

The Foundations of Practical Reason

D. Phil. Thesis

Submitted for Examination in March 2007

by Mario Brandhorst

Merton College
Faculty of Philosophy

Supervisor: Dr. T. W. Child

Address for correspondence:

Mario Brandhorst
Limastr. 12
14163 Berlin
Germany
mbrandh@gwdg.de



For Kel

The difficulty is to realize the groundlessness of our believing.

WITTGENSTEIN

The Foundations of Practical Reason
D. Phil. Thesis
Faculty of Philosophy
Supervisor: Dr William Child

Mario Brandhorst
Merton College

Hilary Term 2007

Abstract

This thesis is an examination of the foundations of practical reason. Building on the later work of Wittgenstein, I argue for a subjectivist view of moral judgment and of judgments about reasons for action. On this view, moral judgments and judgments about reasons for action can be true or false, but they are not objective.

The argument for this view has the form of an inference to the best explanation. Using a distinction between primary and secondary qualities, I suggest that moral judgments and judgments about reasons for action should not be construed as referring to an ethical or normative reality that exists independently of us. There are ethical facts and facts about our reasons, but these facts arise as the result of our involvement in a linguistic practice. This provides a new way of accounting for these judgments that differs both from moral realism and expressivism.

The view of reasons that emerges is closely related to, but not identical with, reasons internalism as described by Bernard Williams. I reject his argument in favour of internalism and provide a new and independent argument to support this view of our reasons. In the course of spelling out that argument, I show why internalism as described by Williams should be modified, and why this does not commit us to externalism.

In the final chapters, I show that there is an important parallel between our practical predicament and the account of our epistemic condition as portrayed by Wittgenstein. The inference to the best explanation is completed by considering a number of objections to subjectivism that are based on the idea that a subjectivist account of moral judgment and of reasons fails to do justice to the ethical phenomena. I reject these objections, and suggest that a subjectivist can both be reflectively aware of his subjectivism and continue to live well.

Contents

Acknowledgments	iii
Introduction	1
1 Perspectivity and Colour	12
1 A Maze of Little Streets and Squares	14
2 Judgments, Propositions, Attitudes	20
3 Truth and Correspondence	23
4 Reality and Objectivity	26
5 How Do I Know That This Colour is Red?	35
6 Objectivity and the Autonomy of Grammar	47
7 Explanation and the Absolute Conception	58
8 Primary and Secondary Qualities	63
2 The Truth in Moral Realism	67
1 The Application to the Ethical	68
2 Minimalist Maxims	72
3 The Inference to the Best Explanation	76
4 What About Expressivism?	84
5 Wittgenstein Again	90
6 Ethics and Aesthetics	97
3 Are There External Reasons?	109
1 The Structure of Internalism	110
2 A Humean Challenge	118
3 Requirements of Rational Deliberation	129
4 The Unsupported Premise	132

5	Back to Basics	137
6	The Problem of Induction	139
7	A Sceptical Conclusion	141
8	Wittgenstein on Reasons	142
9	The Fundamental Facts of Human Nature	146
10	A Misleading Analogy	149
4	A Fresh Approach	158
1	Physics, Fools and Heretics	159
2	The Expansion of π	165
3	Reasons, Bedrock and Foundations	172
4	A Realistic View of Reasons	182
5	From Subjectivism to Internalism	188
6	Reconsidering Reality	193
7	Is This Still Internalism?	201
8	Transcendental and Empirical Idealism	206
5	Frameworks and Foundations	210
1	An Elusive Piece of Writing	211
2	Knowledge or Unshakeable Conviction?	214
3	A Peculiar Logical Role	218
4	Building a World Picture	222
5	Theory and Practice	231
6	Ethical Thought Without Foundations	235
6	Living With Subjectivism	241
1	Stability and Independence	243
2	Making Sense of Moral Motivation	250
3	The Ring of Gyges	260
4	Deliberation and Experiences of Externality	267
5	What Is a Free Exercise of Reason?	270
6	Convergence and Its Explanation	276
7	The Language of Morality	283
8	The Epistemic Status of Morality	291
9	The Charge of Arbitrariness and Power	295
10	Conclusion	303
	Bibliography	307

Acknowledgments

Many people have contributed in many ways to the production of this thesis.

First, I wish to thank the Scatcherd European Trust and the Royal Institute of Philosophy for their generous financial support, without which this thesis could not have been written. I would also like to thank the *Studienstiftung des deutschen Volkes*, which provided generous support when I first came to Oxford.

Second, I wish to thank those individuals who have discussed ideas and drafts with me, given advice, or provided other forms of feedback. These include Peter Bieri, Simon Blackburn, John Broome, Krister Bykvist, David Charles, Roger Crisp, Philippa Foot, Lizzie Fricker, Peter Hacker, Martha Klein, Thomas Krödel, Kirsten Meyer, Derek Parfit, Julius Schälike, Thomas Schmidt, Holmer Steinfath, Helen Steward, Holm Tetens, Jay Wallace, Ralph Wedgwood and Tim Williamson. I am grateful to them all.

Special thanks are due to Roger Crisp, Lizzie Fricker, Ben Morison, Helen Steward and Ralph Wedgwood for letting me teach some of their students. I found this a most rewarding experience.

My greatest intellectual debt, both for encouragement and philosophical instruction, is to my supervisor William Child. He has shaped my thinking about Wittgenstein and ethics, and my outlook on philosophy, much more than he may realize. I am certain that without his patience, good advice and careful critical attention I would not have finished this degree.

It will be obvious to anyone who reads this thesis that it owes a great debt to the work of Bernard Williams. Sadly, Williams died before I had a chance to get to know him better, but both personally and through his published work, he was a most inspiring figure.

On a personal note, I would like to thank my parents and grandparents for years of support and freedom to pursue my studies. I owe a great debt to my friends at Oxford. They helped me when Oxford life was most enchanting, and still more on rainy days when it was not.

My biggest debt of all is to my wife, to whom this work is dedicated.

Introduction

The topic of this thesis is the question what forms the foundations of an ability that is distinctive of human agency. The ability in question may be called the ability to reason practically, and what it consists in will become clearer in due course. But what are we asking when we ask what forms its foundations?

The question does not call for a historical, or sociological, or psychological answer. It is not, or not in the first instance, an empirical question. It is a conceptual one, given a sufficiently wide understanding of that notion. The question is what we need to presuppose if we want to make sense of this aspect of human life, interpret this area of our activities, and arrive at a reflective understanding of the language that we use when we are engaged in, or describe, practical thinking. This is not to say that explanation, including modes of psychological explanation, have no place in the account. It is to say that explanations can take many different forms, and that scientific explanations, for example, have an interest that may go beyond their most immediate significance for science.

Introduction

The picture of practical reason that emerges as the argument unfolds is that our ability to think and reason practically is neither based on reason, nor on sentiment, but on something more basic than that. It may be called a practice, or a form of life, and it manifests itself in the language that we speak. As we learn to use words such as 'reason', 'good', 'wrong', 'useful' or 'desirable', we grow into a social world that is shaped by the uses of these words, as of many others, and by the judgments people make when they make use of the words in question. At the social level, explanation comes to rest with the social practice of using these words. At the level of the individual, explanation comes to rest with the dispositions that enable us to partake in the practice. Either way, explanation comes to rest with fundamental facts of human nature. The fundamental fact in this account of our ability to reason practically is that we are social creatures who have the ability to speak a language. Our language permeates our lives and has evolved in a natural world that exists independently of us. But our ability to reason practically neither has, nor needs, further foundations in objective fact.

In this way, the account of our ability to reason practically can contribute to a naturalistic understanding of human affairs. Its distinctive mark is the denial that our linguistic practices put us in touch with a normative, evaluative or ethical dimension of reality. Using a Wittgensteinian expression, we can describe it by saying that "the way in which some reality corresponds - or conflicts - with a physical theory has

no counterpart here".¹ In virtue of that denial, the view may justly be called a form of subjectivism. Its hallmark is the claim that reasons, values and moral commitments are subjective, where their subjectivity entails that there are no facts about reasons for action, values, and moral matters independently of our involvement in a linguistic practice. In what follows, I will try to explain and muster support for this view.

Chapter 1 establishes the conceptual framework in which the subjectivist account of reasons, values and ethical judgment is located. Based on a distinction between truth and objectivity, the argument establishes that there is room for an intermediate position that differs both from moral realism and expressivism. Moral realists insist that moral judgments can be both true and objective, while expressivists deny that, at the fundamental level of analysis, moral judgments can be either true or false. I agree that moral judgments can be true or false, but deny that they are objective. The model for this intermediate account is found in a subjectivist account of colour. The analogy establishes that there can be truth without objectivity, and suggests that the relevant distinction between truth and objectivity has to do with different styles of explanation.

Chapter 2 develops this analogy and applies it to the moral case. On the view that I describe, there are moral facts, and there can be moral knowledge, but it is not knowledge of a moral or a normative reality. Given that there is room for such an intermediate view, I then argue that it

¹ Cf. R. Rhees, 'Some Developments in Wittgenstein's View of Ethics', p. 24.

has a good chance of being true, because it relies on resources on which any realistic view of our psychologies has to rely anyway. In this respect, it differs from all forms of moral realism, because moral realists rely, and must rely, on something more. The basic argument in favour of subjectivism is that it provides the better explanation of the ethical phenomena.

Chapter 3 turns to the subject matter of practical reasons. One aim of the chapter is to explain the distinction between internal and external reasons, and to get clear about an influential argument by Bernard Williams to the effect that all reasons for action are internal reasons.² I do not think that the argument is sound or can be made sound without begging the question against the externalist. One important truth contained in it, however, is worth stating clearly, which is that internalists do not rely on dubious assumptions about supposedly different explanatory roles of reason and desire. I argue that the argument of 'Internal and External Reasons' is better understood as the attempt to show that there is a basic incoherence in the thought that there might be external reasons. I doubt that the argument shows this, but insist that for all that, internalism may be true.

In the second half of the chapter, I begin to argue independently for a way of thinking about practical reasons that may be called a version of internalism. I start by interpreting a number of remarks in which the later Wittgenstein responds to traditional philosophical puzzles concerning

² Cf. B. Williams, 'Internal and External Reasons', in: B. Williams, *Moral Luck*, essay 8.

Introduction

induction. What Wittgenstein says, I believe, amounts to a specific view of the metaphysics of our reasons, and this view may justly be called subjectivist. The important fact about our reason claims is, as before, that they can be either true or false, but that they are not objective.

Chapter 4 argues that such a subjectivist view of the metaphysics of reasons strongly supports an internalist view of reasons for action. It emerges that this form of internalism is not quite the same as that Williams originally envisaged. It is, however, unmistakably related to internalism as Williams describes it, and it can accommodate a number of important insights that are often thought to be incompatible with it. If this account is on the right lines, it shows that subjectivists can make sense of the phenomenon of acting for a reason, and in this way, further strengthens the subjectivist account.

In chapter 5, I turn to what may seem to be the major obstacle to subjectivist accounts. It is often said to be unclear how a subjectivist can make sense of our practical commitments, in particular, when these commitments have some moral content. In response to this objection, which I find understandable and most important, I develop an analogy between ideas developed by the later Wittgenstein in a book of notes *On Certainty* and our practical commitments. Wittgenstein argues convincingly that the system of belief that guides us in all thought and action neither has, nor needs, a rational foundation. This, I suggest, is equally true of our practical commitments.

Introduction

Chapter 6 addresses a range of objections to subjectivism on the grounds that it fails to make sense of ethical experience. It is the critical experiment for all subjectivist accounts how well they come to terms with these objections. In my view, subjectivism can make better sense of ethical experience than any rival view. This, however, is already more than needs to be established. Given that subjectivism can make sense of ethical experience, and that it relies on less problematic theoretical assumptions, this completes the inference to the best explanation. At the same time, it shows that subjectivism loses much of its initial threat. We can live with subjectivism, and there is good reason to be hopeful that subjectivists can live reflectively, and continue to live well.

Throughout the thesis I rely on Wittgensteinian materials. These are drawn from a wide range of sources, the most prominent of which are *Philosophical Investigations*, conversations with Rush Rhees about ethics dating from 1942 and 1945, 'Lectures on Aesthetics', *Remarks on the Foundations of Mathematics*, *Remarks on Colour* and *On Certainty*. When quoting from these sources, I generally follow the translations in the editions cited. There are, however, frequent changes to the standard translations. Where I found the translation defective, I supplied my own.

I draw on Wittgensteinian resources in a wide range of different contexts. Wittgenstein features prominently in my discussion of colour, ethical judgment, reasons for action and the framework or scaffolding for practical thought.

Introduction

This might suggest that I impose a picture on the Wittgensteinian ideas that is not there to be found, or that the materials I use have no internal connection. This, however, is not so. Insofar as there is a common theme, it is most likely that wherever we may choose to look, we find unsupported practices at the basis of our outlook: “What has to be accepted, the given – one might say – is *forms of life*”.³ Forms of life are the foundation, and this is so with regard to our beliefs about the world around us, as well as with our beliefs about what would be good, or right, or sensible. It is a thought that shows a whole tradition in the philosophy of mathematics to be radically mistaken, and it applies to reasons, too. In my view, there is every reason to extend the application of this view to ethics, and to our practical predicament more generally.

One thing about which I say very little is to what extent this view amounts to something that deserves the name of a moral theory. I doubt that this is a clear question as it stands, partly because it is unclear what a moral theory is or can be, and partly because the notion of the theoretical is itself unclear. This is related to a question that concerns the connections between the idea of philosophy on the one hand, and the idea of producing theories or explanations on the other. Once again, I think that the question is itself unclear, and that when it is made clear, it soon ceases to be of interest.

³ L. Wittgenstein, *Philosophical Investigations*, pt. II, p. 192. All references to this work are to the 3rd edition, Oxford: Blackwell, 2001.

Introduction

The reason for this is that we readily agree that philosophy does not produce theories or explanations that are structurally similar to those of science, and that where we disagree about the aspirations of philosophy having granted that there is an important difference between doing philosophy and doing science, we become immediately engaged in substantial philosophical disputes. These are indeed interesting disputes, but the interest of the question whether philosophy can produce theories or explanations, and if so, in what sense, is then derived from the wider interest of these questions.

Wittgenstein, as is well known, was adamant that it is no part of philosophy to produce 'theories' or 'explanations'. I think that this goes too far. If Wittgenstein is right, there are certain styles of doing theory, or providing explanations, that are no longer available for the purposes of philosophy. There are philosophies that are, and are shown by Wittgenstein to be, nothing but castles in the air, and here it is appropriate to think that the proper task of philosophy is to destroy them.⁴ But this falls far short of the claim that there can be nothing like theory or explanation in philosophy, or that all we can ever aspire to do is to describe the use of words. Given a sufficiently broad understanding of both 'theory' and 'explanation', I do not see what should prevent philosophy from producing either.

Moreover, I do not see why we should not think of Wittgenstein as having made substantial contributions to philosophy so understood. This

⁴ L. Wittgenstein, *Philosophical Investigations*, § 118.

