

KNOWING RIGHT FROM WRONG

KIERAN SETIYA

OXFORD
UNIVERSITY PRESS

CONTENTS

<i>Preface</i>	vii
Introduction	i
1. Plan of this Book	3
2. Ethical Supervenience	8
1. Disagreement	12
1. Equal Weight	15
2. Intuitions, Coherence, Reflective Equilibrium	23
3. Bias towards the Truth	39
2. Reliability	56
1. Ethical Judgement	57
2. Coincidence	66
3. Begging the Question	76
3. Knowledge	85
1. No Accident	89
2. Reflection, Induction, Perception	99
3. The Problem of Ethical Knowledge	111

4. Human Nature	128
1. Knowledge without Convergence	129
2. Justification	138
3. Faith	142
<i>Bibliography</i>	159
<i>Index</i>	169

PREFACE

Though it is not a sequel, this book takes up questions that were not addressed in *Reasons without Rationalism*, and that its argument makes urgent. In that book, I answered a certain sort of ethical sceptic, one who grants that justice and benevolence are virtues, but who asks why he should care. Why bother to act as a virtuous person would? The answer I gave is that the 'should' of reasons cannot be detached from ethical virtue. Its standards are those of good character, applied to practical thought. It follows that, if justice and benevolence are virtues of character, they are responsive to considerations that therefore count as reasons to act. One form of the question 'Why be moral?' is misconceived.

This argument is controversial. Even if it works, however, its limitations are clear. It does not answer a sceptic who grants the authority of virtue but still asks 'Why be moral?' Why believe that justice and benevolence are virtues at all? If they are not, we can act as a virtuous person would while disregarding so-called 'moral virtue' as a sham. Worse yet, there are sceptics who argue that, whatever the truth about virtue and practical reason, our beliefs about them are never justified. In light of disagreement, or problems of reliability and luck, our claims to ethical knowledge are flawed.

It is sceptics of these other kinds that I address in this book. Its arguments are self-contained: they do not rest on my previous work. But there are common themes. One is a denial of normative independence: the standards of reason are not distinct from those of virtue; nor are the standards of epistemology independent

of ethical fact. In each case, this prevents a sceptical question from being asked. Along with this parallel, there is the pervasive influence of Aristotle and Hume, and the provocation of Bernard Williams, first through 'Internal and External Reasons', then through *Ethics and the Limits of Philosophy*.

For reactions to this material in earlier forms, I am grateful to audiences at Brown University, the University of Chicago, Columbia University, the University of Illinois, Urbana-Champaign, Johns Hopkins University, the University of Miami, the University of Missouri, St. Louis, the University of Nebraska, the New York Institute of Philosophy, Oxford University, the University of Pittsburgh, Princeton University, Rutgers University, the University of Tennessee, the University of Texas, and the University of Wisconsin, Milwaukee; to participants in graduate seminars on ethical realism and on moral theory at the University of Pittsburgh; and to Ori Beck, Paul Benacerraf, John Broome, David Christensen, Brad Cokelet, Jonathan Dancy, Cian Dorr, Casey Doyle, Antony Eagle, Adam Elga, Matt Evans, Kit Fine, Marah Gubar, Anil Gupta, Mark Hopwood, Paul Horwich, Anja Jauernig, Tom Kelly, Ben Laurence, Brian Leiter, Peter Lewis, John MacFarlane, Colin Marshall, Adam Marushak, John McDowell, Joe Milburn, Dan Morgan, John Morrison, Jessica Moss, Mike Otsuka, Hille Paakkunainen, Joseph Raz, Mark Richard, Gideon Rosen, Karl Schafer, Joshua Schechter, Stephen Schiffer, Nishi Shah, James Shaw, Michael Slote, David Sosa, Nick Stang, Robert Steel, Greg Strom, Scott Sturgeon, Larry Temkin, Katja Vogt, Brian Weatherson, Ralph Wedgwood, Tim Willenken, Mark Wilson, and three anonymous readers for Oxford University Press. Throughout this process, Peter Momtchiloff has been a marvel of efficiency and editorial wisdom. Thanks to the University of Arkansas Press for permission to reproduce, in chapter one, material from my essay, 'Does Moral Theory Corrupt Youth?' (*Philosophical Topics* 38: 1 [Spring 2010]).

