong to kill children, even if we do not know this truth. This sort of independence from our own attitudes is precisely what makes its truth *objective*. And if objectivists are committed to such truth, they should reject (6). And since the argument never decides whether p if (2'), Street is mistaken to think that p is not objective, if S does not know that p .

The problem so stated is that the second premise above begs the question against the objective truth of moral beliefs: it assumes that the truth conditions of these beliefs aren't objective from the outset. It is worth restating, then, that the original premise (2) above operated on the *insensitivity* of biological processes to the objective fact that p .

Furthermore, this insensitivity is made possible—perhaps probable—by the epistemic priority empirical details occupy in explanations of moral belief. The crucial skeptical step

thus *relies* on the immediacy, obviousness, parsimony, or coherence of the scientific explanation of the belief in comparison to the veridical explanation, which the skeptical explanation suspends, rather than refutes, *for the sake of the argument* (Street 2006: 130).²⁹ Again, skeptics use evaluative standards in deciding whether to believe scientific or moral theories. And according to these evaluative facts, the evaluative *moral* facts of objectivism are difficult to know, *regardless of whether they are true*.³⁰

But that we do not know these facts does not mean that our beliefs are *never* objectively true or that such facts are not objective. It does not follow from our lack of moral knowledge that our moral beliefs are false, and it does not even follow from the unavailable conclusion that these beliefs are not objectively true that morality is not objective! In short, in order to demonstrate that moral truth is not objective, Street must demonstrate that moral knowledge is possible. But in order to demonstrate this, Street must assume that moral truth is not objective. The objectivist should not make this assumption. After all, rejecting objective evaluative truth does not demonstrate an objective epistemic fact that *S* knows that *p*. If anything, as the next chapters will argue, this evaluative truth is also subjective if Street's arguments against objective normativity are successful. So although

²⁹ Joyce in 2001 argues like this as well.

³⁰ I suspect that the success of Street's argument depends on whether she equivocates on terms like "objective" or "real". Like Copp and Shafer-Landau, and unlike Korsgaard and Street, I consider some kinds of constructivism as realist positions in meta-ethical debates, in which case Street's argument may very well apply to Korsgaard's and Rawls' approach, at least insofar as she disputes the ultimate *objectivity* of moral claims (Copp 2015: 58). But this cannot be right—Street argues that constructivisms *of all sorts* are immune to her argument in a way that "realisms", as she uses the term, aren't. But the sense in which some realisms are problematic, whereas others aren't, because they aren't really "realist", is far from clear in Street's argument. In other words, Street's stated targets are at once incredibly ambitious, and ultimately ambiguous, precisely because she refuses to admit the possibility of objective moral facts in her argument against them.

the objectivist may have to give up on moral knowledge, they do not need to give up their approach to moral truth.

Street might defend her conclusion by clarifying that her argument never states that it is *true* that morality is not objective, but that it merely argues that we have more reason to *believe* this claim.³¹ Properly understood, I do think that this is the most precise way of stating Street's argument. But as Chapter 4 will show, this argument assumes that what we have reason to believe does not depend on objective normative facts.³² And since this ultimately begs the question against objective epistemic facts as well, the response I detail here is still available to the moral or epistemic realist or objectivist. The objectivist about normativity does not have reason to believe that there are no objective evaluative truths at all. If objectivists have reason to believe anything at all, they have more reason to doubt the truth of their particular moral beliefs than to doubt objective morality or normativity itself.

In conclusion, Street's dilemma does not demonstrate that we have moral knowledge at all, since it does nothing to demonstrate what moral truth itself looks like. What the dilemma accomplishes, at first, is epistemic: it is actually compatible with the objective fact that we

³¹ Street now seems to characterize her argument in much this way, writing most recently: "The mere fact that realism brings with it some apparent problems—even serious ones—is not enough to defeat the view if the alternative, anti-realism, is even more implausible in comparison [...] we can do nothing but opt for that conviction which, on reflection, seems most plausible all things considered" (2016: 331)

³² Here, I have in mind Street's footnotes to her 2006 article, one which states: "an objector might charge that the anti-realist, in arriving at his or her view on the way in which evaluative truth is a function of our evaluative attitudes, must rely heavily on our substantive evaluative judgments (regarding practical reasons). Since those judgments are contaminated with evolutionary influence (and since the antirealist presumably wishes to say that his or her meta-ethical view is true independently of our evaluative attitudes), the objector might argue that the Darwinian Dilemma threatens anti-realism as much as it does realism. [...] I assume for the sake of the argument that [we] share our judgments regarding epistemic reasons; this assumption is complicated by the fact that (as mentioned in note 2) I believe the Darwinian Dilemma can be extended to apply against realism about epistemic reasons" (164).

do not know that any of our moral beliefs are true. Of course, if it is not an objective fact that we do not know this—if it is a subjective fact that we know some moral truths, then it trivially follows that objective moral knowledge is not possible. Without moral or epistemic objectivity, this conclusion is redundant. The more troubling and compelling conclusion is just the more modest result that moral knowledge is not possible if morality is objective.

3.3: In Defense of Joyce

In the last chapters, I have tried to illustrate some ways in which various arguments against moral knowledge involve an implicit account of moral and epistemic facts. For Street, for example, the relevant account on either point is subjectively normative: it is a subjective epistemic fact, if it is a fact, that our beliefs in objective morality do not count as knowledge. These beliefs are false as objective statements, and so we do not have objective moral knowledge. But just as Joyce concluded that the possibility that a moral belief is false shows that we do not know that it is true, Street might have concluded instead that we do not know whether such truth is objective. Instead, Street variously combines or implies several of these conclusions:

(C1): Beliefs in any objective moral facts are not justified.

(C2): Beliefs in objective morality are not justified.

(C3): Beliefs in any objective moral facts are all false.

(C4): Beliefs in objective morality are all false.

The problem with Street's analysis is that it variously confuses or conflates these four distinct conclusions, although Street would agree that (C1) has the most straightforward

appeal. In fact, I have just argued that (C4) is not supported by her argument anyway, properly understood. I contend that Street cannot help herself simultaneously to all conclusions 1-4, especially since a rejection of something a lot like (C4) is the most compelling reason to accept (C1). After all, an argument that has less to prove is, *ceteris paribus*, more plausible than an argument that employs itself in the service of more conclusions than it really might suggest, even if these conclusions are ultimately true. The advantages of (C1) over (C1)-(C4) are, at any rate, plausible *according to normative considerations*, if at all. So in a more general way, the point here is that the appeal of (C1) is most readily explained by evaluative facts about what counts as knowledge.

On the other hand, I have argued that Joyce restricts his argument to supporting only the first of these four conclusions, although he does *accept* or *believe* (rather than know or conclude) (C2) and (C3) as well, while admitting that he cannot know the truth of these claims. Compare:

(C1): The belief <Napoleon lost Waterloo> is not justified.

(C2): Beliefs about Napoleon are not justified.

(C3): The belief <Napoleon lost Waterloo> is false.

(C4): Beliefs about Napoleon are all false.

One advantage Joyce's solution to the empirical problem has over Street's is just that it needs less warrant in that it proves even less than Street. Even if Street is right on all her points, this would also mean that Joyce is also correct in his more modest claims. At any rate, (C1) is all that the skeptic requires to compel the objectivist to take an empirical problem seriously. For instance, I consider myself an objectivist, and I will readily agree

with Joyce that our moral beliefs are difficult to justify. On the other hand, I cannot accept that the empirical problem in meta-ethics shows all possible beliefs as such to be false or subjectively true instead. I especially take issue with the latter claim that objectivism itself is necessarily false, even granting scientific explanations of moral belief.

This is why I am convinced that Joyce's modest agnosticism should not be confused with denying, as Street does, that the truth of the matter is a feature of the world. This is not surprising—science often explains the physical world in spite of, or despite metaphysical ignorance or indecision: empirical explanations, *by their very nature*, do not decide between abstract metaphysical accounts on their own.³³ Of course, in either case, objective moral *knowledge* is off the table. But after the psychological component of belief is explained away, there is still a question of why those beliefs were wrong, if at all. And Joyce's objective approach gives a straightforward explanation of why our beliefs about morality are wrong, if they are wrong. In short, our beliefs are more or less wrong *insofar as* they are insensitive to real moral facts. And it is only because we do not know that they are *not* wrong, at the same time, that Joyce has not begged any questions against the objectivists or any particular account of moral knowledge.

On the other hand, the conclusion that moral truth is not objective if our scientific explanations are correct also requires a metaphysical argument against the objectivity of

³³ It may be the case, however, that the correct metaphysical account accommodates natural moral facts. Perhaps this account involves scientific properties either as we already use them, or in a way we do not yet understand. Whatever the case, the naturalist still faces an empirical problem, even while granting such an account. Sturgeon (charitably) reads Harman as making just his point—after all, Harman's empirical argument is question-begging if it takes the scientific irrelevance of moral facts to depend just on their supposed irreality. Rather, the problem seems to be that moral facts, even of the natural sort, are of no consequence at all to our best scientific theories of the world, even when we grant the reality of such "facts" (see Sturgeon 237).

the moral and the epistemic domain. While it is true that the objectivist also must provide a metaphysical account in favor of moral truth, I have demonstrated here that the relevance of empirical considerations for moral objectivists are just that our moral beliefs are unjustified, as a matter of objective epistemic fact.

3.4: Moral Irrealism and Normative Truth

In this chapter, I have argued that although objectivists do not have reason to believe in particular moral truths, they do not have reason to believe that morality is not objective. At the very least, the scientific considerations Street discusses do not demonstrate on their own that morality is not objective. If there is reason to believe that morality is not objective, Street must demonstrate that moral reasons or epistemic reasons are not objective. Unfortunately, employing the epistemic reasons at stake in favor of such a conclusion assumes too much. If Street thinks that we have reason to believe that morality is not objective, she must convince us that these reasons are not themselves objective.

This is why I have just suggested that Street isn't just a moral irrealist, she's an evaluative irrealist as well. This is because she cannot appeal to objective epistemic facts in the course of an argument against objective morality. As we will see in the next chapter, Street thinks that there are no objective *evaluative* facts at all—or, at any rate, we cannot know any of the objective evaluative facts that would make our beliefs true, and so we cannot even have *objective* evaluative knowledge, ever! So Street generalizes her dilemma about moral value quite readily to any domain of value, to support the conclusion that particular domains of value, like morality and epistemology, are not objective.