Introduction

immediately raises the question of what he meant when he said that philosophy produces neither theories nor explanations. I have no answer to this question. I am inclined to agree with Williams that among those who think there is room for ongoing philosophical theory, and that Wittgenstein contributed to it, someone owes Wittgenstein an account of why he ceased to see that this was so. But like Williams, I am also moderately confident that such an account might be given. If so, Williams writes,

we might come to understand that if Wittgenstein could no longer see the edifice of an intellectual subject, his sightlessness was not that of Samson, but rather that of Oedipus at Colonus, whose disappearance left behind healing waters.⁵

I think it is a possibility worth taking seriously that to the extent to which Wittgenstein was blind to the possibilities of philosophical reflection, this was due not to blind rage, but rather to especially demanding views of the tasks and possibilities of philosophical reflection. Given such ambitious views about what philosophy would have to be, it is not surprising to find Wittgenstein conclude that there can be no theory, or explanation, or system in philosophy. Perhaps Wittgenstein never entirely freed himself of these conceptions. They were once his own, as the *Tractatus* shows in written words and silence, and they remained, for him, a personal temptation.

⁵ B. Williams, 'Nietzsche's Minimalist Moral Psychology', in: B. Williams, *Making Sense of Humanity*, essay 6, p. 66.

Introduction

It will be enough for the purposes of this investigation if it can be shown that Wittgensteinian ideas can be put to good use in ethics. In particular, I think that it can be shown that Wittgensteinian ideas not only allow, but rather invite and suggest, a range of applications so as to provide a sensible subjectivist account of ethics.

If this can be shown, it allows us to respond to two different kinds of misunderstanding. The first is that Wittgenstein cannot be called upon to support any ethical theory. If this goes beyond the general and unclear point about philosophy and theory, such an attitude is likely to be due to a vague uneasiness about the later Wittgenstein, ethical theory, or both. This is unsurprising, given that Wittgenstein is often seen as a philosopher who had very little time for ethics, and for a long time thought there could be no such topic.

The fact is that there is every reason to believe that to the extent to which Wittgenstein came to allow that there could be such a thing as philosophy, he allowed, and certainly could allow, that there can be moral philosophy. In fact, Wittgenstein himself has a number of interesting things to say about moral philosophy, and the *a priori* objections to the very possibility of ethics that are powerfully expressed in the *Tractatus* and are present to some extent in the transitional 'Lecture on Ethics', gradually fell away with the picture theory of meaning, which provided their central support. But there is a more systematic reason also why the later Wittgenstein should not be seen as being radically opposed to the idea that there could be a philosophical enquiry into, for example, ethics

and aesthetics. The later view makes room for a description, survey, and grammatical investigation of ethical language just like any other. The fact that Wittgenstein himself did not survey this area of language, or did not survey it in detail, should not discourage us from drawing up our own map.

I said that Wittgensteinian ideas not only allow, but invite and suggest a form of subjectivism, and that this point suggests a response to a second misunderstanding. While the first misunderstanding is expressed in the thought that Wittgenstein cannot be called upon to support any ethical theory, the second is expressed in the thought that he suggests a particular ethical theory, namely a realist one. This, I think, is a mistake, and can only seem plausible as an interpretation of his texts when it is imposed on them from the outside, through the spectacles of realist ethical theory. Needless to say, there are insights in Wittgenstein that realists are well advised to take on board, and many realists have done that.⁶ But this is not enough to show that Wittgenstein endorsed a form of realism either in aesthetics or in ethics or in the philosophy of mathematics. If my reading of his writings on these topics is correct, he was right not to endorse it, and has given conclusive reasons why.

⁶ Cf. e.g. P. Johnston, *Wittgenstein and Moral Philosophy*, S. Lovibond, *Realism and Imagination in Ethics*, and J. McDowell, *Mind, Value, and Reality*.

1

Perspectivity and Colour

This chapter prepares the ground for a subjectivist account of moral judgment. I begin by discussing two traditional approaches to questions of truth and objectivity in ethics.

The first is moral realism, which I define as the view that moral judgments are, in the standard situation, (a) either true or false, and (b) objective. The second is expressivism, which denies both that moral judgments are true or false, and that they are objective. According to expressivists, moral judgments have a primarily practical function and should not be construed as representations at all. Expressivists admit, or rather insist, that moral judgments can still be called true or false and even objective. But they think that moral judgments can be called true or false

and objective only in a derivative sense, namely as expressions of moral attitudes. The moral attitude expressed, however, is construed as being neither true nor false, nor could it be objective.

It is a consequence of this way of construing the debate that truth and objectivity are not the same. This is a welcome consequence. We have good independent reason, I believe, to distinguish between truth and objectivity.

There are many ways of drawing a distinction between truth and objectivity, and there is no way of drawing this distinction that is, in itself, correct. Everything depends on our interests and purposes, and on the phenomena to which the distinction is applied. One way of drawing it that I think is relevant for ethics revolves around the relativity of certain facts or truths to certain standpoints or perspectives. To illustrate the contrast, I will discuss the relativity of secondary qualities to the perspective of perceivers. My discussion will be centred around colour, but it equally applies to other secondary qualities. The conclusion of this chapter is that the perspectivity of colour is based on differences in explanation between primary and secondary qualities.

The upshot is that there can be truth without objectivity, and that this is the situation with respect to colour. Despite their subjectivity, there are facts or truths about the colours of physical objects. Indeed, the colour of an object is something that we can come to know. But such truths or

facts are not objective. They essentially involve a reference to our perceptual perspective.

This suggests that in ethics, there is room for a middle position, which is the one that I will explore. On this view, it is true that there are moral truths or facts. But it is not true that they are objective. Thus, the account will share features of both moral realism and expressivism. Given its denial of objectivity, the result will be a form of ethical subjectivism.

What the shape of such a subjectivist account might be and what reasons there are to believe it will, I hope, become clear in due course. But a few preliminary remarks are in order. As we shall see, they already take us a good way towards the main questions with which we will be concerned.

1 A Maze of Little Streets and Squares

First, there is the question of the subject matter that is circumscribed by terms like 'ethical' and 'moral'. Some writers, notably Bernard Williams, distinguish between the subject matter of the ethical and the subject matter of the moral. As they use terms, the 'moral' is a subclass of the more inclusive 'ethical'. Taken in this way, the ethical embraces everything that bears on the question how to live, while the moral is

construed as being centred around a special and demanding notion of moral obligation and characteristic responses of blame and guilt.¹

I will use both terms interchangeably. As Williams himself emphasizes, there is a useful notion of obligation that survives the collapse of the 'morality system', and there is no reason not to call it 'moral obligation'. Indeed, before he introduces his distinction between 'ethical' and 'moral', Williams speaks quite freely of 'moral philosophy', 'moral claims' and 'moral considerations', "using that word in a general way, which corresponds to what is, irremovably, one name for the subject: moral philosophy".² It is this more inclusive and much less demanding sense that I have in mind when I discuss the subject matter of the moral.

With regard to this sense of the word, I doubt that there is more, or much more, of substance to be said about the class of considerations that we are used to calling 'ethical' or 'moral'. This is so at least to the extent to which we are looking for, or are implicitly guided by the idea of, an essence of the moral, something like a set of common features that holds moral phenomena together as a class. There is no reason to believe that there is any such thing, but there is also no reason to be troubled by this. It does no harm that the notion of the ethical is vague, and the same goes for

¹ Cf. B. Williams, *Ethics and the Limits of Philosophy*, esp. chs. 1 and 10. Williams elaborates the contrast further in his Sather Lectures of 1989; cf. B. Williams, *Shame and Necessity*.

² B. Williams, *Ethics and the Limits of Philosophy*, p. 6.

the notion of the moral when we interpret it as I am suggesting.³ In § 71 of

Philosophical Investigations, Wittgenstein writes:

One might say that the concept 'game' is a concept with blurred edges. –
"But is a blurred concept even a *concept*?" – Is a blurred photograph even
a picture of a person? Is it even always an advantage to replace a blurred
picture with a sharp one? Isn't the blurred one often exactly the one we
need?⁴

The question is rhetorical. The answer that Wittgenstein intends, and the
answer that is surely correct, is that very often it is.

If so, we can go one step further. A conclusion that suggests itself at
this point is that while it is not impossible that we discover essences of
basic and central concepts, the odds are very much against it. The reason
is that the meaning of a word depends on the use we make of it, and the
patterns of that use need be neither articulated nor defined. There is no
reason to think that a naturally evolving use will always lead to clear,
orderly patterns.

This is brought out in a beautiful metaphor early in the *Philosophical
Investigations*:

Our language can be seen as an ancient city: a maze of little streets and
squares, of old and new houses, and of houses with additions from
various periods; and this surrounded by a multitude of new boroughs
with straight and regular streets and uniform houses.⁵

³ Cf. B. Williams, *Ethics and the Limits of Philosophy*, p. 7. It is a further question whether in
its more technical sense, the term has, or comes closer to having, an essence. If there is a
special notion of moral obligation employed in the morality system, this may well be its
essential element.

⁴ L. Wittgenstein, *Philosophical Investigations*, § 71.

⁵ L. Wittgenstein, *Philosophical Investigations*, § 18.

This kind of idea is also found in Nietzsche, who famously said that it is impossible to give a precise account of the institution of punishment – for “definable is only that which has no history”.⁶

We can generalize the lesson. Our practices of using words are shaped by our interests and purposes, which is not to say that they were consciously designed or planned. Patterns of use tend to evolve without an overarching reason, and they have evolved in this way over long periods of time. Like ancient cities, the outcome is a complicated structure that defeats any attempt at simple systematization.

If we take this picture seriously, it casts doubt on the idea that there is an essence of the ethical, just as it casts doubt on the idea that ethical concepts themselves can be sharply delineated and defined. Ethical concepts do not lend themselves to systematization, and this has important consequences for the hopes and purposes of moral theory. In particular, the metaphor casts doubt on the idea that ethical concepts could be linked or defined with the kind of precision we find in logic or science. Scientific concepts are at the outskirts of the city, making up much of what Wittgenstein calls the ‘new boroughs with straight regular streets and uniform houses’.

Ethical concepts are at the heart of the ancient city. They are in the very midst of the maze of little streets and squares, of old and new houses,

⁶ F. Nietzsche, ‘On the Genealogy of Morality’, in: ‘*On the Genealogy of Morality and Other Writings*, II, sec. 13.

and of houses with additions from various periods which make up our language. If this is the situation, we should be suspicious of all moral theories that purport to show the ethical to be reducible to one set of ethical considerations, or to something that takes us outside of the ethical domain. There is no reason to expect that 'right' can be defined in terms of 'good' or 'good' in terms of 'right', or to imagine that a single formula could cover every aspect of ethical life or provide the basis for all other moral concepts.⁷ This is equally true of simplistic appeals to 'interest' or 'desire'.

The project of reduction might have a deeper motivation than mere systematization. It may be thought that it reveals the truth about the ethical, insofar as it gives us an account of right and wrong, good and bad, and other aspects of a complex ethical reality. But even if we grant, if only for the sake of argument, that there is a truth about the ethical that can be discovered, it is not clear why that aim should encourage theory or reduction. If there is such a thing as the truth about the ethical, why is there any reason to expect it to be simple?⁸

It is no good to reply that science has established that the simple explanation is most likely to be correct. It is doubtful that this is true of science, but even if it were true, we would only have established the result for science. That nature is uniform and simple, where it is uniform and

⁷ Cf. B. Williams, *Ethics and the Limits of Philosophy*, p. 16: "We use a variety of different ethical considerations, which are genuinely different from one another, and this is what one would expect to find, if only because we are heirs to a long and complex ethical tradition, with many different religious and other social strands".

⁸ Cf. B. Williams, *Ethics and the Limits of Philosophy*, p. 17.

simple, must be confirmed by experience. There is no reason to think that ethical experience is uniform and simple in that way. If anything, experience strongly suggests that it is not.

No matter how we look at it, the idea that ethical concepts can be reduced or shown to have an underlying structure looks deeply implausible. Moreover, conceptual connections and explanatory relations between our ethical concepts that we might find always remain conceptual connections and explanatory relations between *our ethical concepts*. They would decidedly not answer to, or incorporate, some hidden truth about the ethical. The explanation of those conceptual connections and explanatory relations that we might find will run through our linguistic practices, our distinctive form of life and our contingent history.⁹

It should be clear that neither the inherent vagueness nor the radical contingency of the ethical domain constitutes a reason to deny its conceptual autonomy. The notion of the moral is not identical with either the notion of the evaluative, or the notion of the normative, or the notion of the practical, though it overlaps with all of them.¹⁰ Nor is it opposed, in any simple way, to something we might call 'the factual' – or so I shall argue.

⁹ An illuminating study of the ways in which such complex forces interact is B. Williams, *Truth and Truthfulness*.

¹⁰ See B. Williams, *Ethics and the Limits of Philosophy*, pp. 135 f.; for his own, deliberately vague characterization of the ethical, see chs. 1 and 10.

2 Judgments, Propositions, Attitudes

In this chapter, I am concerned to clarify and to defend the claim that moral judgments are subjective. I have pursued a number of questions pertaining to the first feature of the formula, the subject matter of the ethical or moral. Next, something should be said about the use of 'moral judgment', which raises a number of familiar questions.

The expression 'moral judgment' seems to associate the conceptual with the psychological in unhelpful ways. It is natural to think that we should not, or should at any rate not primarily, be interested in moral judgment, but focus on the moral proposition instead.¹¹ The problem is that this requires us to be able to say, in advance and across the board, that what we are dealing with when we discuss moral attitudes are attitudes that have propositional content. This, however, is precisely what we cannot do, because it is a contentious question whether moral attitudes are propositional attitudes. Moreover, even if we agree that they are propositional attitudes, we can still disagree about the correct interpretation of the term 'propositional attitude'. In order not to prejudge any questions concerning propositional content, I will continue to speak of moral judgments, leaving it deliberately open whether moral attitudes are adequately understood as having propositional content, and if they are, what that might mean.

¹¹ Historically, the use of 'proposition' is bound up with the demise of 'judgment' as the result of the logical critique of psychologism; cf. G. Frege, 'On Sense and Meaning' and 'Thoughts', in: G. Frege, *Collected Papers on Mathematics, Logic, and Philosophy*, pp. 157-77, 351-72.

The idea that moral judgments have propositional content has been explicitly questioned by writers in the expressivist tradition, who frequently deny that we should think of the content of a moral judgment on the model of the proposition.¹² Expressivists emphasize the practical dimension of moral thought and language. While the moral realist believes that the attitudes expressed in moral judgment can be either true or false, the expressivist believes that they are neither.

But this claim immediately calls for a further qualification. Expressivists like Blackburn and Gibbard say that moral judgments *can* be true or false, even if they express moral attitudes that can be neither. The idea is that on the expressivist account, predicates like ‘...is true’ and ‘...is false’ amount to nothing more than rhetorical devices for affirming or rejecting certain attitudes, and thus add nothing of substance to the attitude originally held or expressed.

One question that arises immediately about this construction is how it is possible that a moral judgment *can be* true or false if it has no propositional content. The expressivist seems to be saying (1) that moral judgments have no propositional content, (2) that moral judgments can be true or false, from which it follows (3) that they have content that can be either true or false. But (4) content that is either true or false *is* propositional content, so the expressivist is now committed to holding

¹² ‘Expressivism’ is associated in particular with the work of A. Gibbard; *cf.* A. Gibbard, *Wise Choices, Apt Feelings*. The label has been adopted by S. Blackburn; *cf.* *Ruling Passions*, p. 77. Two classic statements of a similar position are A. J. Ayer, *Language, Truth, and Logic*, ch. 6 and C. L. Stevenson, *Ethics and Language*.

that moral judgments have propositional content, which flatly contradicts (1).