It is not uncommon for authors to confess, in the preface to their books, that their writing was a cause of hardship to their

family, to apologize for this, and to express their gratitude. The years in which this book were written were not easy ones, but their adversity did not come from the fact that I was writing it. For me, something like the opposite: being absorbed in philosophy, in this work, was a source of solace in difficult times. But not the most important source. In words attributed to Rilke, ‘people have misunderstood the position love has in life; they have made it into play and pleasure precisely because they thought that play and pleasure are more blissful than work; but there is nothing happier than work, and love, precisely because it is the supreme happiness, can be nothing other than work’. Marah, this book is for you.

Kieran Setiya

Introduction

One way to prompt doubts about objectivity in ethics is to ask a rhetorical question: if the facts about right and wrong, or what there is reason to do, are independent of my beliefs, how could the truth of those beliefs be anything but luck? Had I been brought up in other ways, at other times, in other places, I would have believed, no less emphatically, quite different things: that it is morally permissible to keep slaves or treat women as property; that justice is a shameful vice; that there is reason to do only what is in one's own best interests, regardless of its effects on others, or to maximize utility, even though some are trampled along the way. Though my beliefs would have been different, the facts would not have changed. They do not depend on me. Instead, my beliefs would have been false. Is it simply my good fortune to believe the truth?

The anxiety elicited here can take quite different forms. Why is it disturbing to suppose that our access to ethical truth depends on luck? One answer cites the epistemic significance of disagreement. I do not hold the terrible views described above, but what could I say to someone who does? Our disagreements are sufficiently deep that any argument I could give would beg the question. I have no independent grounds on which to discount his views, or to treat myself as more reliable. How can I then persist in my convictions? How can I think that I am the one who happens to be right?

We can press the point in other ways. If I am reliable in ethics, there is a systematic correlation of fact and belief: of what is right and wrong, or what there is reason to do, on one side, and my psychology, on the other. If the facts are independent of what I think, how could this correlation be explained? It is a mere coincidence, miraculous, convenient, and, on reflection, quite impossible to believe.¹

Nor do I simply aspire to justification. Much of what I hold in ethics I take to be not merely true, and justified, but known. I do not merely think that slavery is wrong, I know it is. I know that women are not the property of men, that there is reason to care about people other than myself, that one should respect the rights of the innocent even at some cost to the greater good. If I know these things, the truth of my beliefs is not an accident; it can somehow be explained. Yet if the facts are independent of me, what could this explanation be?

This book confronts the threat to objectivity in ethics posed by epistemic luck, teasing out its separate threads. There is the problem of ethical disagreement, which is the topic of chapter one. There is the problem of reliability and coincidence, which is the topic of chapter two. And there is the problem of knowledge and accidental truth, which is the topic of chapter three. In solving these problems, I argue against epistemologies of intuition, coherence, and reflective equilibrium; I argue that the basic standards of epistemology in ethics are biased towards the truth; I argue that there is ethical knowledge only if our beliefs are constitutively bound up with the facts; and I argue that we can accept all this, without implausible predictions of convergence or relativity, through a synthesis of Aristotle and Hume.

¹ This style of argument is inspired by Harry Field (Field 1989: 25–30, 230–9).

I. PLAN OF THIS BOOK

Before giving a more detailed outline of my argument, I introduce some claims I won't attempt to argue for. Although my premises are modest, they are not indisputable: they represent the framework of the book and a limitation of its scope. This is not a comprehensive treatise in the metaphysics and epistemology of ethics, though it has a lot to say about both. It is a response to the most challenging and intractable forms of ethical scepticism.

What do I mean by 'ethics'? I want to be inclusive. There are claims of right and wrong, and of what there is reason to do, but there are practical reasons of other kinds—reasons for wanting, respecting, or admiring things—along with claims of virtue and vice, acting well or badly, moral obligation, justice, benevolence, courage, and the rest. A rule of thumb, to which I return below: if in doubt, it's ethical.

I assume at the outset that we can speak of ethical claims or propositions, as I have done; that they can be true or false; and that we aspire to knowledge in our ethical beliefs. Though they suggest a form of 'cognitivism', these assumptions may be shared by contemporary 'expressivists' and 'quasi-realists'.² They do not presuppose a particular interpretation of truth or belief.