But this anti-objectivism does not solve the empirical problem. Rather, moral objectivity is at least (in principle) compatible with the normativity that makes the task of unifying our natural and normative explanations so difficult. I will argue in the next two chapters that if this normativity is to be coherent at all, it must be objectively coherent, and theoretically unified across the practical and theoretical domains. On the other hand, if objectivists are to give up moral knowledge for subjective reasons, an argument against objectivism cannot assume these subjective reasons from the outset. In what follows, I will argue that Street's arguments for epistemic irrealism do not provide objectivists with reasons they should accept. There is no good reason to believe that there are no ultimate, objective reasons at all. If this is right—if there are good reasons to believe or disbelieve anything at all—then it turns out that our lack of moral knowledge is an objective epistemic fact, rather than the conjunction of moral and epistemic irrealism.

4

Skepticism and objective epistemic facts

4.0: Introduction

Much of this thesis discussed ways in which we may have made dramatic doxastic mistakes—ways in which our moral beliefs about an objective moral world *may* miss their mark in, for example, an amoral world they represent inaccurately. I've also described ways in which skeptics argue that science (or inquiry more generally) can correct for epistemic errors of this sort. The previous two chapters examined two ways in which we might correct for these errors: we might give up 1) the justification or 2) the objectivity of our moral beliefs. In each of these ways, Joyce and Street argue that we have reason to give up on objective moral knowledge. In agreeing on this point, both philosophers grant the truth of some epistemic norms: they hope that the facts that convinced them of their own argument(s) also influence others to come to similar conclusions.

In this chapter, I call this normative truth into question. This is because the epistemic norms that motivate the empirical problem in meta-ethics are themselves the target of an analogous meta-normative empirical problem. This argument is not original—I owe it to Street's *Evolution and the Normativity of Epistemic Reasons* (2011), and much of this chapter compares her approach to that of epistemic objectivists.

In §4.1, I outline Street’s epistemic irrealism, and in §4.2, I explain Street’s meta-normative empirical problem. However, I argue that her irrealist considerations do not *coherently* rule out objective epistemic facts. In §4.3, I articulate some advantages of epistemic realism, and introduce an account of “irreducible” epistemic facts, courtesy of Terence Cuneo. After explicating Cuneo’s criticism of epistemic irrealism, I argue that Street’s argument for moral irrealism is unconvincing if her argument for epistemic irrealism is right (§4.4).³⁴ In short, *at most* one of Street’s solutions to these two empirical problems can succeed: reasons for belief or disbelief are most convincing if objective, or self-defeating otherwise.

4.1: Epistemic Irrealism and Normative Truth

In the last chapter, I complained of the ways Street assumes moral irrealism in her analysis of the empirical problem. Here, I outline the sort of epistemic irrealism that both motivates and concludes an analogous meta-normative empirical problem. Consider these positions:

TABLE 2: Comparison of Meta-Epistemic Theories

	Are there reasons for belief?	Is truth <i>necessarily</i> a reason for belief?	Are all our reasons for belief entailed only from our own practical attitudes? ³⁵
Epistemic realism	✓	✓	✗
Epistemic irrealism	✓	✗	✓

³⁴ This argument is disjunctive: if Street’s argument for moral irrealism is necessarily correct because it is true, her argument for epistemic irrealism is incorrect. If epistemic irrealism is correct, her challenge against moral objectivism is not plausible, at least not as a matter of objective epistemic fact.

³⁵ See Street 2011 and Street 2016 for a precise statement of this question. Street targets views according to which we sometimes have practical or theoretical reasons that do not follow from our “practical standpoint”.

In the above, irrealist accounts differ from objective epistemic realism in that they 1) reject truth as a necessary reason for belief because 2) our reasons for belief are only correct if they cohere with our antecedent attitudes and values. For example, Street's irrealism gives epistemic rules authority only insofar as anybody *endorses* the causal or constitutive process of belief they happen to follow. Irrealists maintain that any given norm can be given up or adopted insofar as the believer ultimately desires such change. In this way, epistemic irrealists like Street believe in radical doxastic freedom. Epistemic autonomy of this sort entails the ultimate contingency of any reasons whatsoever—including whether irrealism itself is *really* "worth believing". In this way, her constructivism exonerates even those that don't value survival and prefer what's "amusing, exciting and so forth" rather than truth (2011, 236). In this way, each believer ultimately makes normative mistakes, if at all, with respect to their own values and standards.

So rather than worrying about whether truth really matters as such, epistemic constructivism vindicates the value of truth: that is, truth *does* matter when we make it matter. In other words, the standards of correct belief are subjectively normative, if at all. But the normativity doesn't really come from the world—it isn't true when no one thinks about it—rather, believers give themselves their own rules, as it were, depending on what they want. At the same time, however, Street's *constructive* normativity is compatible with widespread agreement about normative truths: truth-valuers, for instance, share a common etiology for their normative beliefs.³⁶ Some of the norms that the subjectivist

³⁶ Epistemic irrealism might license some "universal" norms. Street doesn't think this universality actually obtains in either the moral or epistemic domains, so I am concerned with the intuitive subjective claim that S ultimately has the reasons S *wants* (if S wants or values anything at all). In other words, there's no easy way of predicting what these reasons are independent of particular desires and circumstances (see Street 2011).

adopts in achieving these aims might coincide with those that the objectivist endorses. But the irrealist and realist disagree on how optional such norms and aims are, as an ultimate matter of fact.³⁷ Again: the irrealist maintains that any given norm can be given up or adopted insofar as the believer ultimately desires such change.

But I will argue in §4.3 that this sort of irrealism misses the whole point of normativity in our explanations—normative objectivists aren't just concerned to explain that *all* of their normative beliefs are *correct*; they also wish to explicate compulsory standards that locate and then excise bad beliefs in this domain. The normative beliefs that destruction is good, or that amusement is ultimately most worth taking into epistemic consideration, are just such candidates for mistaken normative beliefs by the standards I will introduce in §4.3. These norms cannot be entertained by any construal of objective rationality. These are norms, in short, that are self-defeating or unintelligible—or, by analogy, that can't be understood by players as legitimate rules of a game they all purport to “play”.³⁸

³⁷ In particular, the irrealist point that truth is not valuable *per se* or as such diverges from the robust epistemic objectivism I defend, according to which truth is necessarily valuable in virtue of our normative activity of belief. There is a difference between the robust objectivism Street criticizes and the possibility of a “universal”, but attitude-dependent epistemic constructivism. Street criticizes both moral and epistemic objectivism, but thinks that a universal constructivism is at least intelligible in the latter domain. For instance, in 2011, Street writes “I doubt that any robust conclusions about morality follow from within the standpoint of every valuer, no matter what his or her ‘starting set of values’ is. [But] it’s possible that we will discover an objectivity *of the specified sort* in the epistemic case, or perhaps in other cases, even if we can’t find it in the moral case” (2011, 216) Note that “of the specified sort” (emphasis mine), as it were, will at most turn out to be “universal” norms, rather than objective or thoroughly attitude-independent in my sense.

³⁸ Street compares the business of believing with games like tennis. If we agree to a tennis match, then we have to play by the rules; if we believe things, then we will already buy into “truth”. But there is no rule in general that makes tennis compulsory for us all; similarly, there is no ultimate fact according to which we must believe truth. So “merely failing to believe” for Street, is “not in and of itself a mistake, any more than failing to play tennis is in and of itself a mistake” (2011, 231). Compare Cuneo: “a wide range of epistemic reasons stands in sharp contrast to the norms that govern activities such as chess and tennis and bear a closer kinship to the norms of logic and mathematics. Indeed, arguably some of them *are* norms of logic and mathematics” (50, emphasis his). In what follows, I doubt that the norms of belief, logic and math are

4.2: Epistemic Facts and the Meta-Normative Empirical Problem

According to the position above, irrealists maintain that doxastic processes are always *hypothetically* connected to the world. Since the practical aims that regulate such processes require getting some facts right, belief involves truth accordingly. Put differently, that something is true does not make *S* believe it, but that *S* wants anything at all explains why truth matters. For instance, *S* should only believe that an animal is dangerous if the animal *is* a threat and *S* wants to stay safe. Further, if false beliefs about the world help *S* live just as well, then *S* might get by believing, say, that a safe animal is dangerous so long this false belief furthers *S*' survival (or some other goal).

Street illustrates this irreality and contingency of normative reasons like so:

Let's suppose the first two valuing creatures ever were born—remarkably, as it happened, in a fairly sophisticated fashion. Until that moment, nothing had ever consciously valued anything. [...] As it so happened, the first valued its own survival and nothing else, whereas the second valued its own destruction and nothing else [...] when the first creature judged that its own survival was good and the second creature judged that its own survival was bad, the first was not recognizing some normative truth that the second was somehow missing. The first creature survived, of course, but this isn't because its judgment was true, but rather simply because that creature tended to do what promoted its survival. The second creature, in contrast, didn't survive, but that's not because it failed to be sensitive to any normative reality; it's rather simply because that creature sought to destroy itself and succeeded. (2008: 16-17)

So said, Street's argument for epistemic irrealism closely parallels her argument for moral irrealism. In both cases, Street argues that subjective normativity in each domain is more

optional or subjective—the point of the “rationality game”, so to speak, is that we (all of us capable of dispute, anyway) already make mistakes *according to its rules*: we are already “playing”, whether we want to or not.

compatible with a scientific explanation of our world than the objective facts of epistemic and moral realism.³⁹ While objectivists argue that humans can access correct normative truths, Street thinks this confidence doesn't withstand empirical scrutiny since it would be an "implausible coincidence" were biological beings to happen across a truth that isn't empirical in any ordinary way (2011: 233). As she puts it:

Anyone who thinks it's a real fact about reasons that considerations about what's true, as opposed to other kinds of considerations—such as what's amusing, exciting and so forth [...] is confronted with a puzzle about how it could have happened that the causes which shaped our views on this normative matter landed us so close to the independent normative truth. (2011: 236)

Street's alternative, sketched above, is that whatever considerations biological beings take in favour of believing or acting are ultimately true, if at all, in virtue of their valuing those considerations rather than others (2011: 241).⁴⁰ This is how Street avoids the difficulty of explaining why the odds worked out in favour of true, objective, normative belief. On the other hand, objectivists cannot assume a veridical explanation for the relevant norms they endorse. Even if there are objective epistemic facts, Street argues, these can't *vindicate* any of our normative beliefs. Rather, beliefs in objective epistemic facts are mistaken, especially given a biological explanation of the normative beliefs' causal etiology. Or so the argument goes.