The obvious response to this problem is to extend the expressivist analysis to include not only talk of truth and falsity, but to include talk of propositional content as well. Blackburn, for one, seems to have no qualms about talk of the ethical 'proposition'.¹³ But now we are faced with a new kind of problem. If the expressivist says that moral judgments can be true or false, that they have propositional content and that they can be the object of a belief, what exactly is it that, on the expressivist analysis, moral judgments cannot be? Representations? But they are true or false. And if moral judgments can be true or false and have propositional content, how can they fail to be representations? Perhaps they are not representations in the sense that they do not represent the moral facts? But if they are representations, as we are assuming, and if their subject matter is the moral, as we have been assuming all along, how can they fail to represent the moral facts? These claims are all part of the same package.

The suspicion that arises, therefore, is that the claims of expressivism are vacuous. To deflect the charge, however, is the task of the expressivist, and I will not try to decide the issue here. What is clear, I think, is that the disagreement between the expressivist and his realist opponent is not one about the question whether there are moral truths, moral facts, moral beliefs and moral propositions. Everything depends on

¹³ Cf. S. Blackburn, *Ruling Passions*, ch. 4.

what we say *next*. The real question is what kind of theoretical commitment we incur when we speak of moral propositions, moral truths and moral representations. The disagreement will presumably be located here, since for the moral realist, belief is what one might call a basic moral attitude. For the expressivist, by contrast, there will always have to be some deeper grammatical analysis that shows ethical beliefs to be something other than they seem to be. But even if we agree, if only for the sake of argument, that moral judgments can be true or false, and express propositions in the way realists say they do, new questions arise.

3 Truth and Correspondence

There is a need for an expression like 'proposition'. When I say to you 'Tom looks tired', you say to Tom 'You look tired', and Tom says to us 'I look tired', we want to be able to say that these are so many different ways of saying *the same thing*.¹⁴ The point of the expression 'proposition' is that it will let us do that, and when we say that different sentences or statements all say the same thing, that is clearly not to say that there is any one 'thing', abstract or otherwise, that says it. As Wittgenstein points out, we can learn the concept 'proposition' without even mentioning the idea of abstract entities. To say that a proposition is whatever can be true or

¹⁴ Cf. P. F. Strawson, 'Truth', in: P. F. Strawson, *Logico-Linguistic Papers*, p. 192.

false essentially amounts to saying that we call something a proposition when in our language we apply the calculus of truth functions to it.¹⁵

As a working definition of the concept 'proposition' this is perfectly innocuous, as long as it is remembered that there is a close connection between the concept 'proposition' and the concepts 'true' and 'false'. As Wittgenstein points out, there is this basic connection, but for that very reason, giving a definition of 'proposition' in terms of 'that which can be true or false' is not very informative. It merely draws out a conceptual connection that was already in place. We *have* the concepts 'true' and 'false' only if we already have the concept of the proposition, and there is no way of driving a wedge between them.¹⁶

Finally, there is a close connection between the concept of a proposition and the concept of a fact, which is what 'makes the proposition true' if it is true. This connection must be stated with care, because it invites a well-known misunderstanding. Consider Moore, who thinks that truth, in the context of a belief, consists in a relation between the belief that *p* and the fact that *p*, and claims that this relation can be defined as a relation of correspondence. If so,

we can at once assert that 'To say that this belief is true is to say that there is in the Universe a fact to which it corresponds; and to say that that it is false is to say that there is *not* in the Universe any fact to which it

¹⁵ Cf. L. Wittgenstein, *Philosophical Investigations*, §§ 134-6.

¹⁶ The implicit target here is what may be called a metaphysical conception of the nature of the proposition; cf. L. Wittgenstein, *Notebooks 1914-1916*, p. 39: "My *whole* task consists in explaining the nature of the proposition. That is to say, in giving the nature of all facts, whose picture the proposition *is*. In giving the nature of all being" (22. 1. 1915).

corresponds.' And this statement, I think, fulfils all the requirements of what we actually mean by saying that the belief is true or false.¹⁷

But is quite unclear what this could mean. On the one hand, facts are said to be in the world, or as Moore says, 'in the Universe'. But facts are abstract, linguistic entities. They are not 'in the world' in the way in which tables and chairs are. Facts have no history. They leave no trace in space and time. Facts can be remembered, grasped or seen. They cannot be located, or touched or painted or moved about.

The difficulty is systematic. It appears that if the requirements of the correspondence theory are to be met, facts must be in the world if correspondence *is a relation*. But 'in the world' is exactly where they cannot be if we think of them as that which makes a proposition true or false. For what could make 'Brutus stabbed Caesar' true, if not the fact that Brutus stabbed Caesar? If the relation between that fact and that proposition really is a relation of *correspondence*, facts themselves must have a linguistic structure. But what could it mean to say that a fact is 'in the world' or 'in the Universe', and has a linguistic structure?

The best way out of this quandary is to give up the theory of truth as correspondence. This does emphatically *not* mean that we should deny that what is true corresponds to the facts, or that a proposition that is true represents the world correctly. The point is that there is a conceptual connection between 'true' and 'fact', because we can generally replace

¹⁷ G. E. Moore, *Some Main Problems of Philosophy*, pp. 276 f.

'That's true' with 'That's a fact' or 'That is really how things are'. Of course a proposition that is true can be said to correspond to the facts. In fact, it could not be otherwise, precisely because the relation between 'true' and 'fact' is a conceptual connection. Whenever there is a truth, there is a fact that makes it true. If there were no fact, there would be no truth either. As we shall see, this has important implications for understanding both moral realism and expressivism, but it will also prove to be important when we explore the claim that there may well be moral facts, but that such facts are not objective.

The important point is that to say of a proposition that it corresponds to the facts is merely a *variant* on saying that it is true. It does not, and does not need to, provide a substantial explanation of what it is for it to be true. In particular, it does not presuppose any such explanation in terms of a theory of correspondence. There is, I think, a sense in which we want to say, with Wittgenstein: "The world is the totality of facts, not of things".¹⁸ But we will find it less misleading to say, with Strawson: "The world is the totality of things, not of facts", and add that a full description of the world cannot merely list the things, but has to state the facts.¹⁹

4 Reality and Objectivity

How are 'true' and 'false' related to 'objective'? I said that in my view, moral judgments can be true or false, but they are not objective. It is a

¹⁸ L. Wittgenstein, *Tractatus Logico-Philosophicus*, 1.1.

¹⁹ P. F. Strawson, 'Truth', p. 198 n.

consequence of this way of construing the debate that truth and objectivity are not the same.

This can seem surprising. After all, it is natural to think that 'objective' belongs with 'true' and 'false', so that whatever is true is true objectively. The test for truth and falsity would therefore simply *be* the test for objectivity, and moral realism would be the view that moral judgments have propositional content, that is to say, that moral judgments can be either true or false.

This is a common use of 'objective', and I do not think there is anything wrong with this use. The problem is that this way of construing the debate conceals important differences in ethical theory. Many seem to think that the debate about the view that I called moral realism is at bottom a debate about the question whether moral judgments are ever true. Mackie, for example, argues that no moral values are objective, and seems to equate that claim with the view that there are no moral facts. Value, he argues, is not "part of the fabric of the universe". Given his realist interpretation of moral language, the upshot is an error theory. The theory says that while moral judgments are capable of being true, they are in fact all false, because there are no moral facts that make them true.²⁰

I believe this is too simple. It is too simple for the obvious reason that expressivists do not deny that moral judgments can be either true or false, so that we would have to say something more about moral realism.

²⁰ J. L. Mackie, *Ethics*, ch. 1.

The more interesting reason, however, is that this way of construing the debate passes over an important distinction. In this way, it tends to conceal an interesting option in ethical theory from our view.

What is the relevant distinction? As we saw, truth and falsity are conceptually connected with the notion of the moral proposition, which is conceptually connected with the notion of a moral fact. There is a sense in which this, by itself, is more than enough for objectivity. Objectivity entails some kind of independence, in particular, some kind of independence of our thoughts about, and attitudes towards, the facts. Where the subject matter is the moral proposition, truth and falsity themselves require such an independence. After all, on any plausible account of truth and falsity, it must normally be possible for my beliefs about the matter and the facts to come apart.

But this is not the only possible interpretation of 'objective'. In my view, moral objectivity as the realist construes it requires something more. Truth and falsity, when used in moral contexts, entail that there are moral facts. Objectivity entails that the moral facts are as they are independently of us. If moral realism is correct, moral facts are something that we find, discover, or acknowledge. They are not something that we invent, or construct, or create. But so far, these are merely metaphors. What could the relevant difference consist in?

So far, I have discussed the possibility that there may be moral truths, but that those truths are not objective. This approach is useful, but

only up to a point. It is useful insofar as it answers to a basic intuition in this area, namely that moral judgments seem to be capable of being true and false, and that they present themselves as objective. The approach is useful only up to a point because there are ways of construing moral objectivity according to which it is based, not straightforwardly on the objective truth of moral judgments, but on the objective validity of moral judgments. Alternatively, it can be based on the idea that there are certain kinds of feelings and dispositions one must have if one is to be in good shape as a human being, where what it is to be in good shape as a human being can be established as a matter of objective fact.²¹

These strategies are at the heart of Kantian and the Aristotelian approaches to moral objectivity. Their implications for the question whether moral judgments can be said to be true are complex and interesting, but I will not pursue the matter here.

The Kantian programme is essentially the attempt to provide an objective grounding for morality that is not realist. According to Kant, moral claims are based on practical reason itself, rather than being based on features that, for Kant, belong merely to the world of experience. The Aristotelian approach, though very different, similarly attempts to establish objectivity in ethics by appealing to practical reason rather than to 'plain' objective moral truth and our ability to come to know it.

²¹ For this way of drawing the distinction, see B. Williams, 'Ethics and the Fabric of the World', in: B. Williams, *Making Sense of Humanity*, essay 14, esp. pp. 174-6. See also C. Korsgaard, 'Aristotle and Kant on the Source of Value', in: C. Korsgaard, *Creating the Kingdom of Ends*, ch. 8 and B. Williams, *Ethics and the Limits of Philosophy*, chs. 2-4.

The fact that truth and objectivity can come apart and that this independence runs in both directions, once more raises the question what the difference between truth and objectivity could consist in. But it also suggests a tentative answer. Consider the Aristotelian appeal to the notion of practical wisdom. Truth and objectivity remain a central feature in the Aristotelian account of practical reason.²² The basic move is to identify the virtuous action as ‘what the virtuous person would do in the circumstances’. But this will not yield ethical objectivity unless we can stabilize the notion of the virtuous person. After all, there could be various and possibly incompatible conceptions of what it is to be a virtuous person, and as a result, there could be various and possibly incompatible conceptions of virtuous action.

So something more needs to be said about the virtuous person, and on the Aristotelian conception, something more can be said, because the virtuous person is the one who *gets things right* in this domain of thought and practice. This idea, in turn, must no longer be relative to some conception of what is right in this domain of thought and practice or of the virtuous person that is merely subjective. Rather, what the virtuous person says is the virtuous action must be correct absolutely, that is, correct in a sense that does not reflect, but rather provides a criterion of correctness for a particular conception of virtue.

²² The same is true of many who write in the Aristotelian tradition today; cf. J. McDowell, *Mind, Value, and Reality*, esp. ‘Virtue and Reason’ and ‘Projection and Truth in Ethics’, essays 3 and 8.

In this way, the contrast between 'subjective' and 'objective' will concern the ways in which a moral judgment can be said to make a claim on us, in particular, whether it can be said to make a claim on us independently of our believing it, or our accepting it, or our being motivated to comply with it. Given this account of the distinction, the question arises on which side of the divide we should locate the expressivist. The answer seems obvious enough if, as we are assuming, the expressivist claims that at some deeper level of analysis, the attitudes expressed in moral judgments are not pieces of knowledge or belief, but rather preferences or feelings, or some other kind of attitude that is neither true nor false. If there is a fundamental level of analysis at which attitudes that are expressed in moral judgment can be neither true nor false, it is natural to think that the question whether they are true or false objectively does not even arise.

But this is not quite correct. Blackburn, for example, thinks that we can say about many moral facts what we are prepared to say about the facts of nature: "We discover such facts, we do not invent them".²³ It is certainly not merely because of the way we form moral attitudes and beliefs, we want to say, that kicking dogs is wrong. Surely it is wrong to kick dogs whatever we thought or felt about the matter. "Fluctuations in our sentiments only make us better or worse able to appreciate how

²³ S. Blackburn, *Spreading the Word*, p. 217.

wrong it is", as Blackburn puts it.²⁴ This seems to be anything but an overblown requirement of moral realism. It is part of our ordinary way of looking at the world.

Is the expressivist forced to deny it? According to expressivists, 'objective' specifies the *attitude* that is being expressed in moral judgment, and should not be construed as an expression that makes a further claim *about* them. There is clearly a truth in this. What makes it wrong to kick dogs is that kicking dogs is cruel and painful, not that we have certain sentiments about it. *That* input should yield disapproval and indignation as the output.²⁵

Expressivists construe the contrast between 'subjective' and 'objective' so that it runs parallel to the expressivist interpretation of our use of 'true' and 'false'. There is, however, at least one important difference. While 'true' and 'false' add nothing substantial to the content of a moral judgment, 'subjective' and 'objective' at least serve to mark a genuine difference *in attitude*, since not every attitude we might wish to express is objective in this way. What is rude or polite, for example, normally does depend on what we feel about the matter. So even on expressivist accounts, there is a substantial contrast between 'true' and 'objective'.

Truth and objectivity can come apart, and as we have seen, the independence runs in both directions. As the Kantian model of a moral

²⁴ *Ibid.*

²⁵ *Cf. ibid.*, p. 218.

theory makes clear, there is a way of arguing for substantial moral objectivity in the absence of objective moral facts as the realist construes them. Kant himself, however, needs to be interpreted with care. He makes it quite clear that in his view, the foundations of morality are ultimately bound up with the claim that every rational being, insofar as it is capable of moral thought, has 'absolute value'.²⁶ The Aristotelian model is even more closely related to the realist approach, of which it is, in effect, a special version.

More importantly, we have now seen that there is a way of holding on to moral truth, while rejecting claims to moral objectivity. Consider what may be called the subjectivist version of the Aristotelian conception of practical wisdom. There is room for a view that accepts a notion of the virtuous person, but rejects the idea that the virtuous person 'gets things right' in this domain of thought and practice. The thought would be that moral facts are always relative to moral outlooks or to some conception of the virtuous person that not everyone must share. There would thus be room for facts and truths and propositions, but there would be no sense in which what a virtuous person says is correct absolutely, that is, correct in a sense that does not reflect, but rather provides a criterion of correctness for her subjective conception of virtue.

²⁶ Cf. I. Kant, *Groundwork for the Metaphysics of Morals*, pt. II, e.g. AA 4:429: "Rational nature exists as an end in itself" and 4:435, where Kant claims that "morality, and humanity so far as it is capable of morality, is the only thing that has dignity", where dignity is thought of as "inner worth". Having inner value, in turn, is opposed to having a "price" or being of "relative worth".

So it appears that there is room for a view that is located between moral realism nor expressivism, but identical with neither. On the one hand, we have moral realism, which claims that moral judgments can be both true and objective. If moral realism is correct, there are objective moral facts, to which the proper attitude to take would be belief and, given suitable conditions, moral knowledge. On the other hand, we have expressivism with its categorical denial of both. Of course expressivists agree that moral judgment can be both true and objective. But they admit this only if there is a further theoretical analysis available to back up these realistic ways of speaking. This analysis reveals both moral truth and objectivity as something other than they seem to be. Expressivists do not think of belief as a basic moral attitude.