Since we can be wrong in our beliefs, there is scope for ethical disagreement. I assume that this is possible even among those who belong to quite different communities. It is not that we cannot talk past each other, but mostly we do not. Exceptions aside, when we seem to disagree in ethics, we do: at least one of us must be mistaken—though not, perhaps, irrational—in his beliefs. This premise rules out one form of relativism. But it is ecumenical. Some so-called 'relativists' aim to accommodate

² See Blackburn 1993, 1998; Gibbard 1990, 2003. For the same reason, I doubt that the issues raised in Fine 2001 are directly relevant here.

disagreement between different communities and their members; they can accept what I assume.³

A final premise is familiar, but more complex: it is a version of ethical supervenience. The core idea of supervenience is that ethical differences turn on differences of other kinds, that ethical concepts apply in virtue of others. If I acted wrongly and you did not, there must be something to say about why: about the action I performed, its causes or effects, the intention with which it was performed, its content or context, that explains the ethical contrast, facts in virtue of which my action was wrong and yours was not. If there is a reason for me to act which is not a reason for you, there must be a difference between my circumstance and yours from which this ethical difference flows. Likewise for virtue and vice. If you are temperate and I am not, there must be psychological differences between us that underlie this fact. The proper formulation of supervenience involves substantive difficulties, which are taken up below in section 2. What we need for now is the core idea.

With these assumptions, we turn back to our principal theme. This book works through a series of sceptical arguments that appeal to epistemic luck, as it appears in disagreement, in the claim to reliability, and in accidental truth. These arguments question our entitlement to believe, or claim to know, the ethical facts. They do not aim to refute, directly, the existence of such facts.⁴ In answering them, I will ask what must be true about the structure of justification and the nature of ethics if we have ethical knowledge or justified belief. Although I defend these conditions, and believe they can be met, I will not argue for them on independent grounds. The case for my conclusions is

³ I am thinking here of recent work by John MacFarlane (2005, 2007) and Mark Richard (2008: Chs. 4–5).

⁴ In this respect, they differ from the metaphysical version of Mackie's 'argument from queerness', on which there is something problematic in the very idea of an objectively prescriptive fact; see Mackie 1977: 38–42.

that we must accept them if we are to avoid the prospect of ethical scepticism.

Thus, in chapter one, I assume that we should persist in true beliefs even in the face of intractable disagreement. Confronted with those who deny that justice and benevolence are virtues, that there is reason to care about anyone but oneself, the proper response is to hold one's ground. Having clarified this view, I explore its implications for evidence in ethics. One consequence is that the evidence for ethical beliefs does not consist in ethical intuitions; nor is justification a matter of reflective equilibrium or coherence among beliefs. In order to resist the inference from fundamental disagreement to scepticism, we must reject these influential pictures of justification in ethical theory. What we need instead is a conception of justified belief on which the evidence for ethical claims is fundamentally non-ethical: it is evidence for facts on which the truth in ethics supervenes. This view contrasts with a second form of intuitionism, which draws not on intuitions as evidence, but on self-evident principles or non-evidentially justified beliefs. In defending the evidential view, I engage with traditional debates in epistemology, about coherence and foundationalism, and about the relationship of justification to truth.

Despite its scope, chapter one leaves many unanswered questions. To know the evidence for ethical beliefs is not to know when such beliefs are justified, if ever. There is more to justification than evidence. Chapter two advances the discussion in three ways. First, it addresses the objection that forming ethical beliefs on non-ethical grounds is in violation of a Humean constraint: no inferring 'ought' from 'is'.⁵ I argue that the deductive model of justification behind this complaint is generally flawed, and I explore the role of reliability in forming justified beliefs.

⁵ Hume, *Treatise* 3.1.1.27. Whether Hume accepts the constraint is a matter of dispute.

The second task of chapter two is to refute an argument against ‘ethical realism’—roughly, the independence of ethical fact and belief—according to which it makes our reliability in ethics a sheer coincidence.⁶ The correlation of ethical fact and belief is held by the critic to be incredible unless it can be explained, which ethical realism precludes. In response, I argue that we are justified in accepting a coincidence, without the need for explanation, if we have sufficient evidence of its occurrence. According to the theory of chapter one, that is our situation in the ethical case. Finally, I show that there is nothing circular in this account. It does not beg the question to argue for reliability in ethics by way of ethical beliefs.