³⁹ In this way, Street thinks that practical and theoretical reasons are unified on a general normative account, and that the irreality of one domain of reasons counts against the reality of normative reasons in another domain (2016). I will return to this implicit connection between the moral and epistemic below, in Chapter 5.

⁴⁰ See also Street's "Constructivism about Reasons" (esp. p. 34-36). Although the paper concerns meta-ethical freedom about reasons, it clarifies the meta-epistemological analogue elsewhere of her 2011.

Call this the meta-normative empirical problem. A formal representation of the problem runs like so, where *E* stands for a scientific theory of natural selection (and so, a scientific explanation of normative belief) and *N* stands for “normative”:

(1) Causal premise: *E* (science (of natural selection)) explains beliefs in objective *N* facts.

(2) (Meta)epistemic premise: *E* is insensitive to the objective *truth* of *N* facts.

∴ Our beliefs about *N* facts are probably false. (Street: normative objectivism is false)

Again, Street’s reasons for the stronger conclusion in parentheses above is her confidence that many of our particular normative beliefs are actually true instead (2011: 243). We know all sorts of things about knowledge itself, and since we are more or less right about these normative facts, and since we would not be right about these normative facts if objectivism were true, objectivism must be false. At any rate, empirical considerations explain why the belief in objective normativity persists, even if it’s false.

Subjective normative “truths”, in contrast, aren’t in this tension with the biological explanation—at least not in this way. Again, Street takes biological considerations to count *in favour* of her epistemic irrealism. Street calls her solution “much more plausible and parsimonious” than accounts of independent normative truths precisely because it is compatible with whatever norms humans happen to endorse. In illustrating this account, she asks:

Why did creatures who took the fact that all previously observed tigers were carnivorous to be a reason to believe that the next tiger would be carnivorous do better than those who took the same fact to be a reason to believe that the next tiger would not be carnivorous? According to [my] account, the answer is not that the one group was sensitive to a normative fact that the other group missed; it is rather that creatures who made the former

normative judgment tended to believe that the next tiger would be carnivorous, and, as it turned out, the next tiger *was* carnivorous. (2011: 240-241)

On Street's view, the inductive rule at stake in the example isn't true *because* it explains believers' success; rather their success explains both why they take the inductive rule to be true *and* why it is true *for them*. And so, their having used the inductive rule is ultimately what makes their rule a normative rule. Thus, while any account of normativity might hold that "the fact that n/m of the observed Xs have been Y is a reason to believe that n/m of the unobserved Xs are Y", the normative truth of this rule, like the normative truth of a moral principle, ultimately depends on whether believers value it (2011: 218).

One problem for this view is that our attitudes do not seem as relevant to epistemic evaluations as they were in the moral case. Take Street's example. *S* has come across 99 carnivorous tigers, so *S* has evidence that these animals are dangerous. As it happens, *S*, in coming across the next such animal, considers whether it is friendly. While *S* might believe that the tiger is friendly—that is, the cognitive response is *possible*—this possibility seems to depend (at minimum) on the facts about whether the tiger is dangerous. While *S* may not *have to* respond to these facts, doing so (in believing that the tiger is friendly, say) seems to miss some facts. So regardless of *S*' actual response, the degree of responsiveness to the relevant facts is itself (at least) a straightforward fact about *S*' belief, rather than an attitudinal fact constitutive of *S*' particular wants, desires, intentions and so forth.⁴¹

⁴¹ Admittedly, despite the more obvious facts of the matter about what *S* believes and whether *S*' belief is correct, understanding relational facts holding between *S*' belief and *S*' evidence for that belief—between *S* and what *S* sees in the world, as it were—is less straightforward. At the very least, the fact requires a comparative evaluation in the way the other facts do not. For instance, the above example stipulated that *S* might believe that a given tiger is friendly. But it also stipulated that *S* was relevantly aware of the *fact* that tigers are dangerous carnivores. And my point in the above is just that the belief that the tiger is carnivorous

Denying the objectivity of such facts is also at odds with much of traditional epistemology.⁴² Take, for instance, John Locke's imperatives against wishful thinking and willful ignorance (710) or Hume's normative platitude: "a wise man proportions his belief to his evidence" (146). These imperatives are epistemic: they argue that human beliefs are better or worse according to the degree to which they take relevant reasons or evidence for their beliefs into consideration—the better the evidence, the better the belief. Call this epistemic norm EVIDENCE. Although Street agrees with EVIDENCE she disagrees with Locke and Hume that EVIDENCE is objectively true or false. Again, her reasons for disbelieving the ultimate truth of EVIDENCE is a scientific account according to which we would have believed EVIDENCE even were EVIDENCE objectively false or not objective.

But of course, the reasons for believing this scientific account are *themselves* ultimate norms much like EVIDENCE. That is, Street argues that it is simpler to excise objectivity from theories of the world, and this simplicity gives us reason to believe that objectivism is false. But this amounts to saying that we should believe irrealism because it is true, since irrealism saves our normative beliefs, we *ought* to accept epistemic irrealism (2011: 241). This dialectical move *relies* on the immediacy, obviousness, parsimony, or coherence of the scientific explanation of moral belief in comparison to the objectively veridical explanation,

is ultimately correct *because it is true*, and beliefs otherwise, especially in spite of evidence to the contrary, are incorrect not *just* because they are unwarranted, but more importantly, *because they are false*.

⁴² I find that the word "fact" best expresses the sense these evaluative standards make it the case that beliefs are correct. But I do not think that using this *word* in this way is essential to my argument. Compare the "fact" that 2 is the correct answer to 1+1. Once we define 1 and the plus relation, that 1 bears the plus relation to 2 and itself, rather than another number is "exactly the same for everyone who deals with it" (see Frege 72). (Is it a "fact" that "1+1=2"? Does it matter what we call this concept?) Here, I am trying to articulate what "truths" like this amount to. If my use of the word "facts" to express the truths of these cases is distracting, substitute whatever word better fits ("eine Tatsache", say). The important idea is that there are real standards, and real truths about real cases, even merely conceptual ones, independent of our contingent values and beliefs.

which the skeptical explanation could, Street argues, either suspend or abandon entirely *for the sake of the argument*.

So the subjectivist assumes, *for the sake of the argument*, that the inductive rule is not objectively true (2011: 240). It follows that anyone believing that the inductive rule is objectively true is already, automatically, *always* wrong. It also follows that when believers are aware of some natural explanation of our belief in the inductive rule, they should then disbelieve the objective truth of the inductive rule. In sum, since the irrealist interpretation of the inductive rule is currently consistent with the empirical facts, no one, least of all the subjectivist, really knows that the inductive rule obtains as a matter of objective “fact”.

The worry, however, is that Street cannot defend the attitude-dependent normative solution to the meta-normative empirical problem, since her own normative view is at stake in this problem as well. That is, Street does not have an argument for this attitude-dependent or subjective normativity that does not assume the very epistemic norms she is arguing for. Indeed, instead of demonstrating the truth of such norms, Street explicitly appeals to epistemic norms that the subjectivist should not assume on pain of the very norms she assumes! At some point, Street’s argument works because of these norms, but if these norms are themselves the reason her argument works, and if they are not good reasons in general, then her argument against them does not really work at all. So subjectivists *are* doing something wrong by flouting the rules they use, given that they in fact *want* to use them. So long as subjectivism has any say at all what anyone should believe or do—so long as it continues as a candidate in the game of rationality, its own rules make such rationality compulsory. Using a conclusion as a hidden assumption in an argument for

itself cheats *any* construal of these rules, subjective or otherwise. So should anybody be bothered by the skeptical explanation of normative belief? Should they give up normative and moral objectivity on this account?

On *any* answer, Street's argument against epistemic or normative objectivity *should* come across as incoherent, even if similar arguments about morality seemed compelling in comparison. After all, normative beliefs are required to make sense of the conclusion that our normative beliefs are not warranted. But if these beliefs are "not warranted", then we cannot coherently employ them in making sense of *that* very conclusion. But if they are warranted, and if there is a flaw in *this* scientific explanation of normative belief, then subjectivist owes objectivists the "right" explanation (so long as she continues to believe that there are, as a matter of fact, better and worse explanations). Hence, Street's solution is incoherent—in particular, using the meta-normative empirical problem to establish epistemic irrealism is self-defeating.

In conclusion, if we assume subjective normativity from the outset in an argument against objective norms, the problem for an irrealist is that if all access to such normativity is defeated in this way, there is nothing really left to have done the defeating to begin with. Street tries to avoid this self-defeat by denying wholesale that these evaluations are themselves *matters of fact* (2011: 229). Street argues, in effect, that constructive moral and epistemic reasons can only be made to work if they are in some ultimate sense *subjectively* true (2008). And if epistemic evaluations are subjectively true, they are not matters of fact that *must* matter to us all. So even if our epistemic explanations must employ subjective norms instead, even from the outset, it cannot even be used to establish the subjectivity of

these norms. Even a subjectivist *has* to believe something or another, and that the space of normative possibilities open to them is theoretically constrained. That this normative space looks *objective* begins to look a lot less coincidental on a meta-normative account of all possible positions in this debate.

4.3: Objective Epistemic Facts

In what follows, I outline an objective interpretation of the sort of “rules” Street and Joyce employ, however implicitly, in arguments for moral skepticism. I do not take my account to demonstrate any *particular* facts or limits about what to believe. But insofar as believing is an intelligible activity, and insofar as anybody has an interest in continuing or improving this activity, I argue that there are at least facts concerning whether beliefs are more or less responsive to reality. Put simply, these facts are normative, and they are objective.

I try to sketch a position that is compatible with biological considerations. In particular, I will grant Street her biological data in what is suggestive of objective norms. I had already admitted, for example, that the epistemic imperatives canvassed above had explicitly biological consequences. To repeat an example, the evidence that tigers are carnivorous only matters insofar as *S* believes that tigers are carnivorous. So it is true that the *belief* that “n/m of the observed Xs have been Y is a reason to believe that n/m of the unobserved Xs are Y” preserves believers. It is no secret that beliefs in norms like this inductive rule are sufficient for a causal explanation. To grant Street her point, it does not really matter for the natural explanation whether any particular normative beliefs are true—it is enough that merely thinking them true can continue, say, *S*'s (and other's) practice of believing.