It may seem that the theoretical alternatives are exhausted by these options. Either we interpret truth and objectivity as the realist does, who construes moral truths or facts as features of a basic normative dimension of reality, or we interpret them as the expressivist does, who thinks of these words as tools to specify the contents of our attitudes. But there is room for a view that combines elements of either theory. On this view, it could be true that there are moral facts, to which the proper attitude to take would be belief. Consequently, there would be no need to assimilate beliefs, at any level of analysis, to attitudes that are neither true nor false, and it might well be possible to accommodate even a claim to moral knowledge. But moral facts do not need to be thought of as objective.

Expressivists may well be right that there is no such thing as an objective, normative dimension of reality.

In this way, the subjectivity involved in ethics could be usefully compared to the relativity of secondary qualities to different perspectives and perceivers. Perhaps there is a sense in which Hume was right when he suggested that vice and virtue

may be compar'd to sounds, colours, heat and cold, which, according to modern philosophy, are not qualities in objects, but perceptions in the mind: And this discovery in morals, like that other in physics, is to be regarded as a considerable advancement of the speculative sciences; tho', like that too, it has little or no influence on practice.²⁷

The question is what the analogy between ethical thought and secondary qualities like colour could consist in.

5 How Do I Know That This Colour Is Red?

In § 381 of the *Philosophical Investigations*, Wittgenstein writes:

How do I know that this colour is red? – One answer would be: 'I have learnt English'.²⁸

This remark is puzzling. In its brevity, it borders on the obscure. If this is one answer, are there others? If there are, is this the one that is correct?

²⁷ D. Hume, *A Treatise of Human Nature*, bk. III, pt. I, sec. I, p. 469.

²⁸ Cf. L. Wittgenstein, *On Certainty*, §§ 527-8.

What does the answer tell us about colour, and our knowledge of a colour, anyway?

Perhaps we should not assume that Wittgenstein is making a specific point about knowledge of *colour*.²⁹ What he has in mind may be a kind of knowledge that is basic, so that it makes sense to think of it as being due to basic linguistic training.

Perhaps a better way of putting the idea would be to say that if someone were to make a mistake about something so very basic, then we would not react as we do when people make mistakes. If you think that Toronto is the capital city of Canada, then those who know better might say that you have made a mistake. Ottawa is the capital city, and this is simply a piece of information that you lack. Suppose now that you think the red flower here in front of you is black. We would no longer think that this is merely another piece of information that you lack. We would expect something rather unusual. You may be blind; you may be joking; you may be expressing yourself metaphorically; there may have been a slip of the tongue – but nothing more. If everything is normal, and you insist that the rose in front of you is black while looking at it in bright daylight, we would begin to wonder if you had grasped the meanings of the words ‘red’ and ‘black’ correctly.

²⁹ Cf. L. Wittgenstein, *Remarks on the Foundations of Mathematics*, 1-3, where Wittgenstein considers the question ‘How do I know that in working out the series + 2 I must write ‘20004, 20006’, and not ‘20004, 20008’?’, and adds: “The question ‘How do I know that this colour is ‘red’ is similar”.

This indicates that there is a close connection between knowledge of colour and linguistic competence. Moreover, it indicates that there is a close *conceptual* connection, and not merely a *psychological* or *empirical* one. One thing that we can safely assume is that Wittgenstein is not interested in doing science or psychology. What Wittgenstein says may well contribute to what he calls the 'natural history of man', but if it does, it will be from a point of view and with an interest that is more than merely psychological.

Still, there seems to be something special about our knowledge of colour, and that it has something more to do with linguistic training. Developing the Wittgensteinian theme more freely, I will now ask what that special role of linguistic training may consist in.

Evidently, the mere fact that we conceptualize experiences in certain ways, for example, as experiences of red, is not characteristic of our knowledge of colour. Knowledge *that* is always structured by concepts, so the fact that linguistic training is involved in our knowledge of colours *in that way* shows nothing that specific about *colour*. The question is how, when we try to understand knowledge of colour, linguistic training has a special role.

To make some progress, we need to be clear about the kind of linguistic training that is implicitly assumed in the case under discussion. After all, there are very many different ways of learning words for colours. Some of them, but not all of them, are available to those who

cannot distinguish certain colours, and the same is true of those who see nothing at all. They cannot distinguish colours just by looking, so they lack certain techniques of using colour words. But they certainly have colour concepts, and these colour concepts are at least close relatives of ours.³⁰

What Wittgenstein says applies to languages only if they have already opened up the conceptual space of colour. The question is how, once that space is opened, linguistic training has a special role in structuring it and making it accessible. Here it is important to remember that the language game under consideration is the normal one of seeing colours and responding to them by using our colour words. Colour words are used against a background of responses that we can describe as seeing colours normally, as opposed to the responses of those who are, for example, blind or colour-blind.

Those who see colours normally learn most colour words by demonstration. The language game that Wittgenstein examines is the basic one of looking at an object, like a ruby or a rose. Given good light, normal vision and our linguistic competence, we know immediately and just by looking what the colour of the object is. Of course there are many colours that we cannot name, and there may not always be an answer to the question what colour an object is. But this is not normally what we say about red, which is among our elementary colour concepts. Seeing red is

³⁰ Cf. L. Wittgenstein, *Remarks on Colour*, I-75, 77, II-31-2, III-114-24, 278-94.

fundamental for our form of life and for our language. Before questions concerning the borderlines between red and orange, red and purple, or red and pink can even arise, our linguistic competence with 'red' must already be well established.

What, then, is so special about our knowledge that this colour is red, and similarly, our knowledge of other colours? I think the special role of our linguistic training is that colour concepts apply merely on the basis of perceptual responses. The important feature is that colour concepts apply just in virtue of the way an object *looks*, and that is the important difference that sets all knowledge of colours apart from other areas of knowledge.

This idea is liable to being misunderstood. Colour concepts apply primarily *to* physical objects. The property of redness is instantiated by the rose, and not, for example, by my visual impression of it. This is one respect in which the contrast between primary and secondary qualities as Hume draws it in the quoted passage is misleading. Hume suggests that colours are 'perceptions in the mind', and there is a clear sense in which that is precisely what they are not. So as it stands, his claim is liable to being misunderstood. But there is an important truth in what Hume says. It is better expressed by saying that coloured things are objects, but that what it is for an object to be coloured involves an essential reference to perceivers. Colour concepts apply to objects, but they apply to objects just in virtue of the fact that those objects look a certain way to us.

This has an important consequence. Colour concepts apply in virtue of the fact that coloured objects present a characteristic visual appearance. The visual appearance of an object is the necessary and sufficient basis for the application of a colour concept. Consequently, colour words can be applied even if we know nothing whatsoever about a coloured object. Whether a red thing is a rose or a ruby depends in complex ways on its origin, its structure and its chemical and physical characteristics. Whether a red thing is red does not *conceptually* depend on any of these things. Of course the colour of an object will be physically *explained* in terms of surface properties and wavelengths. In that sense, it may not be *physically* possible for a red object to have a radically different physical structure. All the same, it is red if it looks red in certain circumstances, and to certain perceivers, whatever its origin, its structure, and its chemical and physical characteristics may turn out to be. In this way, our knowledge of colours does not, so to speak, reach out into the world. Colour words are applied in virtue of characteristic visual appearances. In this respect, they differ fundamentally from the words we use in chemistry and physics.

There is a sense, then, in which the *essence* of colour lies in *being perceived*. Alternatively, one might say that the *essence* of colour secures its *existence*, so that someone who knows what colour *is* can no longer be in any doubt *as to whether there are colours*. Wittgenstein writes:

God's essence is said to guarantee his existence - what this really means is that here what is at issue is not the existence of something.

For could one not equally say that the essence of colour guarantees its existence? As opposed, say, to the white elephant. For it really only means: I cannot explain what 'colour' is, what the word "colour" means, without the help of a colour sample. So in this case there is no such thing as explaining 'what it *would be like* if colours *were* to exist'.³¹

If such a view of colour is correct, there is a sense in which physical and chemical characteristics drop out as irrelevant:

We are not concerned with the facts of physics here except insofar as they determine the laws governing how things appear.³²

I may know nothing else about the thing in front of me, but there will be no question that if I look at it in good light and my sight is unimpaired, then I have every reason to believe that it not only *appears* to me to be red, but that it *is* red. In such favourable circumstances, I *know* it to be red. I know it to be red, moreover, even if I pretend that the object has a different physical or chemical structure. Indeed, I can imagine that chemical structure constantly changing. As long as the object looks red to me, I see it clearly, and certain other conditions are satisfied, I know that the object is red. If I can be said to know it to be red, it must of course be true that it is red. A systematic error in my knowledge that this is red is not even possible.

³¹ L. Wittgenstein, *Culture and Value*, p. 94.

³² L. Wittgenstein, *Remarks on Colour*, III-180.

This is one facet of the fact that 'red' and other words for colours like it are basic to the human form of life, and why it is not easy to imagine human life without them. Red things stare us in the face. They always have done and they always will, so it is only natural that we should have introduced a word to pick out such things in the way we do. The concept 'red' is central to the human form of life in ways in which no chemical concept ever was or is or can be. Colour words will still primarily apply to objects, not to impressions of those objects. But they characterize these objects in virtue of, and with respect to, the visual impressions they produce in us in certain circumstances. That is what colour words are *for*.³³

Some may be inclined to think that where we draw the boundaries of 'red' has to reflect some deeper, underlying similarity between red things that we may have yet to discover - but why should that be so? There may be, but there may also not be, a story about surface properties and their ways of reflecting light that explains why things look red to us if and only if they have those surface properties. It is certainly no objection to the secondary quality account that we cannot specify what it is for something to look red independently of that very concept. There is no *natural* divide between the look of red things and the look of those that border on the orange, or between the look of those things that look red and the look of those that border on the purple. We understand perfectly well that colours form a spectrum or continuum, and we know that we

³³ Cf. L. Wittgenstein, *Remarks on Colour*, III-64.

could choose to cut up this continuum in quite different ways. Similarly, we have many concepts that distinguish shades of red, and form subdivisions within the more inclusive concept. It is not that groups of coloured things already have fixed boundaries around them that our colour concepts match.

Indeed, this is precisely why the linguistic training we receive has something to *explain*. Red things are not orange and not purple, but can be of many different shades between. They come in degrees of brightness and saturation, without ceasing to be red. That this is so is an observation about the grammar of the word 'red', and as such reflects our linguistic training.

It is a consequence of this that there could be languages with very different words for colours. Moreover, there could be words for colours that we do not even see. Our colour concepts respond to the way things look to us. Other beings could have colour concepts that respond to the way things look to them.

We must be clear, however, that not just any concept we might come across would count as a colour concept. On the one hand, there will be a point at which we would no longer recognize the words that others use *as colour concepts*. There are some general constraints on what we count as colour concepts, but there is a more specific limitation, too. What counts as a colour concept is not something that is fixed independently of

our colour concepts. Our colour concepts partly *define* the conceptual space of colour.³⁴

If so, it is not obvious that we *must* think of a concept that cannot be matched with any area of our colour sphere as a colour concept.³⁵ But note that the emphasis can shift, so that we can also say that what counts as a colour concept is only *partly* defined by our colour concepts:

We will, therefore, have to ask ourselves: what would it look like if people knew colours which even our people with normal vision do not know? This question will not in general admit of an unambiguous answer. For it is by no means clear that we would *have to* say of such abnormal people that they knew other *colours*. For there is no generally accepted criterion for what is a colour, unless it is one of ours.

And yet we could think of circumstances under which we would say, "These people see other colours in addition to ours."³⁶

In a later draft of the same passage, Wittgenstein omits the final sentence. One can only speculate about his reasons, but I think that conceptually, this first draft makes a move in the right direction. If there are no generally accepted criteria for what is and what is not a colour, then that means precisely that – there are no generally accepted criteria, and it would seem arbitrary if we were to draw a line around our colour

³⁴ Cf. L. Wittgenstein, *Remarks on Colour*, I-66, I-77, III-32, 86, 155.

³⁵ For the idea of the colour sphere, see P. O. Runge, *Farbenkugel*. The treatise was first published in 1810. Wittgenstein frequently refers to it in his *Remarks on Colour*, reading Runge as making grammatical remarks (cf. I-21 f.). In the colour sphere, the three primary colours red, blue and yellow are placed around the equator, with the mixed colours orange, green and purple in between. White and black are at the poles, thus opening a new dimension in which any mixture of the primary colours is combined with either black or white.

³⁶ L. Wittgenstein, *Remarks on Colour*, III-42; cf. III-127, where Wittgenstein adds an interesting analogy with 'light': "This is like speaking of infrared 'light'; there is good reason for doing so, but one can also declare it a misuse".

concepts and to declare by *fiat* that these colour concepts are the only ones there could possibly be. That would misrepresent the grammar of our colour concepts.

Wittgenstein comes back to the question of a different 'geometry of colour' elsewhere:

"Can't one imagine that certain people have a geometry of colour very different from ours?" Now what that means is: Can't one imagine people who have colour concepts different from ours? And that in turn means: Can't one imagine that people do *not* have our colour concepts, and that they have concepts that are related to our colour concepts in such a way that we would also call them "colour concepts"?³⁷

There is no answer to this question in the passage, but the drift appears to be a 'yes' with some qualifications, and this seems to me to be the correct reply. There could certainly be concepts that are similar enough to ours to qualify as colour concepts, but at the same time different enough from ours to count as different colour concepts. There is nothing inherently problematic about different 'geometries of colour'.

This raises the question of the perspectivity of colour concepts, and with it, the question of their objectivity. We learn colour words by demonstration, and demonstration works against a background of shared sensory responses. But there is no reason to think that our sensory responses are the only ones. They are certainly not the only possible responses, in fact they are not even the only actual ones. This is evident

³⁷ L. Wittgenstein, *Remarks on Colour*, I-66; cf. III-86, 154.

when we consider people who see certain colours differently, cannot see certain colours, or do not see anything at all. We do not want to say that people who see colours differently cannot have our colour concepts. Even so, colour words are used against a background of responses that can be described as seeing colours normally, as opposed to the responses of those who are, for example, blind or colour-blind.

The linguistic training that explains how the conceptual space of colour, once it has been opened, is then structured or divided up will therefore necessarily reflect what may be called the human standpoint, insofar as it will reflect the normal human way of seeing colours. That there is such a normal human way of seeing colours is one of the unfounded foundations of the language game we play with words for colours. Our colour concepts reflect the way things look *to us*, prior to their conceptualization, and this is an important fact about them.

Now it is natural to think that colour concepts reflect something more than that. In addition to the shared responses that make up what I referred to as the normal human way of seeing colours, our conceptual repertoire might well reflect more local interests in discriminating colours. Those interests might differ between cultures or communities. They might even differ between individuals. So if *our* knowledge of colour is to be explained by reference to the linguistic training *we* received, who is that 'we', and what is the explanatory role of that training?

It has been observed that generally, Wittgenstein is not so much concerned with language 'narrowly construed', but usually takes 'language' in a wide sense that suggests something as inclusive as a world view.³⁸ This is true, but it should not blind us to the fact that there are often local differences, and that at times, these local differences make more than a local difference to the bigger picture. Our colour concept 'red' reflects our linguistic practice with all its contingencies and accidents, but also with its background of brute shared experience of the way things look. That is what gives content to the claim that colour concepts reflect the way things look, not to just anybody, but to us.