The sceptical problems considered in the first two chapters turn on misconceptions in epistemology. They can be solved in epistemic terms: by determining what counts as evidence in ethics, and what is involved in begging the question. These solutions have no metaphysical upshot. They are consistent with the irreducibility of ethical facts, their causal impotence, and their constitutive independence of us. In chapter three, we confront a less tractable argument, about the conditions of knowledge as non-accidentally true belief. When *S* knows that *p*, I argue, she knows it by a reliable method, and her reliability is no accident. There is a connection between her using method *m* and its being reliable. The challenge in ethics is to meet this condition. What does the reliability of our methods have to do with the fact that we use them?

This demand has metaphysical implications. Unless our beliefs are constitutively bound to ethical facts—or aligned with them by an accommodating God—our use of reliable methods in ethics cannot be explained. So, at least, I attempt to show. This conclusion leaves room for constitutive theories of opposing kinds, some ‘constructivist’, explaining the facts in terms of our beliefs, others ‘externalist’, explaining our beliefs partly in terms of a tendency

⁶ Here I am concerned especially with the work of Sharon Street (2006, ms.).

to match the facts. The fundamental question for such accounts is how to connect fact with belief, in either direction, without falling into relativism or predicting an implausible measure of convergence in ethical thought. Here disagreement returns as a constraint on theory. If there is a common subject matter about which individuals and communities disagree, sometimes in radical ways, the facts involved cannot be too closely bound to what any of them believe. At the same time, the connection of fact and attitude must be strong enough to explain the reliability of some.

In chapter four, I try to reconcile these conflicting claims. If ethical facts are bound to our beliefs through the natural history of human life, I argue, it is no accident that we are ethically reliable, when we are. At the same time, there is room for individuals and even whole societies to go astray. If ethical knowledge is possible, despite the limits of convergence, it is explained by human nature. The argument of the book thus leads, in a new way, to a traditional view: that ethical knowledge needs foundations, and that unless they are merely social, they must be found in human nature, in the nature of reason, or in God. My approach is secular, and I argue against foundations in reason. If we hope to save ethical objectivity, human nature is the only place to turn.

The closing sections of the book explore the metaphysics of human nature, along with its empirical study. Does the doctrine of what I call 'natural reliability' conflict with what is empirically known about the diversity of ethical belief? I argue that it does not, though my argument rests in part on the poverty of existing work. Our empirical knowledge of human nature is thin. I also confront a final issue, about knowledge and justification. Unlike the arguments of chapters one and two, the argument from non-accidental truth bears specifically on what we know. Is there room to give up on ethical knowledge while saving justified belief? Not on the assumption, which I defend, that justified belief turns on the capacity to know. If ethical knowledge is impossible, we are not entitled to ethical belief. The task of this book is to expose and

defend the conditions under which, against the threat of epistemic luck, we can lay claim to both.

2. ETHICAL SUPERVENIENCE

The supervenience of the ethical lies in the background of our discussion, especially in chapter one. According to the doctrine of supervenience, ethical differences turn on differences of other kinds; ethical concepts apply in virtue of others. But what exactly does this mean? The idea of ethical supervenience can be developed in several ways, some of which raise serious problems. The purpose of this section is to work these problems out.⁷

It might be thought, to begin with, that the ethical supervenes on the 'natural': no difference in the ethical properties of an act or agent without a corresponding difference in their natural properties. The trouble is to find an account of the natural, a way of classifying properties as natural or not, on which it is safe to assume that ethical properties are not themselves natural. If natural properties are those with causal powers, it is controversial to suppose that ethical properties are not: that they are causally inert. If natural properties are the subject matter of the natural and social sciences, reductive naturalists will hold that they include the ethical.⁸ If being right is maximizing pleasure, the property of being right is ethical and natural at once. The idea that ethical properties are natural properties is not an objection to supervenience. The problem is rather that it makes the supervenience of the ethical on the natural entirely trivial. It fails to capture the

⁷ In writing it, I have been helped by Sturgeon 2009, which provides more detailed references and deals with some distinctions I neglect.