This is why Street insists that biology does not vindicate the epistemic facts I defend here. And certainly, there are plenty of biological beings that get along well enough without cognition of the sort outlined above.⁴³ But biological processes also select for our capacity to see the world in terms of “truth”, and that we understand our cognitive representation of reality as corresponding to the way things really are, whether or not we successfully represent reality at all (Hoffman 8). Again, this does not mean that we are *successful* in achieving this goal in all worlds in which we are biological beings of some sort.⁴⁴ But it does mean that in worlds in which we understand ourselves to have beliefs, we are in worlds where there are beliefs; and in worlds where there are beliefs, there are normative facts of the sort that it is *better* to believe something or the other *correctly*.⁴⁵

Ultimately, whether anyone *can* ignore or dismiss rational arguments or evidence for belief is beside the point; the very concept of belief presupposes that people *should respond* to an obvious and important truth *by believing it*. Insofar as we understand belief *as such*, we cannot both continue our cognitive lives and *stop believing*.⁴⁶ No less can anyone coherently

⁴³ At least, most discussions of the empirical problem agree that *if* our moral beliefs are insensitive to truth (which is what they usually disagree about), then we simply have to or should stop believing what we used to (although we can continue to accept it). In what follows, I disagree with these irrealist arguments.

⁴⁴ Here, I am gesturing very deliberately towards Justin Clarke-Doane’s evolutionary debunking of mathematical realism (2012) as well as Donald Hoffman’s evolutionary debunking of physicalism (or, more modestly, perceptual realism) given that perceptual beliefs about the material world, because adaptive, are all false (as Hoffman puts it, “natural selection drives true perception to swift extinction”) (2009).

⁴⁵ So long as there are worlds at all, there are truths, and so long as we recognize that there are truths in any worlds, we cannot be mistaken in seeing that we should recognize it. So the sort of cognition I just outlined requires that beliefs always involve a (better or worse) representation of reality, and our recognition of this fact just consists in our being sensitive to *that* fundamental epistemic fact as such.

⁴⁶ Here, the move to reject truth as a necessary reason for belief deserves Samuel Johnson’s or Thomas Reid’s responses to analogous arguments for idealism or perceptual skepticism: “If a man pretends to be a skeptic with regard to information of the sense” says Reid, “and yet prudently kept out of harm’s way as other men do, he must excuse my suspicion, that either he acts the hypocrite, or imposes upon himself” (170).

endorse a “stop-believing” norm since the principle requires doxastic uptake of the sort it rejects. So the problem with Street’s argument so stated was that it confused normativity with necessity. Street did not grant the objectivist any common ground, since her solution to the meta-normative empirical problem assumed that objective norms are not essential to epistemic explanation. And the reason for this assumption is that belief in objective epistemic facts is not required for knowledge. But it does not follow from the fact that we *can* believe such facts are not objective that they aren’t objective!

For example, even if *S* could *conceivably* maintain that the last 99 tigers are carnivorous *and* that the next one isn’t—even when the contrary is evidenced by yet another unfortunate victim—there is something objectively right, insofar as anything in epistemology is right, about criticizing *S* for this “belief” (disbelief). The correct belief that the tiger is carnivorous just tracks the relevant facts about *S*’ environment and doxastic history *better* than thinking of the tiger as a friendly herbivore. Such beliefs are meant to capture whether, for instance, the tiger is *in fact* carnivorous—whether they are *correct* in thinking themselves safe from it (Shah 199). For this reason, giving such believers further evidence of, say, the danger they are in is typically sufficient to change their mind with respect to whether they are safe. Ignoring such evidence results in false beliefs, or only accidentally true beliefs. It misses the relevant facts, and a belief that misses some fact *necessarily* fails in its aim.

There are distinct advantages to epistemic objectivism so conceived. For one, it solves the fundamental epistemic problem of what to believe once and *for all*. If I am right, the correct answer to the question of what to believe is just whatever happens to be true, simply

because it *is* true that it is correct to believe truth. (In much the same way that the belief <I exist> is immune to doubt, the negation of <true beliefs are correct beliefs> is unbelievable, since <false beliefs are correct beliefs> is either false, and so incorrect, or else incorrect if true).⁴⁷ Descartes concluded that the skeptic cannot coherently doubt his own existence (105), and I would argue that skepticism about the epistemic axiom <true beliefs are correct beliefs> is unbelievable for similar reasons. As Nishi Shah puts it, every question about whether to believe that *p* necessarily collapses into the question of whether *p* (201).⁴⁸ And since the truth of <true beliefs are correct beliefs> is our reason to believe it, the truth of some other fact is also a reason to believe that fact, according to the axiom just considered. Truth is *not* a normative property of beliefs, but beliefs that express truth do have a normative property: they are correct. Perhaps the similar statement “true beliefs are better than false beliefs” better expresses the difference between truth and what it is *best* to believe.

Objective epistemic facts so considered also provide absolute, definite answers to other epistemological problems downstream from the fundamental problem above. Objective epistemic reasons make it the case that there are better or worse answers to epistemic problems. And it follows from the objective epistemic fact that true beliefs are correct

⁴⁷ As Hazlett puts it: “the proposition that I exist turns out to have an amazing property: I can’t believe it and be wrong. That’s what Descartes means when he says that the proposition that I exist is necessarily true whenever it is put forward by me or conceived in my mind. And, obviously, this applies not only to me, nor only to Descartes, but to anyone: it is impossible for someone to believe that she exists and for her belief to be false. This proposition cannot be falsely believed” (2014: 79).

⁴⁸ Shah adds (elsewhere): “‘correct’ does not mean ‘true’. Strictly speaking, propositions, not beliefs, are true or false; beliefs are psychological states whose contents are true or false but are themselves neither true nor false. The claim above states that beliefs are normatively assessable as correct or incorrect in virtue of the non-normative property that their propositional contents have of being true or false” (2011b: 100).

beliefs that there are better and worse ways to believe, or better and worse beliefs, according to some further epistemic facts with respect to the truth, or the way the world is. In much the same way that Descartes proceeded from very limited cases of indubitable beliefs to evaluate and justify other beliefs, so the discovery that true beliefs are correct beliefs motivates our evaluation of all beliefs.

In this, I am in agreement with Terence Cuneo and Nishi Shah, who construct their normative ontologies from the very activity of belief. Both Cuneo and Shah think that our beliefs in normativity, whether true or false, already exhibit the most fundamental normative properties. Such beliefs are fundamentally evaluative and cognitive—by definition, belief just is the activity of thinking *of* something or the other (as true), with more or less success. Cuneo and Shah point out that thinking of at least one thing as true in this way is compulsory of all cognitive agents: belief is not optional for believers, and this in itself is a fundamental truth about *us*. For Cuneo, for instance, “epistemic discourse [...] consists only in the sincere utterance of epistemic sentences and at least one such sentence is used explicitly to present an epistemic proposition” and for Shah “believing anything at all makes it the case that there is at least one true normative proposition” (193).⁴⁹

It follows from our status as persons that *in fact* believe at least one thing or another that we are committed to at least some epistemic norms, and that we are committed to thinking

⁴⁹ Shah argues that the *belief* that there is a normative proposition *entails* at least one true normative proposition: if this belief is true, the normative proposition that the belief is correct is true, and if it is false, the normative proposition that the belief is incorrect is also true. Hence, the belief that <there is at least one *true* normative proposition> cannot be false. As Shah puts it: “If anyone ever wonders whether or assumes or desires that there is at least one true normative proposition, then there is at least one true normative proposition [...] it follows from the fact that anyone has any of these attitudes that there is at least one true normative proposition, because having any of these attitudes requires having some beliefs, and believing anything entails the existence of some normative property” (194).

of these as true, if we consciously think of anything else as true at all.⁵⁰ Now, it is certainly an open question which of the norms we commit to are the right norms. Indeed, it is according to its normative status that we take this question seriously. And there are also good questions about what the nature of these norms themselves are. For instance, Cuneo argues, somewhat controversially, that these norms are “irreducible” facts that satisfy “a commonsensical conception of epistemic facts, fundamental to which are platitudes that concern the content and authority of such facts” such that epistemic realism is true “if and only if such facts exist that satisfy the content and authority platitudes fundamental to our commonsensical conception of epistemic facts” (56).

The reasons we have for thinking this account is true ultimately depend on our analysis of epistemic concepts themselves. Whatever else it isn't, knowledge is *at least* an actual event in the world: it happens when someone believes things correctly for the sorts of reasons that generally guarantee correct believing (“good” reasons, like the inductive rule).

Certainly, there is a descriptive, psychological ring to rationality so construed. But this does not show that the normative facts about knowledge are merely descriptive. If rationality is constitutive of knowing anything at all, or believing at all, that humans are so constituted

⁵⁰ Or, alternatively, if we think at all. An interesting philosophical puzzle, no doubt, is whether it's possible for someone much like us to behave much like us without believing anything at all. Such a person would arguably be something very much like David Chalmers' “zombies”—that is, entities that behave like us, without our inner mental life (a machine that does not really “think”—see Chalmers (106) for this thought experiment). But I'll dismiss this digression quite readily: if you are reading this sentence, you are likely understanding it (if not, it is probably my fault, not yours), and this understanding consists in your thinking it true, or false, and so on. In this precise moment, you are already reacting to these sentences and reflecting on what you read. And this is just the sort of reflection characteristic of what it is to “think” at all, which is just to respond and represent things around us *veridically*. From a very young age, most of us are like this. We are not machines. We are biological and cognitive beings with a rich inner mental life. That this is so, I argue, is all we need to get epistemic normativity off the ground.

does not count against objective normativity. It is hardly any coincidence that humans recognize themselves in characterizing the structure of rationality, for it is the human activities of belief, knowledge, and all the rest that rationality explains and improves. In the shared human experience, beliefs *are* responsive to truth and evidence; they are collected, refined, and evaluated in the very epistemic processes constitutive of rich and rewarding cognitive lives. These objective features explain both how believers come to know what they do, and what they should believe next, given what they already know and believe. Epistemology, in short, invokes the tools humans already use to evaluate and change existing beliefs or behaviour in the service of furthering and improving cognitive activity. This is just why deliberation or reflection about what to believe often presupposes normative conditions for rationality.