6 Objectivity and the Autonomy of Grammar

If my account of what Wittgenstein says in § 381 of *Philosophical Investigations* is on the right lines, we can sum up what Wittgenstein is getting at as follows: (1) Language, and the linguistic training we have received, have a special role in explaining our 'geometry of colour'; (2) Colour concepts embody a perceptual perspective and are only adequately understood in those perspectival terms; (3) Colour concepts are tied to a specific point of view, namely our way of perceiving colour; (4) Despite this inherent perspectivity, the correct application of a colour concept can, in favourable circumstances, be or become knowledge of colour; (5) So there are truths or facts about the colours of physical objects.

³⁸ Cf. B. Williams, 'Wittgenstein and Idealism', in: B. Williams, *Moral Luck*, essay 12, esp. sec. 3.

The colour of a flower, for example, is something that we can be said to know.

There are two kinds of worry about this account. One is that if we accept it, we can no longer distinguish between knowledge of colour, which is subjective, and knowledge of objective fact. This is a threat of idealism. Another worry is that even if we can make room for the distinction between subjective and objective, the result is an implausible idealism about the propositions that fall on the subjective side of the divide. I will take these up in turn.

Propositions like 'The moon is smaller than the earth', 'The earth revolves around the sun' or 'Rubies contain chromium oxide' seem to be true independently of us. The moon is smaller than the earth whatever we happen to think about the matter. It is true that rubies contain chromium oxide, and it was true before any human beings existed. But how can they be true independently of us if this kind of knowledge involves our concepts, and we could have classified these things and their relations differently?

'Sun' and 'moon' and 'ruby' are certainly words we use because we find them useful, and we can easily imagine beings who conceptualize large chunks of matter differently or do not think in those terms at all. This is one aspect of the phenomenon that Wittgenstein calls the 'autonomy of grammar'. Grammar answers to the purposes of language.

Language, in turn, answers to the purposes of those who use it. As Wittgenstein puts it at *Philosophical Investigations*, § 497:

The rules of grammar may be called “arbitrary”, if that is to mean that the *purpose* of grammar is nothing but that of language.

The target of this thought comes out more clearly in a passage in *Philosophical Grammar*:

Grammar is not accountable to any reality. It is grammatical rules that determine meaning (constitute it) and so they themselves are not answerable to any meaning and to that extent are arbitrary.³⁹

If this is true, it is true for words that have a use in physics as well as for words for colours. But it does not follow that the propositions we express when using those words are all relative to human interests or perspectives. Colour concepts involve a specific relativity, and so do the facts of colour. But the reason for that is not to be found in the autonomy of grammar, but rather in the way in which a colour word is used. Colour words are used to pick things out according to the way they visually appear. That is why the fact that grass is green involves a relativity to our perspective. ‘The moon is smaller than the earth’ and ‘The earth revolves around the sun’, by contrast, involve no such relativity. It would

³⁹ L. Wittgenstein, *Philosophical Grammar*, p. 184. For a detailed discussion of the notion and its relevance in Wittgenstein, see M. N. Forster, *Wittgenstein on the Arbitrariness of Grammar*.

misrepresent their grammar if they were construed on the model of the language game of colour.

This suggests that when we move from physical or chemical characteristics of an object to its colour, we cross a line between objective and subjective, and this is so even if we maintain, as we surely want to do, that physical and chemical characteristics of objects underlie, determine and explain their colour. Colour is in principle explicable by reference to such things as surface properties of objects, light reflection, chemical processes occurring in the retina, and their complex effects on the brain. Wittgenstein, for one, has no objection: "That colours have their characteristic causes and effects, that we know".⁴⁰ The point is that all such connections are contingent. What it is for something to be red is not to have a certain set of surface properties, but to look like red things look to us.

So there is no unacceptable idealism here. There is of course an element in this account of colour that makes it appropriate to call it a form of idealism. But this amounts to no more than the denial that the facts about the colour of an object are what they are independently of us. The account still leaves room for a form of objectivity, both with respect to facts that involve no reference to a perceiver, and with respect to facts that do.

⁴⁰ L. Wittgenstein, *Remarks on Colour*, III-82.

This is as it should be, because colours are not ascribed to objects in virtue of the way they look to me, and here and now. This is evident given that we can make mistakes about the colour of an object, so there remains a difference between what we think the colour of an object is and the facts about its colour. The account preserves substantial contrasts between knowing and believing, being right and making a mistake, appearance and reality. On the present account, the facts about the colour of an object are normally independent of *particular* experiences of it. They are not independent of colour experience *in general*, which is why there is a crucial difference between primary and secondary qualities.

Even if we grant this point, some will insist that while this is some form of objectivity, it is not objectivity enough. After all, we do not want to say that human agreement decides what the colours of objects are. That would be a deeply problematic conclusion, and it can seem that the appeal to shared responses commits us to this view.

The fact is that this conclusion does not follow. Agreement in our judgments is in general neither a necessary nor a sufficient condition of truth and falsity in any domain of discourse. It is a grammatical point that 'It seems to be so' and 'It is so' can come apart:

From its *seeming* to me - or to everyone - to be so it does not follow that it is so.⁴¹

⁴¹ L. Wittgenstein, *On Certainty*, § 2.

This belongs with *Philosophical Investigations* § 241:

“So you are saying that human agreement decides what is true and what is false?” – True and false is what people *say*; and people agree in their *language*. This is not an agreement in opinions, but in form of life.

Elsewhere, Wittgenstein explicitly applies the point to colour:

Does human agreement *decide* what is red? Is it decided by appeal to the majority? Were we taught to determine colour in *that* way?⁴²

The answer is clearly ‘no’. We learn the names of colours by looking and by listening to what the others say. But once we have learnt the colours of objects, we only need to look. It does not matter what the others say, and what they say need not force us to revise our judgment. Wittgenstein points out that this is reflected in our methods of instruction. When we teach colour words, our basic tool is demonstration:

Colour-words are explained like *this*: “That’s red” e.g. – Our language game only works, of course, when a certain agreement prevails, but the concept of agreement does not enter into the language-game. If agreement were universal, we could be quite unacquainted with its concept.⁴³

So there is no suggestion that human agreement *determines* the facts about colour in any case. But this can be equally problematic. After all, we want

⁴² L. Wittgenstein, *Zettel*, § 431.

⁴³ L. Wittgenstein, *Zettel*, § 430.

to say that there is some conceptual connection between our shared responses and the colour of the object to which we respond. In *Remarks on Colour*, Wittgenstein puts the point as follows:

From its *seeming* to me – or to everyone – to be so it does not follow that it *is* so.

Therefore: From its seeming to all of us that this table is brown, it does not follow that it is brown. But just what does it mean: “In the end, this table is not really brown”? – So it really *does* follow from the fact that it seems brown to us that it is brown?⁴⁴

No, it does not *follow*. But it is not easy to imagine circumstances where it is true *both* that our judgment concerning the colour of an object is at variance with its colour, *and* that we see the object in good light, have unimpaired vision, feel absolutely sure about its colour, and so on.

Two kinds of grammatical connection must be kept apart. First, there is a *general* connection between the way things look to us in certain circumstances and the language game of colour. This kind of connection would allow exceptions, so it is not, in every instance, a necessary link. Appearances of colour and true colours could be said to come apart.

There could, however, also be a more *specific* link between the way things look to us and the language game of colour. Where the specific link holds we can say that it *follows* from the fact that this table seems brown to me or to us that it is brown. Unlike the first, this kind of grammatical connection would not admit of exceptions. After all, to say that ‘This is

⁴⁴ L. Wittgenstein, *Remarks on Colour*, III-96.

brown' *follows* from 'This seems brown to me' or 'This seems brown to us' is tantamount to saying that the two can never come apart.

The point is that even if we accept that there is a general grammatical connection between colours and appearances of colour, we are not committed to accepting the stronger, more specific link. While it makes no sense to ask whether a table that seems to all of us, in normal circumstances, to be brown *really is* brown, we can still say that it does not strictly speaking *follow* from the fact that it seems brown to us that it is brown, because this would express a more specific, and unwarrantedly strong, grammatical connection. This seems to me to be the point that Wittgenstein is after:

Do we not *call* precisely that table brown that appears to normal-sighted people under certain circumstances to be brown? Of course we could conceive of someone to whom things seemed sometimes to be coloured like this, sometimes like that, independently of their colour.

That it seems so to people is their criterion for its *being* so.

Seeming so and being so may of course in exceptional circumstances be independent of one another, but that does not make them logically independent; the language-game does not reside in the exception.⁴⁵

That is why it is important that the vast majority of people agree about the colour of an object.⁴⁶

⁴⁵ L. Wittgenstein, *Remarks on Colour*, III-97-9.

⁴⁶ Cf. L. Wittgenstein, *Remarks on the Foundations of Mathematics*, VI-39.

It may still seem as if the account leaves no room for disagreement where disagreement is to be found, and vital. On a Wittgensteinian view, language is bound up with practices or customs. Is regular dissent from the considered judgment of all other speakers even possible?

Our intuitive reply is that most certainly it is, and I think that Wittgenstein does not have to deny it. He may even imply it in passages such as the following:

Is what we call "following a rule" something that it would be possible for only *one* person to do, only *once* in his life? – And this is of course a note on the *grammar* of the expression "to follow the rule".

The answer is that:

It is not possible that there should have been only one occasion on which only one person followed a rule. It is not possible that there should have been only one occasion on which a report was made, an order given or understood; and so on. – To follow a rule, to make a report, to give an order, to play a game of chess, are *customs* (practices, institutions).⁴⁷

The possibility in question is grammatical. Compare a parallel passage from the *Remarks on the Foundations of Mathematics*:

The application of the concept 'following a rule' presupposes a custom. Hence it would be nonsense to say: just once in the history of the world someone followed a rule (or a signpost; played a game, uttered a sentence, or understood one; and so on).⁴⁸

⁴⁷ L. Wittgenstein, *Philosophical Investigations*, § 199.

⁴⁸ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, VI-21.

Wittgenstein concludes:

What, in a complicated surrounding, we call “following a rule” we should certainly not call that if it stood in isolation.

Language, I should like to say, relates to a *way* of living.

In order to describe the phenomenon of language, one must describe a practice, not something that happens once, *no matter of what kind*.

It is very hard to realize this.⁴⁹

As before, there is no reference to communities. The qualification Wittgenstein describes concerns contexts, practices or customs.

It might be replied that talk of practices or customs already involves a reference to communities. But if it does, this needs a further argument. On the face of it, there is no reason to deny that individuals can be said to be involved in customs or practices in isolation. It is perhaps no accident that Wittgenstein does not ask whether it is possible for only *one* person to follow a rule. Rather, Wittgenstein asks whether it is possible for only *one* person to follow a rule only *once* in his life. He replies by saying that it is not possible. This suggests at least the possibility that an isolated individual can be said to follow a rule, provided that this is not an isolated event, something done only once, on a single occasion. If Wittgenstein meant to rule that out, why did he not say so? Wittgenstein *could* have

⁴⁹ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, VI-33-4.

said that it is impossible for one person to follow a rule if he had thought it true. But he did not do that.⁵⁰

Essentially the same, and one might say evasive, answer to the question is found in a discussion of calculating. That there is consensus in performing calculations “is part of the phenomenon of our calculating”.⁵¹ But that does not mean that isolated individuals cannot perform a calculation:

But what about this consensus — doesn't it mean that *one* human being by himself could not calculate? Well, *one* human being could at any rate not calculate just *once* in his life.⁵²

That this point has an important consequence for ethics is illustrated by an argument once contemplated by Blackburn. We want to say that a person may have everyone against him in a moral dispute and be right, and it is obvious that this point is of more than merely theoretical importance.⁵³

What Wittgenstein says about rule following, however, does not exclude that possibility. It would rule out systematic disagreement if and only if following a rule were something that could only be done in communal settings, so that regular dissent over the correct application of a

⁵⁰ Cf. H. Putnam, 'Was Wittgenstein Really an Anti-realist about Mathematics?', pp. 143-9; see also B. Stroud, 'Wittgenstein's "Treatment" of the Quest for "a language which describes my inner experiences and which only I myself can understand"', and 'Private Objects, Physical Objects, and Ostension' in: B. Stroud, *Meaning, Understanding, and Practice*, chs. 5 and 13.

⁵¹ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, III-67.

⁵² L. Wittgenstein, *Remarks on the Foundations of Mathematics*, III-67; cf. III-69.

⁵³ S. Blackburn, 'Reply: Rule-Following and Moral Realism', in: S. Holtzman and C. Leich (eds.), *Wittgenstein: To Follow a Rule*, essay 6, p. 173; cf. S. Blackburn, *Ruling Passions*, p. 74.

term would not be so much as possible. But there is no reason to accept this.

The root of the misunderstanding lies precisely in the thought that following a rule is tied to a communal practice. Thus, Blackburn suggests that “Wittgenstein regarded consensus as a necessary as well as perhaps sufficient foundation for the notion of correctness”.⁵⁴ If my interpretation is correct, Wittgenstein says no such thing. Agreement or consensus are clearly the rule rather than the exception, and it is important that they are. But agreement is not presupposed in the idea of practices or regularities or customs, including those involved in following a rule. There is no suggestion that consensus is either necessary or sufficient for the application of the notion of correctness.⁵⁵

7 Explanation and the Absolute Conception

It is a natural thought that such knowledge as there may be “is knowledge of what is there *anyway*”.⁵⁶ Knowledge of what is there anyway, in turn, is naturally construed as knowledge of what is true independently of us. At the limit, knowledge of what is there anyway is knowledge of the world as it is independently of any particular point of view or experience of it.

We have seen that knowledge of colour calls for an important qualification to that claim, and the same goes for all the other properties

⁵⁴ S. Blackburn, ‘Reply: Rule-Following and Moral Realism’, p. 173.

⁵⁵ Cf. L. Wittgenstein, *Zettel*, § 429.

⁵⁶ B. Williams, *Descartes*, p. 64.

that essentially involve a reference to perceivers. Colour concepts are grammatically connected with the way things look to us. If knowledge of the world as it is anyway is knowledge of the world as it is independently of our experience, knowledge of colour is not knowledge of the world as it is anyway. But as we have also seen, this does not mean that in the particular case of an unusual perceiver or unusual circumstances, appearance and reality could not be said to come apart.

This raises the question whether there is such a thing as knowledge of the world as it is anyway, and if there is, how that knowledge relates to knowledge that is perspectival. In his study of Descartes, Williams argues that in order to explain our knowledge, both perspectival and otherwise, we need what he calls the 'absolute conception of reality'. What we are looking for is

a conception of reality corrected for the special situation or other peculiarity of various observers, and that line of thought leads eventually to a conception of the world as it is independently of the peculiarities of any observers.⁵⁷

Provided that we are not idealists, this line of thought will yield the conception that we need, namely a conception of the world as it is independently of *all* observers.

As Williams is the first to point out, however, there is a very real risk that the requirements of the absolute conception cannot in fact be met.

⁵⁷ *Ibid.*, p. 241.

For what we need is not merely a conception of the world without observers. What we need is conception that includes (a) a conception of the world as it is without observers, (b) a conception of its relations to different observers, and (c) a conception that explains, for every observer, not only his share of knowledge of the world as it is anyway, but also his or her local representations. This leads to a demand of complete transparency, because the absolute conception must make it possible to explain how it itself can exist:

This conception is not something transcendental, but is an historical product of consciousness in the world, and it must at least yield a comprehension of men and of other rational creatures as capable of achieving that conception. It thus involves a theory of knowledge and of error: it serves to explain how members of these species might come to have or fail to have a true conception of themselves and of the world.⁵⁸

This is a tall order.