⁸ What do I mean by 'reductive naturalism'? Since I do not use the term in the principal chapters, and since I drop the idea of a natural property as unhelpful, I leave the details for others. A rough conception will suffice.

intended thought: that ethical differences depend on differences of other kinds.

A more promising line is that the ethical supervenes on the non-ethical, but this too will need revision. The difficulty stems, again, from reductive naturalism. If being right is maximizing pleasure, the property of maximizing pleasure should count as an ethical property. It is picked out by an ethical concept, the concept of being right. It follows, if ethical properties supervene on others, that there can be no difference in what maximizes pleasure without a difference of some other kind. There is, however, no reason to assume, as an implication of ethical supervenience and reductive naturalism, that maximizing pleasure supervenes on anything else. If the doctrine of supervenience is to accommodate this, it must be formulated in some other way.⁹

The solution, I think, is to shift from properties to concepts, constituents of propositions.¹⁰ Even if being F is being G, the proposition that *x* is F may be distinct from the proposition that *x* is G, and the corresponding concepts not the same. We begin, then, with ethical concepts: *right, wrong, practical reason, acting well or badly, moral obligation, justice, benevolence, courage*. As before, the safest rule of classification is: if in doubt, it's ethical. A case of particular importance, and a partial exception to this: concepts of non-factive attitudes, like belief, with ethical content, are not ethical. In the relevant sense, that someone believes that slavery is wrong is not an ethical proposition, unlike the fact that he knows this to be true. Ethical beliefs are among the facts on which the ethical supervenes: they are relevant to virtue and vice.

Ethical concepts in hand, we can formulate a principle of possible variation:

⁹ This point is made, at greater length, by Nicholas Sturgeon (2009: 69–73).

¹⁰ Where propositions are Fregean thoughts, individuated by sense or mode of presentation. Russellian theorists must adapt this solution to their view.

There can be no difference in the ethical concepts that apply to acts or agents without a difference in the non-ethical concepts that apply to them.

But there is a risk of triviality, again. If the concept, *thinking it is wrong*, is not an ethical concept, and the concept *truth* is not, what about the concept, *being such that if someone had believed that it was wrong, their belief would have been true*? If this too is non-ethical, we can manufacture non-ethical concepts to ensure the truth of supervenience without respecting its intended force. There are various ways in which one might respond to this, for instance, by insisting that a concept can be ethical even when its component concepts are not. A simpler point is that the idea of ethical supervenience is an idea of dependence, not mere correlation, and the application of the counterfactual concept above is never what makes an action wrong. What is true, without being trivial, is this:

ETHICAL SUPERVENIENCE: If an act or agent falls under ethical concept, E, it does so in virtue of falling under non-ethical concepts, N, such that necessarily, what falls under N falls under E.

This principle does preclude a certain form of naturalism. It is consistent with Ethical Supervenience that the nature of ethical properties can be fully expressed with non-ethical concepts, as in the doctrine that being right is maximizing pleasure. More generally, there is room for:

REDUCTIONISM: For every ethical concept, E, there are non-ethical concepts, N, with which we can say what it is to have the property picked out by E.

What is not consistent with Supervenience, or what would make it problematic, is that the concept, *being right*, just is the concept, *maximizing pleasure*. That would make the concept, *maximizing pleasure*, ethical, and Ethical Supervenience would claim that its application supervenes on something else. That is not required for the supervenience of the ethical. If we want to allow for reductive

naturalism at the level of concepts, not just properties, we must revise our principle again.

My response is to deny the antecedent. We began with a generous list of ethical concepts: *right*, *wrong*, *practical reason*, *acting well or badly*, *moral obligation*, *justice*, *benevolence*, *courage*, and more. That these concepts are identical to ones that fall outside the list is quite implausible. The proposition that *x* is right is not the proposition that it maximizes pleasure: these are different objects of belief. Moore's open question argument, according to which it is intelligible to ask, for any such equation, 'It is N, but is it E?' may not do much to prove this.¹¹ (Is intelligibility a sure test of concept identity? How can we generalize from the intelligibility of the questions we have considered so far to the intelligibility of all?) But the relevant conclusion holds: ethical concepts form a limited class, divined by our inclusive list; and what is not on the list is not an ethical concept. It is to ethical concepts, so conceived, that Supervenience applies.

¹¹ Moore 1903: Ch. 1.