I take it that realists and irrealists alike must admit that the only approach to epistemology (or philosophy) in question presupposes compulsory or mandatory standards for human belief. And on the considerations introduced above, these standards are not just a matter of preference—they are objective characteristics of sophisticated cognitive lives.⁵¹ According to the account I sketch, there are epistemic facts that amount to how well believers are connected to their world. Believers that encounter such facts about themselves have an epistemic advantage insofar as they can become responsive to the kinds of norms that

⁵¹ The account I defend here is hardly original. Cuneo points out that philosophers in the Anglo-American tradition take epistemic objectivism to be the “default” view in the history of philosophy (hence, his term “realism of a *paradigmatic* sort”). Here, I simply tease out the normative implications of the standard canonical claim with which Aristotle begins his *Metaphysics*: “*By nature*, all men long to know” (980a, 4). According to Aristotle’s claim, for instance, our rational nature makes all of inquiry possible (see also Hazlett 21 for an illuminating discussion of Aristotle’s position). If there is any point to thinking of anything at all, it lies in responding to the world—of leading more or less *appropriate* cognitive lives. And insofar as anything is human or cognitive, it does not choose itself as such. Some things just *are*.

ideally, if not typically, constitute or guide a responsive doxastic interaction with the world. Furthermore, this interaction hardly seems optional: believing anything at all only makes sense once we take the relevant constraints or connections with the world into account. On the other hand, denying truth as a *necessary* reason for belief gives up on the entire business of believing that the constraints make intelligible.

Of course, it may ultimately be true that there *is* no reason to believe anything, but in response to this, Cuneo states:

A realist conception of the world is something that, all else being equal, ordinary mature human agents whose cognitive faculties are functioning adequately in a world such as ours take for granted in their everyday doings and believings [...] so deeply ingressed is this way of viewing reality in our shared world picture, that it is difficult to find propositions more obvious than those constitutive of realism that can be employed in a non-question-begging argument in favor of it. This [...] is why positive arguments for realism regarding the external world are so difficult to come by. And it is also why, when arguments are offered for realism, they generally have the form of teasing out the implausible implications of its rivals rather than citing the theoretically positive characteristics of realism itself. (10)

In this “argument [...] from elimination”, Cuneo contends that according to his objectivism “we have strong reasons to believe that epistemic facts exist”, and that according to alternatives, “there is no epistemic reason to believe anything” (8). So Cuneo’s commitment to the “commonsensical conceptions” of a normative worldview leads him to reject epistemic irrationalism as an informative alternative to the default view.

Says Cuneo:

For fundamental to our commonsensical conception of the world, I’ve suggested, is the assumption that radical epistemological skepticism is false. In their everyday doings and believings, ordinary agents implicitly assume that any position according to which there are

no epistemic reasons, no beliefs that display epistemic merits and demerits, and no arguments should be rejected. (123)

Cuneo argues that since ordinary agents are *in fact* “minimally rational” (they “assume that radical epistemological skepticism is false”), epistemic realism preserves what Cuneo calls “our commonsensical conception of the world” better than does irrealism. In particular, “accepting epistemic nihilism is incompatible with being a minimally rational agent, for epistemic nihilism implies radical epistemological skepticism” (123).

This is, of course, just the sort of strategy that I pursued against epistemic irrealism in §4.2. This strategy does not show that irrealism is false, but illustrates instead that irrealism cannot compete with realism as an ultimate account of rationality. For one thing, Cuneo points out that irrealism is self-defeating, in that it is explicitly committed to the following “Moorean paradox” according to which the mere acceptance of irrealism “guarantees that it is unjustified or unwarranted” (117):

MP: Epistemic nihilism is true, but there is no reason to believe it.

But of course, there isn’t much sense to believing both conjuncts of MP, since “the acceptance of the proposition expressed in its second conjunct undercuts the justification or warrant we might have for the proposition expressed in its first conjunct” (118). In other words, as I’ve argued above, the assumed absence of epistemic reasons or “facts” necessarily counts against any “argument” against such reasons or facts themselves.

Thus, Cuneo argues that “either epistemic nihilism is self-defeating or it implies that there could be no arguments for anything” and this itself flouts the “commonsensical conception of the world” that objectivism outlines and defends (120). So although irrealism does not

make sense of “minimally rational” agents, this objection does not amount to a knock-down argument against irrealism, since it is the denial of such minimal rationalism that is central to the irrealist account. This is why Cuneo concedes that “there is no plausible non-question-begging argument to establish that epistemic nihilism is false” and that his argument simply draws out the implications of *their* view instead. Ultimately, Cuneo merely concedes that epistemic irrealism *may* have the “unenviable status of both being false and committing those who accept it to accepting Moorean paradoxical positions” (118).

This is why a conclusion about what to believe does not follow from a premise about what is true on Street’s account—or on any irrealist account, for that matter! Belief is a thoroughly, irreducibly normative concept, and the “truths” irrealists cite in their premises are not. Indeed, what makes irrealism compelling, if it is compelling, is just that normativity is *not* entailed simply by the “truths” irrealists cite, but by an agent’s own attitudes. But it’s beyond belief itself to think *with any coherence* that such truth need not matter at all to the standards of evaluation as such. If belief is simply thinking something true, and if knowledge involves correctly thinking of something as true *for good reasons*, it is hard to provide an account of normative knowledge according to which there is no objective fact as to whether knowledge so construed is itself a true account.⁵² When we give up on objectivity, we risk forgoing our epistemic concepts of belief, justification, warrant, knowledge and so on. And there simply isn’t any way we can coherently give up these

⁵² Although I think Street’s and Hazlett’s irrealist epistemologies collapse into this sort of account, it’s worth noting that the value nihilism I accuse them of is decidedly unpopular and incredibly rare. Cuneo supposes that Quine endorses this view, but writes “epistemic nihilism is an unpopular position in contemporary philosophy. One is, in fact, hard pressed to find a philosopher who explicitly embraces the view. [...] These positions are of interest not because of their popularity, but because they promise to teach us something very important about a given domain. I want to suggest that epistemic nihilism is a position such as this” (117).

concepts. We know this, if we know anything at all. Calling these very definitions into question, and doubting normativity itself, are just the sorts of things we cannot believe (at least, not for “good reasons”). The denial of objectivity at the level of normativity itself is, as it were, an unbelievable doubt.⁵³

4.4: Meta-Normative Explanations and Moral Facts

I will conclude this chapter with a dilemma for the moral and epistemic irrealist: if irrealism is true, there is no (ultimate) reason to believe that it is true; on the other hand, if truth is ultimately and necessarily our reason for belief, then irrealism is false. In particular, if we take the truth of empirical claims seriously, our reasons for this are straightforwardly epistemic, and objectively so. And if these epistemic reasons are objective, normativity of this sort makes it the case that there are better or worse answers to epistemic problems. More specifically, the irrealist solution to epistemic problems in morality are only convincing insofar as there are ultimate reasons to believe something or the other about morality. If there are no ultimate reasons of this sort, there is no ultimate reason to abandon our moral beliefs.

Recall that the empirical problem for objective moral knowledge called either 1) the justification or 2) the objectivity of our moral beliefs into question. Again, this because the

⁵³ Again, this claim is formally analogous to Descartes’ skeptical discovery: that reasons for belief are real is just as certain as that we ourselves are real. We can *say* that we do not exist, or we can *say* that we do not believe, but we cannot really *believe* this. If you doubt your existence, you believe you exist. And if you doubt your reasons to believe that you exist, you are saying that it is true that there are no reasons, and you therefore take the truth of this as the reason to think that there are no reasons at all. But then, of course, the truth of the statement “there are no reasons” is just the reason for this belief. You would think it’s true if you “believe” it. But you cannot *correctly* believe it if it’s not true. And so, there are reasons, and I exist, and so on.

truth of these beliefs are not part of the causal or scientific explanation the beliefs themselves. The commitment to truth, then, is a background assumption to solutions to empirical problems, including those of the irrealist variety examined so far. That is, in addition to the claim that our moral beliefs do not track truth is an implicit claim that our moral beliefs *should* track truth. And if it is the case that biological processes, rather than the truth of the matter, causes our moral beliefs, and if it is also the case that *beliefs that do not track truth are bad beliefs*, then it is also the case that we should give up our moral beliefs *since* they are insensitive to truth.⁵⁴

But if there is no fundamentally correct aim of belief itself, and if our intuitions about what to believe are biologically inevitable, there is no point in pointing out that a certain class of beliefs do not track truth, for it is far from obvious on an irrealist account that these beliefs should *therefore* be given up as a result. But Street argues that moral objectivism is false and states that this is a *reason* against *believing* objectivism. But if truth is *not* a reason for belief *as an ultimate matter of fact*, there is no ultimate reason for anyone to accept this conclusion. In other words, in dismissing an argument against objective moral knowledge, the stubborn (dis)believer disagrees with the understanding of belief that the argument employs. Neither Joyce nor Street think such ignorance constitutes a sufficient response to the considerations they raise: rather, rejecting the rules their arguments use also discredits

⁵⁴ This assumes that “if I’m justified in believing that I’d have believed *p* even if it had been false, then I’m not justified in believing *p*”. Roger White calls this principle TRUTH SENSITIVITY. But White challenges this principle and its close cousin EXPLANATORY RELEVANCE (“if I’m justified in believing that the ultimate explanation of my belief that *p* makes no reference to whether *p*, then my belief that *p* is not justified”). According to White, most of our beliefs are not justified according to either principle, in which case both Street and Joyce may have doubted far *too little*. If White is right, however, even though it is the case that beliefs that track truth reliably are correct, some beliefs that don’t are also warranted, including perhaps moral beliefs (2010: 9).

all the beliefs such rules are responsible for. Such rules *guarantee* that believers get some things right, and Joyce, Street, or any other philosopher appealing to any rules at all do so insofar as they believe something or the other *is* right, and they think the fact that it is right makes it correct to believe just this.

And this is just the case we find ourselves in with respect to morality: we find ourselves with reason to believe some things, and not others, based on our reasons for thinking certain beliefs are *correct*. Recall that in the case of epistemic problems in morality, Joyce argued that it is an epistemic fact that we do not know whether *p*, since we do not have good reasons to think that *p* (§2.3). Again, rather than tracking an indubitable truth as to whether <killing children is wrong>, for instance, the biological causes of beliefs such as <killing children is wrong> increase fitness outcomes rather than truth. Unfortunately, however, it isn't the case that <adaptive beliefs are correct beliefs> unless there is an argument that links this belief to the premise that true beliefs are better than false beliefs (earlier: <true beliefs are correct beliefs>).⁵⁵ And so, Joyce doubts moral facts because he believes in objective epistemic facts.