I will not, however, enter into a detailed discussion of the absolute conception or our prospects of attaining it.⁵⁹ Instead, I will try to extract from it a way of understanding the contrast between 'subjective' and 'objective' as it features in the distinction between primary and secondary qualities.

⁵⁸ *Ibid.*, p. 246.

⁵⁹ For a comprehensive discussion of the absolute conception and the distinction between 'subjective' and 'objective' that it suggests, see A. W. Moore, *Points of View*. For further clarification and defense of the original idea, especially as it applies to ethics, see also B. Williams, *Ethics and the Limits of Philosophy*, ch. 8, and 'Philosophy as a Humanistic Discipline', in: B. Williams, *Philosophy as a Humanistic Discipline*, ch. 16.

It is obvious that the absolute conception suggests a way of drawing a distinction between primary and secondary qualities.⁶⁰ As we shall see, this distinction does not in fact presuppose the absolute conception. The distinction will, however, still enable us to understand the difference between objective elements or features of reality such as those described by physics, and those that fall short of objectivity in this ambitious sense. Ethics, I shall claim, falls into the latter, essentially subjective category.

Consider realism about colour. Imagine that colours are basic features of reality that we perceive when we see colours, and that these features are basic in the sense that they cannot be explained as local, perspectival, and essentially subjective responses to an objective reality that is described by physics. In other words, imagine that colours exist in addition to primary qualities of objects that already feature in the absolute conception. The suggestion would not have to be that colours feature in the absolute conception *as* physical qualities. Nor would the objectivity of colour commit us in any way to the view that *like* physical qualities, colours must in principle be visible from any point of view. The facts about colour may only be accessible to those with colour vision, exist independently of all physical properties of objects, and resist all further explanation. Still, if realism about colour were correct, colours *would be*

⁶⁰ Cf. B. Williams, *Descartes*, pp. 237-46.

there to be seen – they would be there to be right or wrong about, and in that sense, they would have become part of the absolute conception.

The point is that this is simply not the way we think of colour. Putting it crudely, the account of colour that we accept goes like this.⁶¹ Physical objects absorb and reflect light, showing a particular pattern. Some of these patterns are detected and discriminated by our sensory apparatus, in particular, the nerve cells on the retina. In these cells, light of different wavelengths triggers different chemical and physical reactions which affect the nervous system in certain regular and systematic ways. On the basis of encoded information about wavelengths that is processed in the retina and passed on through the optic nerve, the brain eventually creates a visual image of the coloured object.

It goes without saying that the story is vastly more complicated than that and far from being fully understood. But for our purposes the details do not matter. What does matter is that this process can be described, at the most abstract level, as the interaction of a physical reality that is objective – involving such things as molecular structures, reflections of light and wavelengths – with a subject like the human being that responds to physical inputs of this kind in characteristic ways. We produce visual representations of objects as coloured in response to light and surface properties. What explains colour perception is that objects

⁶¹ For some critical queries about this account, see B. Stroud, *The Quest for Reality*. In my view, none of the criticisms Stroud presents ultimately undermine the story, and with it, the subjective view of colour.

reflect light that hits their surface in regular and systematic ways. The point is that in a sense, *this is the whole story about colour*, even though we must of course admit that all the details remain to be filled in.

Given the possibility of explanation, there is every reason to believe that the absolute conception has no room for secondary qualities such as colours. Secondary qualities depend on psychological factors. Moreover, they depend on psychological factors in such a way as to make clear that secondary qualities are a function not just of consciousness, but of the peculiarities of an individual or species.⁶² Thus, we can easily understand the thought that a thing could seem one colour to one person, another colour to another, and that it could seem coloured in one way to members of one species, coloured in a different way to members of another. In understanding, even sketchily, and at a general reflective level, why various things appear variously coloured to various observers, we find that we have left behind any idea that, in some way which transcends those facts, they 'really' have one colour rather than another.⁶³

8 Primary and Secondary Qualities

The important result that emerges from this discussion is that the distinction between primary and secondary qualities turns on more than the idea of different perceptual perspectives. It is backed up by a substantial notion of explanation. In particular, it is backed up by a

⁶² B. Williams, *Descartes*, p. 241.

⁶³ *Ibid.*, p. 242.

substantial notion of scientific explanation. This notion provides the absolute conception with the explanatory power that it has in understanding secondary qualities like colour. As Williams says, the picture offered by natural science would explain the phenomena. It would explain them, moreover, even as they present themselves in other, more local representations. It is this consideration that gives the content to the idea, essential to the traditional distinction, that the scientific picture presents the reality of which the secondary qualities, as perceived, are appearances.⁶⁴

What I now want to suggest is that this way of drawing a distinction between primary and secondary qualities is independent of the absolute conception. Suppose the absolute conception is impossible. Williams considers this possibility:

If scientific enquiry turned out not to yield what the present line of thought requires it to yield, then perhaps our everyday distinctions will turn out to go as deep as anything goes. But we cannot assume that that will be so. Moreover, paradoxically, it would be an affront to other parts of our everyday thought if it did turn out to be so.⁶⁵

We should not simply assume, in other words, that if the absolute conception is unattainable, our everyday distinctions will go as deep as anything goes. One of the things we must accept, unless we are idealists, is that there is a physical reality of which we are a part. If we have to take for

⁶⁴ *Ibid.*, pp. 244 f.

⁶⁵ *Ibid.*, p. 242.

granted that there is a physical reality, why do we have to accept anything more?

It is true that if we do not want to accept anything more, we have to be able to *explain*, if only in outline, how other, problematic features of reality are the product of a complex interaction of a physical reality with perceiving subjects. This is not a trivial task, but there is no reason to think that this explanation would be ruled out even if we suspend judgment on the possibility of forming the inclusive, absolute conception. So there is every reason to believe that we can hold on to a substantial distinction between primary and secondary qualities, and that we can hold on to this distinction in the terms in which it could be drawn against the background of the absolute conception.

If this is correct, there is every reason to believe that there is a substantial difference between primary and secondary qualities, and that this difference does not presuppose the possibility of the absolute conception. What it presupposes is a conception of the world as physical, and a set of reliable methods of scientific explanation. This idea gives us a handle, just as Descartes thought it did, on the distinction between primary and secondary qualities, and with it, on the distinction between subjective and objective. This is the sense in which it might turn out to be true what Democritus said in the 5th century BC:

Perspectivity and Colour

sweet is by convention and bitter by convention, hot by convention, cold by convention, colour by convention; in truth are atoms and void.⁶⁶

In the next chapter, I will explore how this distinction could apply to ethics.

⁶⁶ C. Bailey, *The Greek Atomists and Epicurus*, pp. 178 f.; cf. B. Williams, *Descartes*, p. 242.

2

The Truth in Moral Realism

Given the distinction between primary and secondary qualities, I will now apply the view of perspectival concepts outlined in the previous chapter to ethical concepts. This completes my basic argument in favour of subjectivism. The upshot of the argument is that the facts or truths we come to know in ethics are the result of our actions and ways of looking at the world, as opposed to being independent features of the world we look at. There can be moral truth and moral knowledge, but there is no moral objectivity.

The argument in favour of subjectivism is the inference to the best explanation. The basic problem with a claim to objectivity in ethics is that there is no independent reason to believe it. The subjectivist hypothesis explains the relevant phenomena equally well, but does so without

recourse to inexplicable objective moral qualities. Indeed, subjectivism can explain many phenomena better than the realist alternative.

As for expressivism, even if the inference is accepted, there is no need to deny that there are moral facts, because there is no longer anything mysterious or metaphysically suspicious about them. Given our linguistic practices, and some further conditions that we have good reason to think are satisfied, moral facts arise. These facts are not objective, but for all that, they are moral facts. That, I believe, is the truth in moral realism.

1 The Application to the Ethical

Following Bernard Williams, I suggested that the traditional distinction between primary and secondary qualities can be spelled out in terms of the different styles of explanation. Knowledge of colour is explicable in terms that make no reference to objectively existing colours, whereas knowledge of physical reality, unless we contemplate idealist hypotheses, is not. How would this account of secondary qualities apply to ethics? Certainly not directly, which I take to be established by the following related reasons.

(1) The subjective conception of secondary qualities is built on the idea that colour perception, for example, is explicable in terms of our perceptual psychology on the one hand, and its response to a world as characterized in terms of primary qualities on the other. If subjectivism

holds that 'ethical qualities' are 'added to' or 'projected onto' the world, then 'the world' has to be already construed in a psychologically and socially richer sense than the world onto which secondary qualities are projected.¹ What you do may be wrong because it hurts, or helpful because it is encouraging. There is no hope of making sense of ethics if all we are given to make sense of it is a world described in terms of physics.

(2) There is a direct and sufficient basis for the application of a colour concept in what may be called the phenomenology of colour. Colour words apply to things in virtue of the fact that they look a certain way to us. There is no clear parallel to this in ethics. There is no direct basis in experience for terms such as 'good', 'courageous', 'wrong', or 'cruel' because there is no such thing as their distinct phenomenology. Wrong things are not picked out just in virtue of the way in which they visually appear. 'Wrong' applies to certain acts and states of mind or character, but there is no reason to expect that they have some distinctive feature, other than that they look, precisely, wrong.

(3) This brings us to a third and even more important point, namely that the grammar of ethical words is *layered* in a way in which the geometry of colour is not. Things look either red or green or white or black, but not red and green and white and black all over in the same place at the same time. What you do, however, may be shameful but obligatory, courageous but premature, kind but inappropriate all at the

¹ Cf. B. Williams, 'Who Needs Ethical Knowledge?', in: B. Williams, *Making Sense of Humanity*, essay 17, p. 177.

same time. Moreover, what you do may be inappropriate precisely because it is kind, for example because in the circumstances a little rudeness was called for.² Similarly, what you do may well be bad in one respect, good in another. In all these ways, the rules that govern the use of ethical language are vastly more complex than the rules that govern the use of colour words. If so, there is no interesting sense in which the answer to 'How do I know that this action is wrong?' would be 'I have learnt English'. Even on the most inclusive interpretation of what it is to be a competent speaker of English, that cannot be enough.

(4) These grammatical complexities are also reflected in the meaning of our words. In the case of colour, we can explain what we mean by 'magenta' or 'aquamarine' by producing samples. We can add samples of colours that are nearby in the spectrum but no longer fall into those categories, but once we have done this, we can only wait for the response and hope for the best. There is no such thing as a further *explanation* of colour words.³

The situation is not at all like that in ethics. Here, our resources for the explanation of ethical terms are endless and as varied as ethical life itself. We can tell stories, pull faces, draw pictures, act parts. We can give more or less detailed descriptions, turn to history, or translate terms into more familiar idioms. We can describe feelings, actions and reactions, traits of characters, features of situations, and so on. Indeed, what

² Cf. P. Foot, 'Moral Arguments', in: P. Foot, *Virtues and Vices*, essay 7, pp. 102-6.

³ Cf. L. Wittgenstein, *Remarks on Colour*, I-68.

someone means by 'shady character' or 'insult', for example, will be revealed not only in applications of these words, but even more clearly in the kinds of stories or examples that the person gives to illustrate their meaning.

(5) Finally, there is a difference in the grammar of ethical words compared to words for colours in their relations to perceivers. As Blackburn points out, it is plausible to think that if we were to change so that everything that had appeared blue came to appear red to us, this is what it would be for the world to cease to contain blue things, and to come to contain new red things. But if everyone comes to think that it is permissible to mistreat animals, this does nothing at all to make it permissible. It just means that everyone has become insensitive to the fact that it is wrong to mistreat animals.⁴

So the secondary quality analogy will not apply directly. At the same time, there is a deeper and much more important parallel. The parallel consists in style of explanation. Blackburn himself seems to wish to hold on to the idea that there is such a parallel, given his Humean picture of the nature of morality:

On this view we have sentiments and other reactions caused by natural features of things, and we 'gild or stain' the world by describing it as if it contained features answering to these sentiments, in the way that the niceness of an ice cream answers to the pleasure it gives us.⁵

⁴ S. Blackburn, 'Errors and the Phenomenology of Value', in: S. Blackburn, *Essays in Quasi-Realism*, essay 8, p. 160.

⁵ *Ibid.*, p. 152.

In ethics as well as in perceptual psychology, we must ask what need there is to go beyond the physical. The question of their objectivity has already been answered negatively in the case of colours. The phenomenon of seeing colours is in principle susceptible to scientific explanation, even if we do not know all the details and probably never will. Once we accept that there is a complex physical reality, and accept that this reality includes psychological, social and historical dimensions, the question arises why we should be willing to admit anything more. Colours are no part of the world as it is anyway. Moreover, they are not objective in a sense that would essentially involve a relativity to a perceiver, but still support the idea that colours are a fundamental feature of reality that resists all scientific explanation. Should we not say the same in ethics? One thing that is clear in itself, and has been reinforced by the discussion, is that the appropriate explanations will be vastly more complex, both in extent and structure, than the explanations we can give of colours and the way they visually appear. But that is not, in itself, a reason to reject the possibility of explanation.

2 Minimalist Maxims

I suggested that there is an intermediate position in ethical theory that can usefully be distinguished from both moral realism, which is committed to truth as well as objectivity, and from expressivism, with its qualified denial of both. The idea around which this account revolves is that moral

facts, like facts about colours, can be explained by reference to local linguistic practices. The linguistic instruction we receive is bound up with standards for the application of a range of moral words, and when the standards for the application of such words are satisfied, moral facts arise. These facts are perspectival, because they do not answer to some objective moral or evaluative dimension of reality. The linguistic practice that gives rise to moral facts is located in one natural and social world as opposed to another, and embodies complex cultural traditions.

It is important that this is not at all the view suggested by some writers in the realist tradition. Wiggins and McDowell, for example, also think of their view as an intermediate position between overly ambitious forms of platonism and subjectivism.⁶ They admit that what is good or bad or right or wrong is not something that is fixed independently of us, but insist that this does not undermine a basic claim to objectivity. This, however, makes them realists by my classification, just as my view is clearly subjectivist by theirs. Insofar as my view is an intermediate one at all, it is better described as intermediate between realism and expressivism. There is no compromise about the crucial issue, that of objectivity in ethics.

If this view is correct – and I have only sketched it in its barest outline – there could be moral truth as well as moral knowledge. But there would be no moral objectivity, if objectivity involves something that goes

⁶ Cf. J. McDowell, *Mind, Value, and Reality*, and *Mind and World*; D. Wiggins, *Needs, Values Truth, and Ethics*.

beyond the bare admission that a moral judgment may be true or false.

Why should we believe this?

One reason for believing it is this kind of theory is, in a substantial sense, naturalistic. Bernard Williams has described the kind of attitude that motivates the naturalist approach as follows:

First, to the question 'how much should our accounts of distinctively moral activity add to our accounts of other human activity?' it replies 'as little as possible', and the more that some moral understanding of human beings seems to call on materials that specially serve the purposes of morality - certain conceptions of the will, for instance - the more reason we have to ask whether there may not be a more illuminating account that rests only on conceptions that we use anyway elsewhere.⁷

A naturalistic attitude of this kind is found in Hume, and, as Williams points out, in Nietzsche. The idea is to try to make sense of morality in terms that are already established and in use elsewhere, and to treat those elements of moral theory that serve moral purposes with special scepticism.