Street, on the other hand, has doubted objective epistemic facts. I have just argued that hers is an unbelievable doubt. Even if it isn't, however, her rejection of objective morality is only believable if her epistemic reasons are believable. And if her epistemic reasons are not objective, it isn't clear that her particular epistemic reasons are believable for us even if they are true. After all, the *belief* that objective morality does not exist is ultimately and only required for those whose attitudes endorse this claim, whatever its truth. Hence the

⁵⁵ See Wielenberg (2011) for arguments to this end.

dilemma: either epistemic reasons are objective, and epistemic irrealism is false, or epistemic reasons are subjective, in which case they are not good (enough) reasons for the rest of us to take Street's moral irrealism seriously.

In other words, according to the first horn of the dilemma, conceding the irrealists their epistemic point only weakens their arguments in other domains (including but not limited to morality). So long as irrealists are in the business of telling others what to believe, their background views substantially weaken their case for convincing others of any particular beliefs, including those that make up the empirical problem in meta-ethics.⁵⁶ Ultimately, the moral objectivist will only concede moral irrealism or subjectivism for epistemic reasons if these reasons themselves are objective. So if the meta-normative empirical problem implies irrealism, the irrealist solution to the empirical problem in meta-ethics is ultimately useless. On the other hand, if there is an empirical problem for moral beliefs, this is because there are objective epistemic facts about what to believe in this case.

I think moral skeptics of all sorts should accept these broad implications at face value. After all, if any meta-normative approach leads to the same conclusions, skeptics need not provide a meta-normative account to make their point. This amounts to a concise suggestion for skeptics, subjectivists, and nihilists across their domains of (dis)interest. The approach I employ here urges against arguing for particular normative considerations on

⁵⁶ That said, I do not take our epistemic intuitions as authoritative, in and of themselves. I only take issue with the irrealist assumption that their background view is obvious or plausible to others. If this is false—if I am the only one who doesn't already buy into irrealism in the first place (say, without realizing it)—I do not deny that the irrealists will convince everyone else. I deny instead that the closet irrealists themselves have ultimate reasons, *according to irrealism*, to believe it. At any rate, these are hardly reasons for *me* to accept irrealism. Also, were I to accept irrealism because I *want* to, this hardly shows that irrealism is therefore true.

anything less than the most plausible meta-normative account. I argue that objectivism provides the most plausible meta-normative account, and so the most serious doubts for the truth of our moral beliefs.

5

Skepticism and objective normative facts

5.0: Introduction

In the last chapter, I compared an objective and irrealist epistemic account. In making this comparison, I argued that skeptical arguments in domains like morality require objective epistemic facts. I concluded that epistemic objectivism gives us more compelling reason to doubt some moral beliefs than does an irrealist solution to the empirical problem. This is because objectivism guarantees that our beliefs are correct iff true, and if our moral beliefs are not true, then they are not correct, as an ultimate matter of objective epistemic fact. In this chapter, I will argue that if epistemic objectivism is true in the way I've described, then there are also good reasons to think that there are objective moral facts.

The argument is hypothetically syllogistic: if there are no *ultimate* reasons to doubt objective epistemic facts, arguments against objective moral facts are incomplete. These arguments are incomplete if they fail to demonstrate that moral facts are uniquely objectionable. For instance, I've already outlined how epistemic facts may be very real, even if they play no obvious role in scientific explanation. Therefore, their absence from our best scientific explanations was no good reason to doubt their reality. Likewise, in §5.1, I argue that the absence of moral facts from our scientific explanations of moral belief is

consistent with their reality, although it is some reason to think that our particular moral beliefs are not (yet) warranted. Finally, I conclude in §5.2 that the reality of values is compatible with our many mistakes concerning what they are, and an apparent lack of moral knowledge need not demonstrate the irreality of morals. The solution I suggest embraces objectivity in a unified normative account about our epistemic and moral obligations. A complex understanding of normative evaluations reveals basic normative facts according to which we are more or less right or wrong in recognizing or ignoring good or bad evidence for true or false beliefs. In these evaluations, I think the distinction between practical and theoretical reasons is nuanced: after all, our real, everyday situations require careful epistemic and moral evaluation of persons with better or worse characters. In short, if we are to make sufficient sense of what to think and do, we cannot rule out epistemic and moral facts. The thesis, distilled in this conclusion, is simply that we should expect real value—objective normativity, as it were—to evade easy understanding. I conclude: there are real moral and epistemic reasons.

5.1: Normative Explanations and Objective Moral Facts

So far, I've disagreed with Street as to whether the apparent causal or explanatory absence of moral and epistemic facts from natural or non-normative explanations are reasons to doubt these facts. In this chapter, I will agree with Street that normative truths are unified: for better or for worse, moral and epistemic approaches either stand together or fall together. Of course, since I think there are objective epistemic facts, I also think that there must be objective moral facts. Here, I will not rehearse my arguments against irrealist theories of normative truth. My point above was that these theories do not provide any

ultimate reason to believe anything particular about anything at all, including morality. The alternative that at least some normativity is objective is enough to introduce and then defend the possibility that such objective normativity is inclusive of morality.

In fact, I take it that epistemic objectivity, if real, makes real moral objectivity more plausible. A corollary of my agreement with Street is my disagreement with Joyce over whether there are moral facts at all. Instead, I will argue that if there is at least one objective epistemic fact, then there is at least one objective moral fact.

In arguing for this conclusion, I will borrow a crucial hypothetical premise from Street's arguments against normative facts. In particular, I borrow the first of the following premises implicit to Street's argument:

- (1) If moral facts are not objective, then epistemic facts are not objective.⁵⁷
- (2) (From 1) Normativity is unified across the practical and epistemic domains (Street 2011, 213).⁵⁸
- (3) Moral facts are not objective (moral irrealism)
- (4) Epistemic facts are not objective (From 1 and 3) (epistemic irrealism)

⁵⁷ Street does not always argue for this premise, but it is implicit in much of her work (2006; 2008b; 2009; 2011; 2016). For instance, Street agrees with Dworkin and Korsgaard that there is no hard and fast distinction between the theoretical and practical domains of normativity (2016). Here, I agree: our practical and epistemic reasons overlap, so an argument against the objectivity of at least some of these reasons is *ipso facto* an argument against the objectivity of normative reasons in general. (Hence, this chapter).

⁵⁸ Indeed, it is an open question whether there is a principled distinction between epistemic and practical normativity, for there are at least some cases where our intuitions take practical reasons to count for belief, in quite the same way moral normativity does (Cuneo makes this argument, and I will conclude this section with just this point). Allan Hazlett, an epistemic irrealist, also doubts the distinction in several cases, writing: "We should not (from the start) identify the 'epistemic' with the theoretical, nor with the intellectual, nor with the cognitive, nor with the doxastic..." (2013: 28). For irrealists, it is often the practical nature of epistemic reasons that counts *against* their reality. However, if such reasons are real, practical normativity may very well be real, leaving such irrealists with even less to say against the reality of moral reasons in particular.

∴ Normative objectivism is false.

Following Cuneo, call (1) the “parity premise” (89). Cuneo has many reasons for accepting the parity premise, but the one which concerns Street and myself is his fourth argument from *The Normative Web*, or what I have been calling an “empirical argument”. Says Cuneo:

This argument starts with the observation that we don’t have any explanatorily informative story about how we could gain access to moral facts. [...] But I take it to be fairly plain that the argument poses an equally serious challenge to epistemic realism. For if it is the normative nature of moral facts that is supposed to debar us from an informative story about how we grasp them, then the same holds for epistemic facts. And if it is the lack of an informative story of how we grasp moral facts that implies that we ought not to admit them into our ontology, then the same holds for epistemic facts. Considerations of the same sort counsel against the acceptance of the existence of both moral and epistemic facts. (101)

Cuneo’s argument for the parity premise, in other words, grants that we do not know how moral facts explain our moral beliefs. Similarly, epistemic facts are not relevant to scientific explanations of our epistemic beliefs. So, if science does not require moral facts, and if this feature of such “facts” counts against them, this scientific dispensability counts against epistemic facts as well. As Cuneo puts it, either epistemic facts are objectionable for the same reasons moral facts are objectionable, or they are not objectionable, and moral facts are also not objectionable. Presumably, if there are no moral facts, it is because they have features that do not figure into the right ontology of our world, and “if nothing has the objectionable features, then epistemic facts do not exist” (225).

Street believes that nothing has the objectionable features, since moral and epistemic reasons are subjective rather than objectively factive. Again, Street argues that normative principles (coherence, parsimony, and so on) are consistent with the irreality of the moral

facts the subjectivist disputes. QED, objective moral facts are not worth believing. But nothing, I submit, could be further from the truth. The normative principles that motivate the empirical problem are best understood as either objectively true or false. As it happens, I think that if some epistemic principles can survive the meta-normative empirical problem, we may *therefore* think that some moral facts are also immune to the empirical problem. In fact, by contraposition, I agree with the ‘parity premise’ Street employs:

(1) If there are objective epistemic facts, there are objective moral facts.

Of course, this denies the antecedent of Street’s parity premise—it is equivalent to and can be derived from: “if there are no objective moral facts, there are no objective epistemic facts” (Cuneo 231). (Note that, so far, the parity premise permits moral facts without epistemic facts, and is distinct from (2) in this way.) Again, Street explicitly states that her reasons for doubting moral facts are also arguments against normative facts in general, and in this, Cuneo and I think her commitments are exactly right. Simply put, irrealism about normative reasons in the moral domain correctly embraces irrealism about normative reasons across the board. But in Chapter 4, I argued that irrealism about normative reasons is incoherent even if granted, and so incorrect. Normative reasons in one domain make possible normative reasons in other domains. It follows that if some epistemic facts are objective, moral facts must be objective as well.

Now, I take it that since Street’s own argument depends on the parity premise, she will not be willing to pull apart moral and epistemic facts in order to isolate only one or the other as problematic (for Street, both are objectionable, for quite the same reasons). The obvious rejoinder from Mackie or Joyce, however, is to deny the parity premise and argue that

moral facts, unlike epistemic facts, are *uniquely* objectionable. Again, Mackie and Joyce do not see how moral facts explain our moral beliefs—they are not required in an explanation for why we take ourselves to have moral reasons. The idea, again, is that “[moral] categorical reasons are naturalistically unacceptable because they fail to pull their explanatory weight in the sciences”, whereas epistemic reasons, presumably, are scientifically necessary (Cuneo 219). Against this, it is instructive to rehearse Street’s argument at length:

Whenever we make a normative judgment, we regard that judgment as true. [...] On the other hand, we view our normative judgments as subject to causal explanation, just like any other psychological state. The general point about a dual standpoint on our normative judgments of course holds just as much with regard to judgments about reasons for belief as it does for any other type of normative judgment. On the other hand, we – and certainly the realist – think it’s true that considerations having to do with the truth of *P*, rather than other kinds of considerations, are normative reasons for and against the attitude of belief that *P*. Yet we may also view those very same views on reasons for belief as having been shaped by causal forces [...] This coincidence between normative truth, on the one hand, and what causes led us to think about normative matters, on the other, demands an explanation. The problem is that the realist about normativity has no acceptable one (2011, 233).

If Mackie and Joyce take the cold, hard world of science, so to speak, to count against any entities that are not necessary in scientific explanations of the world, including the psychological states it includes, then they must deny objective epistemic facts. But I argued, in Chapter 4, that the scientific dispensability of such facts need not count against them. I also developed an account in Chapter 2 according to which error theorists like Mackie and Joyce think that epistemic facts are real, and I think the considerations of the last chapter support the intuition Mackie and Joyce share about epistemic facts. Such epistemic objectivity also counts *against* the error theoretic intuition about moral facts, because such

facts might survive an explanation of how psychological beings like us believe and behave if epistemic facts are indeed real, for Mackie or Joyce, on an error theoretic account. So, I urge the error theorist to take their own acceptance of epistemic facts to provide *prima facie* reason to accept moral facts as well, and give up error theory instead.

It is worth recalling that I also argued, in Chapter 2, that Joyce *does* take just this sort of modest position on moral facts, at times, for some of the reasons I outline above. That is, Joyce's most recent response to the empirical problem does not rule out moral facts at all, but merely takes the science to complicate our beliefs, regardless of whatever the moral facts may or may not be. Of course, this position will not satisfy the objectivist about moral knowledge, but it does provide common ground on which such objectivists can agree with a moral skeptic. The moral skeptics and committed objectivists will agree, for example, that there are epistemic facts according to which at least *some* beliefs in moral facts are not justified. The dispute between the two is over *which* moral beliefs to doubt.

In this, it is useful to consider whether there are any ultimate moral facts about us that closely parallel the ultimate epistemic fact I argued for in Chapter 4. I argued then that the belief <true beliefs are correct beliefs> corresponds to the objective normative property specified in the content of the belief itself (its being correct). One basic moral fact may just be that good actions are correct actions, or actions we ought to do. But belief in this fact would not demonstrate its truth in the way that belief in the normativity of true beliefs *must* be true in virtue of this very belief.⁵⁹ Worse, while the concept of (true) belief is fairly

⁵⁹ That said, according to Kant and his interpreters (esp. Korsgaard), reflection on the nature of agency and action reveal normative standards of right action in the way I propose that reflection on the nature of rationality and belief reveal standards of correct belief. For example, if Kant is correct, moral facts may be

uncontroversial, it is not clear at all what counts as a “good” action.⁶⁰ But once we are open to the possibility that there are normative facts in the moral domain, it must be true that good actions, whatever they are, are generally required of us.

The moral of this thesis is therefore a moral about belief: to wit, that our normative evaluations of true beliefs and actions require epistemic and moral facts. At least, if they require one, they require the other. This is just what the parity premise proposes. Either there are normative facts, or there are not. And if there are no normative facts, there are no epistemic facts, and Chapter 4 was wrong. Conversely, for those sympathetic to the arguments of Chapter 4, there are also reasons to believe moral facts as well.

Further, normative evaluations in general require considerations that do not fall neatly into either one or the other domain—after all, “some epistemic reasons are practical in nature, favoring the implementation of action plans of certain kinds. And some moral reasons are theoretical in character, favoring the formation, maintenance, and surrendering of propositional attitudes of certain kinds” (Cuneo 74). This means that the theoretical and practical domains of evaluation may very well overlap in some cases. If this is true, Mackie and Joyce are wrong, and the following premise is true:

(2) Normative objectivism is unified across the practical and epistemic domains.

derivative on some ultimate or basic conceptual fact: i.e. the concept of a *good will* itself (56). For the record, I sympathize with Kantian approaches, but won’t *assume* their points in the following discussion.

⁶⁰ Note, however, that we do not need an exact theory of goodness to evaluate particular claims about goodness. Here, I am in agreement with Allan Hazlett, who argues that “it’s wrong, in general, that particular claims about *x* can’t be evaluated until a philosophical account of *x* has been given” (2013: 5). For instance, says Hazlett, “you don’t need to settle for a theory of wellbeing to know that your life would go worse if, while everything else remains the same, you were to be tortured every day for the rest of your life” (2013: 5).

Not only does this premise suggest that the normativity of practical reasons and epistemic reasons are unified, it is inclusive of at least some epistemic reasons that are practical in quite the same way our moral reasons are. So the above premise permits this complexity of our normative evaluations.⁶¹ But even if this premise is false, there are still reasons to accept moral truth. After all, whether or not there are both epistemic and moral facts (or neither), that they are unified according to the parity premise has important consequences.

Consider the pair:

- (1) If there are objective epistemic facts, there are objective moral facts.
- (2) Objective normativity is unified across the practical and epistemic domains.

Irrealists like Street and Hazlett argue for and accept the above premises. Mackie and Joyce do not (at least, in the latter case, not always). But Mackie and Joyce agree with realists like FitzPatrick and Cuneo and Shah on the following:

- (3) There is at least one objective epistemic fact.

⁶¹ There is an important objection lurking in the background here, and it would be remiss of me to leave it out altogether. The objection runs thus: if true beliefs are correct beliefs, but if there are practical reasons for belief, is it not possible that some false beliefs are also correct beliefs? There are many reasons that I doubt that false beliefs are correct beliefs *as a rule*: I do not think that there are patterns of false believing that are good or correct, all things considered. But there may be exceptions to the rule of correct believing—ways in which we may be excused from believing particular truths. For example, it may be possible to count the blades of grass on my lawn, and so form a true belief about whether the number of blades on my lawn is odd or even. But the time I spend studying the grass is time I could be using to believe other, typically more important or valuable things, instead. So it is not required of us to believe *any* truth at all, since some truths are more valuable to believe than others. And I take it that practical normativity has some role to play in evaluating these true beliefs vis-à-vis each other, according to at least some instrumental considerations. As the objection proposes, such instrumental or practical reasons might count in favour of some false belief, and although I cannot call these false beliefs “correct beliefs” per se, I will also grant that these (generically incorrect) beliefs are *sometimes* recommended by practical reason, albeit not conclusively (e.g. that I will have a happy marriage, that I will survive my cancer, that there is an after-life etc.).

If (1), (2) and (3) are true—if Street and Hazlett and Cuneo are right about the parity premise, and Mackie and Joyce and Shah and FitzPatrick are right about (3), the following *must* hold:

(4) There is at least one objective moral fact.

Street and Joyce dispute (4). My complaint against the former was the truth of (3). My complaint against the latter was the truth of (1) (which is sufficient, I think, for (2)). Here, I merely want to point out that the truth of (1) through (4) does nothing to guarantee objective moral knowledge. It guarantees objective normative facts, instead. Hence:

∴ Normative objectivism is true.

Street rejects this conclusion. Now, I may be wrong about the truth of (1) through (4), but my argument is deductively valid, and it formally clarifies the epistemic puzzle about morality even if false. The puzzle is just that (4) is so difficult to defend, even though it must be true if (1) through (3) are true. Worse, (1) through (4) are compatible with a thoroughgoing skepticism about objective moral knowledge. Worse still, moral irrealism also trivially entails such skepticism, and so the two are easily conflated, even though doubting objective moral knowledge is not sufficient for doubting objective normativity. This is also why establishing objective normativity is not sufficient for moral knowledge.

I think many of us already accept the upshot of all this: a thoroughgoing skepticism about objective moral knowledge is unconvincing if (3) is not true.⁶² The most convincing forms

⁶² It is, at any rate, unconvincing according to the standard analysis I've defended throughout this thesis. According to the relevant scientific theories, most of us are in the grips of said "standard analysis" (see footnote 1), and an argument that rejects the apparent reality of normative reasons needs to do more than demonstrate that a scientific account is itself a sufficient explanation of the analysis as such. Instead, irrealist

of skepticism are ones that accept premises (1) through (4). For example, while Joyce has independent arguments against the reality of moral facts, he has recently admitted that his argument does not refute (4), but instead grants (1) through (4) (2016). This strikes me as a more sensible argument than denying (4) while granting the parity premise. After all, if the parity premise is true, a denial of (4) is a denial of (3), and vice versa. And if I am right, a denial of (3) makes an argument for moral skepticism, or anything at all, incoherent.

So there must be at least one objective epistemic fact if moral skepticism makes sense.

Moral skepticism makes sense in at least some circumstances, according to the reality of epistemic reasons, and this is what we would expect if epistemology is objective. Likewise, the reality of moral reasons makes moral skepticism intractable, and this is what we would expect if morality is objective. In conclusion, we cannot coherently doubt objective reasons altogether. To doubt the reality of reasons altogether is to assume real reasons to think that it is true that there are no reasons. It is not possible to coherently believe this. It is an unbelievable doubt.

5.2: Conclusion

The conclusion of this thesis is compatible with the skeptic's conclusion against moral knowledge. It is entirely possible that we do not have moral knowledge.

Importantly, if moral facts are objective, then objective moral knowledge is only possible if there are such facts. I have provided a theoretical defense of objective normativity, I cannot guarantee that the details of this picture are such that we turn

arguments need to provide reasons for thinking that (3) is not true that are not question-begging—reasons that do not already assume the conjunction of science and an antecedent irrealist account of reasons.

out to know moral facts. If I am right, whether we know moral facts and what this knowledge amounts to ultimately depends on the beliefs of real people in real circumstances, rather than an abstract, idealized theory according to which they might know moral truths, were they correctly situated in the world of objective normative facts.