It is important to be clear that this demand for minimalism in moral psychology is not just an application of the Occamist desire for economy.⁸ The perspective from which explanations are assessed as adequate or economical must be holistic. We need a guiding sense of what materials to use, but the task is, and remains, to make sense of moral life as we experience it. What is at issue is not the application of an already defined

⁷ B. Williams, 'Nietzsche's Minimalist Moral Psychology', in: B. Williams, *Making Sense of Humanity*, essay 6, p. 68.

⁸ *Ibid.*

scientific programme, but rather an informed interpretation of some human experiences and activities in relation to others.⁹ For all its openness and indeterminacy, and indeed because of it, this seems to me a helpful attitude to take.

What are the consequences of adopting it? We start with the idea that the world of which we are a part is physical, whatever that may mean. We know that it has complex psychological dimensions, and that it is a world with social and historical dimensions. Moreover, we know that we are involved in linguistic practices, and that we constantly interact with other human beings.

Now it is simply not obvious that in order to make sense of ethics at this most general level we need anything more than, or radically different from, those kinds of materials. We have become used to the idea that mastery of language is essentially the mastery of a technique. These techniques enable speakers to internalize and follow rules, and such rules govern the use of moral words as much as that of any other.

But mastery of moral words enables us to find our way around a *particular* social world.¹⁰ The use of such words will accordingly reflect some attitudes and values that not everyone must share, and external observers may, on that account, be unwilling to use them. ('Chastity' and 'honour' are examples of ethical words that many of us no longer use as our own, or the use of which has changed dramatically.)

⁹ *Ibid.*

¹⁰ Cf. B. Williams, *Ethics and the Limits of Philosophy*, p. 150.

All the same, for those who speak the language, and speak it not as detached observers, but use ethical conceptions as their own, linguistic rules will serve to establish standards. These are standards both for action and for the application of ethical expressions. Once those standards are met, and other conditions are satisfied, there can consequently be both moral truth and moral knowledge, but so far, there has been no suggestion that the language game of ethics answers to or incorporates a normative dimension of reality.

Still, it may be felt that the appeal to naturalism is not quite enough. After all, the naturalist perspective is highly controversial, so it seems desirable to be able to support subjectivism without taking recourse to such controversial notions. This, I want to argue, can be done. The argument to which I think we should appeal is the inference to the best explanation.

Ultimately, naturalism and the inference to the best explanation are not to be kept apart. As Williams describes it, naturalism is itself a version of the inference. But there is still something to be gained from using strategies that are successfully employed elsewhere, and the inference is one of those strategies.

3 The Inference to the Best Explanation

What is the structure of the inference? As Gilbert Harman points out, in making the inference to the best explanation one infers from the fact that a

certain hypothesis would explain the evidence to the truth of that hypothesis. In general, there will be several hypotheses which might explain the evidence, so one must be able to reject all such alternative hypotheses before one is warranted in making the inference. Thus one infers, from the premise that a given hypothesis would provide a 'better' explanation for the evidence than any other hypothesis, to the conclusion that the hypothesis is true.¹¹

How could the inference to the best explanation be applied in the present context? The conclusion is that subjectivist accounts provide a better explanation of the ethical phenomena. In particular, I think, the subjectivist account of moral judgment provides a better explanation (a) of the idea that there are *moral facts*, (b) of our *knowledge* of those facts, and (c) of their *holistic structure*. Let me briefly say something about all three.

First, the subjectivist can agree with the moral realist that there are such things as moral facts in moral contexts. That does not, by itself, put him in a favourable position compared to the realist. Both accept the notion of a moral truth or fact. The difference and advantage comes in at a different level, namely that of explanation. The subjectivist can account for moral facts in a way in which the realist cannot, and this is one essential difference between the two positions.

For the subjectivist, moral truths or facts are susceptible to explanation. Moreover, they are explicable in terms that we already use elsewhere. As we have seen, he appeals to such obvious facts as that we

¹¹ G. Harman, 'The Inference to the Best Explanation', p. 89.

live in a complex social world, that there are distinctive human needs and interests, that we form evaluative attitudes and crucially, that we are involved in complex linguistic practices. On any realist view, by contrast, moral facts are simply there, independently of our established ways of using moral language. This is unsatisfactory, given that they can be explained in more familiar terms by the subjectivist model.

Some realists take comfort in the fact that even the existence of the physical is something we must ultimately simply take for granted. But if that is so, the argument continues, why should we not simply accept the fact that there is a normative order in the universe, just as we must ultimately simply accept that there is such a thing as physical reality?

The answer is that even if our reasons must run out at some point, they do not run out that soon. Given that we have to accept that there is a physical order, it is still an open question whether or not we have to accept anything *more*. If the explanation of the language game of ethics I have sketched can do without a further unexplained and inexplicable dimension of reality, as it claims to do, then that is, in that respect, a reason to accept it.

A similar point applies to moral knowledge. On the subjectivist view, moral knowledge is something we have, where we have it, in virtue of the fact that we master moral language and certain other conditions are satisfied. Our moral knowledge will be limited, because moral concepts can have truth conditions that are not easily satisfied. Some, for example the most abstract ones like 'right' and 'wrong' and 'good' and 'bad' or

'useful' or 'desirable', may not often qualify as candidates for knowledge at all.¹² But even here, our mastery of the rules that govern their application is something on which we can count despite the fact that there will be more room for doubt and disagreement, so that there remains a difference between knowing that a given act is right or overall permissible and knowing that a given act is brave, or treacherous, or noble.

On the realist view, by contrast, it must always remain doubtful that we have moral knowledge, because there is always the possibility that even firmly held beliefs that reflect well-established linguistic practices and the moral facts can come apart. This is even true in the case of Kantian and Aristotelian attempts to provide a foundation for morality in practical reason. There is always the possibility that we misread the demands of practical reason, or misapply its general demands in the particular case at hand. With the possibility of objective moral knowledge comes the unsettling risk that we may get things wrong.

The subjectivist hypothesis, by contrast, leaves no room for global moral skepticism. This is an advantage, but it is an advantage not so much because it would yield moral reassurance. This is something that a moral theory is extremely unlikely to do, and the fact that moral realism hopes for some such reassurance should make us more suspicious, and not less.

¹² Cf. B. Williams, *Ethics and the Limits of Philosophy*, chs. 7-8. Note that my doubts about the availability of knowledge that involves 'thin' concepts such as 'right' or 'wrong' are not doubts about the possibility of knowledge that involves thin concepts *per se*. In this respect, the account I envisage differs from that suggested by Williams. I think there is no obstacle, in principle, to knowledge that involves thin concepts. Everything depends on the situation in which these concepts apply, and on the evaluative perspective from which they are applied.

Realist views of the physical are inescapable because there is no credible alternative. Realist views of ethics are eminently open to dispute.

Finally, we need to consider the conceptual relations between moral judgments themselves. We take it for granted, for example, that justice is a virtue, that promises should be kept unless there is good reason not to, and that the fact that an act would have beneficial consequences counts as a reason in its favour, although it need not be the decisive one. These are platitudes, but they are interesting platitudes for the following reason. *How can we be so sure that justice is a virtue? How can we be so sure that keeping promises is the right thing to do unless there is some good reason not to? How can we be so sure that beneficial consequences count in favour of an action?*

The problem is more general. The question is: What follows from the fact that one moral judgment is true for the truth of another? If someone is just, we assume that insofar as he is just, he is a virtuous person. On the subjectivist hypothesis, this is not a deep difficulty. Moral knowledge is knowledge of rules, and the mastery of such rules sometimes licenses such inferences. No doubt these rules are highly complex. But that is not, as such, a good objection to the claim that we make moves, sometimes imaginative, original, creative moves, within a complicated system of moral considerations and concepts that we have mastered and internalized.

On the realist hypothesis, by contrast, there is first of all the question of how we can know what is beneficial, kind or just, given that

these facts are to be thought of as objective. But there is a further question. Given that an act is beneficial, kind or just, what are we entitled to infer from that? Once again, objectivity inevitably entails the possibility of error.

Now the point is not, or at least not quite simply, that subjectivists can hope to rule out moral error. It is certainly possible for someone to master large segments of moral language, yet to lack the required skills of judgment, or experience of life, or knowledge in other areas that go into the idea of the good judge in moral matters. There is, and remains, the possibility of moral error, which is one fact of our ethical experience that must simply be accepted.

All the same, and regardless of the fact that the application of moral terms requires complicated skills of judgment and builds on substantial knowledge and experience, there are legitimate claims to ethical knowledge and legitimate formal and informal inferential moves within the ethical framework. The point is that once we are entitled to those moves, there is no *further* question to be answered. We frequently give reasons why we think that certain aims are worth pursuing, or why some things are good or right or reasonable to do. Reasons of this kind frequently include further ethical considerations. We might say, for example, that part of the reason why justice is valuable lies in its instrumental value, that is, in the undoubted importance of justice for our welfare. We can say this, moreover, without having to reduce all fundamental reasons that we recognize in ethics to a single type, for

example those of welfare. Of course welfare *is* part of the reason why justice is so important. It features prominently in the set of reasons to which we appeal when someone challenges our belief that justice is a virtue.

For the realist, by contrast, there will always be a question of how, if it is just a basic fact that this act is just and just acts are to be pursued, reasons of those other kinds can so much as get a hearing. We can think of the situation in terms of a dilemma. Either the fact that justice is valuable admits of further explanation, or it does not. If it does not, then we must admit that the contribution justice makes to general welfare is not, after all, part of the explanation of why justice is valuable. We are left with the blank assertion that some things are just and, in virtue of the fact that they are just, to be pursued. This seems to cut off justification where our actual justifications have not yet run out. But if we admit that there can be justification, and that the fact that justice is valuable is partly explained by the more basic fact that just acts contribute to general welfare, then this kind of explanation needs an explanation. On the realist hypothesis, it is not at all obvious what that explanation could possibly be.

On the subjectivist hypothesis, by contrast, inference and justification are just what they appear to be, namely a thoroughly holistic matter. We can cite reasons for ethical judgment, and there is no further question why the fact that justice makes a contribution to our welfare counts in favour of just acts and institutions. General welfare is one of the things *that count* in favour of just acts and institutions, and this kind of

linguistic knowledge is articulated when we say that part of the reason why justice is worth aiming for lies in its contribution to something quite different that we value. In this way, the subjectivist hypothesis can rule out what may be called inferential moral scepticism, and to that extent, it provides a better explanation of the ethical phenomena.

There is a further point concerning holism and ethical justification. It is a plain fact of ethical experience that the moral features of an act or object are related to the facts of nature, but the realist leaves it entirely open how that kind of relation is to be understood. Thus, J. L. Mackie asks:

What is the connection between the natural fact that an action is a piece of deliberate cruelty – say, causing pain just for fun – and the moral fact that it is wrong? It cannot be an entailment, a logical or semantic necessity. Yet it is not merely that the two features occur together. The wrongness must somehow be ‘consequential’ or ‘supervenient’; it is wrong because it is a piece of deliberate cruelty. But just what *in the world* is signified by this ‘because’?¹³

On the view that I am suggesting, this is not a deep difficulty. Certain actions *count* as wrong, and there is no reason why a natural fact should not count as a reason that speaks for or against that action. It will of course be difficult to give detailed descriptions of the use of any moral term. But there is no *further* difficulty involved in understanding the explanatory link between ‘This causes pain’ and ‘There is reason not to do it’, or between ‘This is a piece of deliberate cruelty’ and ‘This is wrong’. The realist, by contrast, will ultimately have to say that this is simply how things are. This is not so much an explanation as an outright rejection of

¹³ J. L. Mackie, *Ethics*, p. 41.

attempts at explanation. This is unsatisfactory, given that an explanation is available.

4 What About Expressivism?

This leaves us with a choice between subjectivism as I have described it and expressivism. It may be said that so far, my subjectivist is not at a clear advantage over the expressivist. Both eschew all deeper metaphysical commitments while preserving talk of truth and even talk of objectivity, provided that we read the latter expression as specifying the content of a moral attitude. So on both views, ordinary moral discourse will, to a large extent, remain intact. Moreover, both appeal to arguments from explanation to support their view.¹⁴

I confess that I am not too clear about the relations between my version of subjectivism and expressivism, mainly because I am not too clear about the range of claims essential to expressivism. From the point of view of moral realism, they will certainly seem similar, if not the same. But there is at least one important difference. While the subjectivist denies the claims of moral realism, he does not say, across the board, that we must think of moral facts, moral beliefs and moral knowledge as something other than they seem to be. Moreover, subjectivism seems to me to be more flexible in its account of what we do when using moral

¹⁴ Cf. S. Blackburn, *Ruling Passions*, p. 80: "The fact that there is a cannonball on the cushion explains why it is sagging in the middle. The fact that kindness is good explains no such kind of thing".

language. Thus, he is not committed to construing all ethical thought as the expression of an attitude that is neither true nor false.

As we have seen, expressivists accommodate both truth and objectivity. But as we have also seen, they do so in a special way. It is not that talk of truth and objectivity can only be taken with a pinch of salt. Our ethical commitments can be called true or false and even objective, but at a deeper level of analysis, they turn out to be something different than they seem to be. Thus, Blackburn writes:

Just *because* of minimalism about truth and representation, there is no objection to tossing them in for free, at the end. But the commitments must first be understood in other terms.¹⁵

The final clause is crucial. If we were to give it up, expressivism would collapse into moral realism.

Nor can the expressivist turn the argument around and say that unless we accept the expressivist analysis, we cannot even articulate a realist alternative. For if he were to do that, he would end up pulling the carpet out from under his feet. The very basis of expressivism is a distinction between attitudes that have, and attitudes that lack, objective truth conditions. That is why expressivists say that before we can pull off the deflationist manoeuvre, ethical commitments need to be understood 'in other terms'. But if we are to think of them 'in other terms' before we can call them true or false, then there had better be other terms for us to

¹⁵ *Ibid.*

use to think about them. It is certainly not enough to say that on the expressivist account, moral judgments express moral attitudes. This is a platitude, as long as we cannot say more precisely what, if anything, distinguishes the attitudes in question from representations.

Sometimes Blackburn seems to think that we can draw the relevant distinction by appealing to the notion of representation. He credits the important insight to the later Wittgenstein:

Nobody can deny that Wittgenstein tried to understand many areas of discourse in terms other than those of 'representing the facts'. According to Wittgenstein, whole areas of language that look as if they are dedicated to describing how things are must be understood in other terms.¹⁶

One might also think of the *Remarks on Colour*:

To say that a sentence which has the form of information has a use, is not yet to say anything about the *kind* of use it has.¹⁷

This is a familiar theme. '2 + 2 = 4', 'One plays patience by oneself', or 'This (pointing at the standard metre) is exactly 1 m long' all look like genuine propositions with representational content. It is a philosophical achievement to see them for what they are, namely rules of representation rather than representations. They are standards of measurement, not things measured.

¹⁶ *Ibid.*, p. 77.

¹⁷ L. Wittgenstein, *Remarks on Colour*, III-336.

But at the same time, there is no objection to calling them 'true' or 'false' if we apply the expressivist analysis, and it is certainly possible to have the *belief* that $2 + 2 = 4$ or that one plays patience by oneself. But if these attitudes can be beliefs, one would expect that they can be knowledge. Wittgenstein, for one, does not deny that:

"I know that he arrived yesterday" – "I know that that $2 \times 2 = 4$ " – "I know that he was in pain" – "I know that there is a table standing there".