For what it's worth, I also suspect that at least some of our normative beliefs, including moral beliefs, are true. And even if we don't have moral knowledge, this is because our moral beliefs are, at minimum, false, or unjustified, according either to the relevant objective moral or epistemic facts, or both (Shafer-Landau 2004: 57). If our moral beliefs are false, it is not because the relevant facts are not objective. On the other hand, that our beliefs are unjustified is an objective epistemic fact. None of this is terribly new—this “standard analysis” of moral knowledge is quite similar to those of centuries past (Plato 309). However, the science of the last two centuries does pose a significant epistemological puzzle for this analysis. I admit that I have not solved this puzzle. My more modest aim, instead, has been to clarify the puzzle in two important respects.

First, in the first two chapters, I emphasized a classical distinction between a belief that is false and a belief that is not justified. It is sometimes very difficult to show that beliefs are not false, and depending on the epistemic facts of the matter, beliefs like this may not be justified as a result. However, that we cannot show that some belief is not false does not mean the belief is not true. Although moral skeptics appreciate this distinction, it does not always play an obvious role in their arguments. Sharon Street takes the difficulty of proving objective moral beliefs to show that these beliefs are a mistake. In this thesis, I have argued that this approach conflates beliefs without warrant with beliefs that are not true. It is true

that some of our moral beliefs are not justified. It is false that they are *therefore* false, if they are objective.

Second, in the second part of the thesis, I argued that it is *objectively* true that some of our moral beliefs are not justified. Richard Joyce agrees with this conclusion. Sharon Street does not. Street does not think there are objective epistemic facts of any sort at all. Street also does not think that our moral beliefs are not ever warranted—this is because an irrealist account can justify our moral beliefs instead. In clarifying these epistemic and metaphysical commitments, I show that Joyce and Street challenge objective moral knowledge for very different reasons. Ultimately the responses are individually inadequate, but taken together, they substantially clarify the empirical problem. Joyce was right to doubt moral truth, and Street was right to see this doubt as evidence of some epistemic, and therefore, moral facts. There are objective epistemic facts, and this exactly why we have reason to doubt our moral beliefs, even though we cannot decisively rule out objective moral facts. The implications of Street's irrealism are not coherently believable. The correct skeptical discovery in this case is simply that moral facts are objective as well.

This is why I have concluded that irrealism does not solve the empirical problem. If we think true beliefs are correct, and if we do not know that our moral beliefs are true, we do not know that they are correct. If we do not think that true beliefs are generally, objectively correct, we have no ultimate reason to count the lack of evidence or facts against beliefs. Put simply, without objective normativity, skeptics have no ultimate reason to doubt anything in particular, moral or otherwise. Thus, on pain of incoherence, moral skeptics should doubt the truth, rather than the objective normativity, of moral belief.

Bibliography

- Aristotle. *The Metaphysics*. Trans. Hugh Lawson-Tancred. London: Penguin Books, 2004. Print.
- Chalmers, David J. *The Character of Consciousness*. Oxford: Oxford UP, 2010. Print.
- Clarke-Doane, Justin. "Morality and Mathematics: The Evolutionary Challenge." *Ethics* 122.2 (2012): 313-40. Web.
- Copp, David. "Darwinian Skepticism about Moral Realism." *Philosophical Issues* 18.1 (2008): 186-206. Print.
- . "Explaining Normativity." *Proceeding and Addresses of the American Philosophical Association* 89.1 (2015): 48-73. Print.
- . *Morality in a Natural World: Selected Essays in Metaethics*. Cambridge: Cambridge UP, 2007. Print.
- Cuneo, Terence. *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press, 2007. Print
- Das, Ramon. "Evolutionary Debunking of Morality: Epistemological or Metaphysical?" *Philosophical Studies* 173.2 (2015): 417-35. Web.
- Darwall, Stephen L, Allan Gibbard, and Peter A. Railton. *Moral Discourse and Practice: Some Philosophical Approaches*. New York: Oxford University Press, 1997. Print.
- Descartes, René. *Meditationes De Prima Philosophia / Meditations on First Philosophy*. (G. Heffernan, Trans.) Notre Dame, IN: University of Notre Dame, 1990. Print.
- Dworkin, Ronald. "Objectivity and Truth: You'd Better Believe It." *Philosophy & Public Affairs* 25.2 (1996): 87-139. Web.
- Enoch, David. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford UP, 2011. Print.
- . "The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It." *Philosophical Studies* 148.3 (2009): 413-38.

- Evers, Daan. "Street on Evolution and the Normativity of Epistemic Reasons." *Synthese* 192.11 (2015): 3663-676. Web.
- Fitzpatrick, William J. "Debunking Evolutionary Debunking of Ethical Realism." *Philosophical Studies* 172.4 (2014): 883-904. Web.
- . "Why There Is No Darwinian Dilemma for Ethical Realism." *Challenges to Moral and Religious Belief: Disagreement and Evolution*. By Michael Bergmann and Patrick Kain. Oxford, UK: Oxford UP, 2014. 237-55. Print.
- Frege, Gottlob. *The Foundations of Arithmetic; a logico-mathematical enquiry into the concept of number*. 2nd ed. Oxford: Blackwell, 1953. Print.
- Harman, Gilbert. *The Nature of Morality: An Introduction to Ethics*. New York: Oxford UP, 1997. Print.
- Hanson, Louise "The Real Problem with Evolutionary Debunking Arguments." 0.0 *Philosophical Quarterly* (2016): 1-26. Web.
- Hazlett, Allan. *A Luxury of the Understanding: On the Value of True Belief*. Oxford: Oxford UP, 2013. Print.
- . *A Critical Introduction to Skepticism*. London: Bloomsbury, 2014. Print.
- Hume, David. *An Enquiry Concerning Human Understanding*. Ed. Lorne Falkenstein. Peterborough, Ont: Boradview Press, 2011. Print.
- Hoffman, Donald D. "The Interface Theory of Perception: Natural Selection Drives True Perception to Swift Extinction." *Object Categorization: Computer and Human Vision Perspectives*. Ed. Sven Dickinson, Ales Leonardis, and Bernt Schiele, and Michael Tarr. Cambridge, UK: Cambridge UP, 2009. 148-265. Print.
- Joyce, Richard. "Evolution, Truth-Tracking, and Moral Skepticism." *Essays in Moral Skepticism* (2016): 142-58. Web.
- . "Irrealism and the Genealogy of Morals." *Irrealism in Ethics* (2014): 1-22. Web.
- . "Moral Anti-Realism", *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition) Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/win2016/entries/moral->

- anti-realism/. Accessed 30 Sept. 2016.
- . "Taking Moral Skepticism Seriously." *Philosophical Studies* 168.3 (2013): 843-51. Web.
- . *The Evolution of Morality*. Cambridge & London: MIT Press, 2006. Print.
- . *The Myth of Morality*. Cambridge, UK: Cambridge UP, 2001. Print.
- Kahane, Guy. "Evolutionary Debunking Arguments." *Noûs* 45.1 (2010): 103-25. Web.
- Kant, Immanuel. *Groundwork for the Metaphysics of Morals*. Ed. Lara Denis. Peterborough, Ont: Broadview Press, 2005. Print.
- Kitcher, Philip. 2006a, "Biology and Ethics," in Copp, D. ed., *The Oxford Handbook of Ethical Theory*, Oxford: Oxford University Press, pp. 163-8
- Korsgaard, Christine M. *The Sources of Normativity*. Cambridge: Cambridge UP, 1996. Print.
- Locke, John. *An Essay Concerning Human Understanding*. Ed. Peter H. Niddich. Oxford: Clarendon Press, 1975. Print.
- Mackie, J. L. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin, 1977. Print.
- Moore, G.E. *Principia Ethica*. Cambridge: Cambridge University Press, 1965. Print.
- Nozick, Robert. *Philosophical Explanations*. Cambridge, MA: The Belknap Press of Harvard UP, 1981. Print.
- Plantinga, Alvin. *Warrant and Proper Function*. New York: Oxford UP, 1993. Print.
- Plato. *The Republic*. Trans. H.D.P Lee. Harmondsworth: Penguin, 1974. Print.
- Reid, Thomas. *An Inquiry into the Human Mind: On the Principles of Common Sense*. Ed. Derek R. Brookes. University Park, PA: Pennsylvania State University Press, 1997. Print.
- Ruse, Michael. *Darwinism and its Discontents*. New York: Cambridge University Press, 2006.
- Shah, Nishi. "Can reasons for belief be debunked?" *Reasons for Belief*. Eds. Andrew Reisner and Asbjørn Steglich-Petersen. New York: Cambridge UP, 2011. 94-111. Print.
- . "The Normativity of Belief and Self-Fulfilling Normative Beliefs." *Canadian Journal of Philosophy* 39.Sup1 (2009): 189-212. Print.
- Shafer-Landau, Russ. "Evolutionary Debunking, Moral Realism, and Moral Knowledge."

- Journal of Ethics and Social Philosophy* 7.1 (2012): 1-37. Web.
- . *Whatever Happened to Good and Evil?* New York: Oxford UP, 2004. Print.
- Sinnott-Armstrong, Walter. *Moral Skepticisms*. Oxford: Oxford UP, 2006. Print.
- . "Moral Skepticism", *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition), Edward N. Zalta (ed.), <http://plato.stanford.edu/archives/fall2015/entries/skepticism-moral/>. Accessed 30 Sept. 2016
- Street, Sharon. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127.1 (2006): 109-166.
- . "Coming to Terms with Contingency: Humean Constructivism about Practical Reason." *Constructivism in Practical Philosophy* (2012): 40-59. Web.
- . "Constructivism about Reasons." *Oxford Studies in Metaethics* 3 (2008): 207-45. Web.
- . "Evolution and the Normativity of Epistemic Reasons." *Canadian Journal of Philosophy* 39.Sup1 (2009): 213-48. Print.
- . "Objectivity and Truth: You'd Better Re-think It." *Oxford Studies in Metaethics* 11 (2016): 293-335. Print.
- . "What Is Constructivism in Ethics and Metaethics?" *Philosophy Compass* 5.5 (2010): 363-84. Web.
- Sturgeon, Nicholas. "Moral Explanations." *Essays on Moral Realism*. Ed. Geoffrey Sayde-McCord. Cornell: Cornell UP, 1988. 229-81. Print.
- White, Roger. "You Just Believe That Because...." *Philosophical Perspectives* 24.1
- Wielenberg, Erik J. "On the Evolutionary Debunking of Morality." *Ethics* 120.3 (2010): 441-64. Web.