In each case I know, only something different? *Of course* – but the language games are far more different than we are conscious of, considering these sentences.¹⁸

And now there is a problem. If there is knowledge that $2 + 2 = 4$, or that one plays patience by oneself, then surely something is known, and that something must be a propositional content. But if that is so, there is no objection to saying that this area of discourse has representational content and is used to 'state the facts'. So how can we say that some propositional expressions are, while others are not, used to represent the facts?

The answer must be that at a deeper level of analysis, ethical propositions are something other than what they seem to be. As before, the question is how to capture that difference. Sometimes, Blackburn appeals to differences in explanation, or a difference in 'adaptive mechanism'.¹⁹ Generally, Blackburn argues that the contrast might be

¹⁸ L. Wittgenstein, *Remarks on Colour*, III-311 f.

¹⁹ S. Blackburn, *Ruling Passions*, p. 80.

drawn by means of a functionalist theory of the mind that identifies different causal roles for beliefs and other kinds of attitude.²⁰

Here it seems to me to be important that subjectivism as I understand it has no need to get involved in these debates. There is no general objection to the claim that there are moral facts, and consequently there is no need to buttress a propositional analysis with a theory that is supposed to show that moral attitudes are something other than they seem to be. There are important differences in explanation, but they do not call for a revision of the fundamental premise that a moral judgment can be true or false.

This observation also helps us to provide a promising alternative to a different way in which subjectivism has been spelled out. Crispin Wright argues that while there are no objective moral facts, we can hold on to the idea that moral judgments are frequently true or false. Once again, the motivation for this view is that there is something suspicious about moral features of reality. The alternative is to construe the content of 'true' as it is applied in moral contexts as something Wright calls superassertibility.²¹ The idea is that once we accept that truth in moral contexts does not involve reference to objective standards, but remains internal to the practice of asserting moral judgments, we no longer have to choose between expressivism and a more ambitious form of moral realism.

²⁰ Cf. S. Blackburn, *Ruling Passions*, ch. 3, *passim*.

²¹ Cf. C. Wright, *Truth and Objectivity*.

But this is not quite so easy. We can think of the situation in terms of a dilemma. We start from the thought that there are no objective moral facts. Now we can move in two different directions.

First, we might put the emphasis on the idea that there are no moral *facts*. This makes it comparatively easy to arrive at superassertibility, because truth must have become something quite different compared to what it used to be. The problem is that truth in moral contexts seems to be precisely what it used to be, namely something that is intimately bound up with the notion of a fact. If we are to revise that connection, this kind of revision needs some independent argument. Other things being equal, it seems clearly preferable to leave the unity of truth intact.

Second, we could put the emphasis on the idea that there are no *objective* moral facts. Now the position is more plausible. But it is no longer clear how we get from the claim that there are no objective moral facts to the analysis of truth in terms of superassertibility. After all, we can hold on to the idea that there are moral truths, and with them, moral facts, without having to think of these facts as objective. Moreover, this way of describing the position already relies on a distinction between truth and objectivity, and if this enables us to capture the distinction between moral realism and subjectivism, as it surely seems to do, there is no need to go further and manipulate the content of 'true' as it features in moral discourse by making it relative to a specific point of view. At the very least, this reinterpretation needs some independent motivation. Once

again, it seems preferable to leave truth to one side and preserve the unity of the truth predicate across all contexts.

Both expressivism and subjectivism as Wright wishes to construe it ultimately make the same mistake, which is to think of moral truths or facts as something inherently problematic. Both respond to this apparent problem by reinterpreting the contents of our attitudes. While expressivists leave truth intact, they think that they must reinterpret moral attitudes. While Wright leaves moral attitudes intact, he thinks that he must reinterpret our attitude to truth. But this is a false alternative, and that is why it seems to me to be a genuine advantage of subjectivism as I have described it that it makes neither demand. Once we accept the subjectivist account of moral facts, there is no longer anything mysterious about them. Given our linguistic practices, and some further conditions that we have good reason to think are satisfied, moral facts arise. These facts are not objective, but that is no reason not to speak of moral truths or facts. That, I have argued, is the truth in moral realism.

5 Wittgenstein Again

By way of conclusion of the argument, I will now consider some of the few passages in which the later Wittgenstein discusses ethical language. What Wittgenstein says is illuminating in the context of the present problem.

As Rush Rhees reports, Wittgenstein discussed the subject matter of ethics with him on several occasions. In conversations in 1942, Rhees brought up the problem facing a man who has come to the conclusion that he must either leave his wife or abandon his work of cancer research. What Wittgenstein says about this problem is worth quoting in full:

Such a man's attitude will vary at different times. Suppose I am his friend, and I say to him, 'Look, you've taken this girl out of her home, and now, by God, you've got to stick to her.' This would be called taking up an ethical attitude. He may reply, 'But what of suffering humanity? how can I abandon my research?' In saying this he may be making it easy for himself: he wants to carry on that work anyway. (I may have reminded him that there are others who can carry it on if he gives up.) And he may be inclined to view the effect on his wife relatively easily: 'It probably won't be fatal for her. She'll get over it, probably marry again,' and so on. On the other hand it may not be this way. It may be that he has a deep love for her. And yet he may think that if he were to give up his work he would be no husband for her. That is his life, and if he gives that up he will drag her down. Here we may say that we have all the materials of a tragedy; and we could only say: 'Well, God help you.'²²

This kind of approach to ethics is already quite far removed, both in substance and in style, from the highly abstract and aporetic treatments of the topic that we find in the *Tractatus* and the 'Lecture on Ethics' of 1929. There is no longer any suggestion that "it is impossible for there to be propositions of ethics".²³ Wittgenstein has brought ethical words back from their metaphysical to their everyday uses. Ethics no longer involves us in virtue of its very nature in something 'higher', 'transcendental',

²² R. Rhees, 'Some Developments in Wittgenstein's View of Ethics', pp. 22-3.

²³ L. Wittgenstein, *Tractatus Logico-Philosophicus*, 6.42.

'supernatural', or 'mystical'.²⁴ It has fallen into place as a chapter in the natural history of man.

One striking fact about this passage is that Wittgenstein construes the situation and the possible responses to it in terms of the different attitudes the husband or a friend may take. Depending on his attitudes, the choice the husband faces can be a most difficult dilemma. But as Wittgenstein points out, there is at least a possibility that it is not. Once more, everything depends on the husband's attitudes:

If he has, say, the Christian ethics, then he may say it is absolutely clear: he has got to stick to her come what may. And then his problem is different. It is: how to make the best of this situation, what he should do to be a decent husband in these greatly altered circumstances, and so forth. The question 'Should I leave her or not?' is not a problem here.²⁵

In other words, not only the solution to the problem, but the answer to the question whether there is so much as a problem will depend on the commitments the husband already has.

At this point, a familiar realist impulse sets in. Surely what the man should do is not up to him, or entirely dependent to his attitudes? We want to say that one of the choices he faces must be the *right one*, and one of the many attitudes he may take towards the situation is the *right one*,

²⁴ Cf. L. Wittgenstein, *Tractatus Logico-Philosophicus*, 6.42, 6.421, 6.54, 6.522, and L. Wittgenstein, 'A Lecture on Ethics', p. 7: "Ethics, if it is anything, is supernatural and our words will only express facts; as a teacup will only hold a teacup full of water and if I were to pour a gallon over it".

²⁵ R. Rhees, 'Some Developments in Wittgenstein's View of Ethics', p. 23.

the one he should take. What Wittgenstein goes on to say about this kind of question is extremely interesting:

Someone might ask whether the treatment of such a question in Christian ethics is *right* or not. I want to say that this question does not make sense. The man who asks it might say: 'Suppose I view his problem with a different ethics - perhaps Nietzsche's' - and I say: "No, it is not clear that he must stick to her; on the contrary, ... and so forth." Surely one of the two answers must be the right one. It must be possible to decide which of them is right and which is wrong.

But we do not know what this decision would be like - how it would be determined, what sort of criteria would be used, and so on. Compare saying that it must be possible to decide which of two standards of accuracy is the right one. We do not even know what a person who asks this question is after.²⁶

One may reply that we are dealing with two different kinds of question.

On the one hand, there is the question whether one of the ethical outlooks is the *right* one. On the other hand, there is the question whether it is possible to *decide* which one, if any, is the right one, and if so, *how* such a decision could be rationally arrived at. Thus, there may *be* a right answer, even if we cannot *determine* it.

This is an important distinction, but what Wittgenstein is getting at is clear precisely for the reason that Wittgenstein does not distinguish clearly between these two different questions. Thus, the reason why we have no clear idea of what such a decision would be like, what sort of criteria would be used, and so on, might be that we have no clear idea of the methods of determining the truth of any given ethical outlook. But the reason may also be that there is no truth to be determined. That the latter

²⁶ *Ibid.*

is the correct interpretation is suggested by the fact that Wittgenstein rejects the very question whether there is such a thing as a 'right' ethical outlook, not only the claim that there is any method or criterion that we could use to find it. Thus, Wittgenstein rejects the very question to which moral realism gives the answer as senseless. That is certainly an overstatement, but it makes clear that Wittgenstein rejects all claims to moral objectivity.

This interpretation is confirmed by the comparison with standards of measurement. The metaphoric field of measuring is in fact one of the familiar tools employed by Wittgenstein to illustrate the contrast between the empirical and the conceptual. On the one hand, there is something to be measured, on the other hand there is something to do the measuring. Just as in logic or in mathematics, this encourages an interpretation of the ethical that likens it not to propositions that are used to make empirical statements but to those that are used to formulate a rule, or to set up a framework of description. This is decidedly not all there is to it according to the moral realist, and perhaps this is the point where it becomes clear that Wittgenstein should not be counted as a moral realist.

That moral objectivity is the target is further confirmed by conversations with Rush Rhees in 1945. Here, Wittgenstein says that the goal of ethical enquiry naturally appears to be to gain insight into the "true nature of goodness or of duty":

The Truth in Moral Realism

People have had the notion of an ethical theory – the idea of finding the true nature of goodness or of duty. Plato wanted to do this – to set ethical inquiry in the direction of finding the true nature of goodness – so as to achieve objectivity and avoid relativity.²⁷

Relativity, in turn, “is to be avoided at all costs, since it would destroy the *imperative* in morality”.²⁸

The truth is that ethical language, like mathematical language, is normative.²⁹ Rules may *look* like empirical statements – ‘one plays patience by oneself’ – but they have a very different function in the language. Their proper task is not to represent, but to articulate norms of representation.³⁰

If this is the picture, it can easily seem that in ethics, there is no room for representation, truth and knowledge after all. But this conclusion does not follow. Wittgenstein admits truth and related notions into ethical discourse, and he does this by explicitly appealing to a minimalist, or deflationist, conception of truth:

Someone may say, ‘There is still the difference between truth and falsity. Any ethical judgment in whatever system may be true or false.’ Remember that ‘*p* is true’ means simply ‘*p*.’ If I say: ‘Although I believe that so and so is good, I may be wrong’: this says no more than that what I assert may be denied.

Or suppose someone says, ‘One of the ethical systems must be the right one – or nearer to the right one.’ Well, suppose I say Christian ethics is the right one. Then I am making a judgment of value. It amounts to *adopting* Christian ethics. It is not like saying that one of these physical theories must be the right one. The way in which some reality corresponds – or conflicts – with a physical theory has no counterpart here.³¹

²⁷ *Ibid.*

²⁸ *Ibid.*

²⁹ Cf. L. Wittgenstein, *Remarks on the Foundations of Mathematics*, VII-61: “What I am saying comes to this, that mathematics is *normative*. But “norm” does not mean the same thing as “ideal””.

³⁰ Cf. R. Rhees, ‘Some Developments in Wittgenstein’s View of Ethics’, pp. 23-4.

³¹ *Ibid.*, p. 24.

This is a significant passage, and it raises very sharply the question how to construe the difference between ethics and physical theory with respect to their relations to the world.

One thing that is clear, and is frequently emphasized by Wittgenstein, is that we should not be misled by the linguistic surface, which makes ethical and physical statements look deceptively alike. Both take the form of declarative sentences with subject and predicate. But according to Wittgenstein, whole areas of language that look as if they are dedicated to describing how things are must be understood in other terms.³² As Blackburn puts it: “Even at a cursory glance, his later work is shot through with warnings against taking surface uniformity as a safe guide to linguistic function”.³³ In *Philosophical Investigations*, Wittgenstein writes that “we remain unconscious of the prodigious diversity of all the everyday language games because the clothing of our language makes everything alike”.³⁴ To cure ourselves from philosophical confusion, we need to

make a radical break with the idea that language always functions in one way, always serves the same purpose: to convey thoughts – which may be thoughts about houses, pains, good and evil, or anything else you please.³⁵

³² Cf. S. Blackburn, *Ruling Passions*, p. 77.

³³ S. Blackburn, *Truth*, p. 130.

³⁴ L. Wittgenstein, *Philosophical Investigations*, pt. II, sec. xi, p. 191; cf. S. Blackburn, *Ruling Passions*, p. 80, and *Truth: A Guide for the Perplexed*, p. 130.

³⁵ L. Wittgenstein, *Philosophical Investigations*, § 304.

All this suggests that Wittgenstein puts forward something like an expressivist interpretation of moral language. The most telling passage is perhaps the one where Wittgenstein rejects a realist interpretation of the claim that moral judgments may be true, which is coupled with the claim that to call a Christian ethics 'true' is to adopt it.

But note that there is room for a different interpretation. It is possible to reject the claim to objectivity in ethics, so as to be able to say that the way in which some reality corresponds - or conflicts - with a physical theory has no counterpart here. It does not follow that there are no moral truths or facts, and this is fully compatible with the view that to call a Christian ethics true is to adopt it. Indeed, now a substantial contrast would emerge between saying that a moral judgment may be true given a certain moral outlook and a custom or practice of using moral words, which is not objectionable, and saying that a moral outlook may itself be either true or false. Unlike the former, the latter is not a useful expression, unless it serves to reaffirm that moral outlook. This contrast is at least implicit in the passage, and on the subjectivist view as I understand it, can be preserved.

6 Ethics and Aesthetics

What exactly Wittgenstein would have said had he addressed the question will remain an open question, and the textual evidence is slim. The important point is that Wittgenstein clearly rejects all realist

interpretations, which is equally confirmed by lecture notes on philosophical aesthetics that were taken by some of his students. These are worth discussing because they help to undermine a different, but related, form of realism.

The lectures start with a dramatic statement:

The subject (Aesthetics) is very big and entirely misunderstood as far as I can see. The use of such a word as 'beautiful' is even more apt to be misunderstood if you look at the linguistic form of sentences in which it occurs than most other words. 'Beautiful' is an adjective, so you are inclined to say: "This has a certain quality, that of being beautiful".³⁶

There is of course a sense in which it is perfectly true that beautiful things *have* a certain quality, namely that of being beautiful, just as there is a sense in which it is true that those things that are good have a certain quality, namely that of being good. This is just a variation on 'It is true that these things are beautiful' or 'These things are good'. The important point is that this does not introduce an item or a quality *in the sense in which the realist construes it*. To think of beautiful objects in terms of a feature 'beauty' that objects either have or lack, and that exists alongside its natural qualities and is there to be right or wrong about as they are is the root of the mistake. If so, then the mistake of philosophical aesthetics runs exactly parallel to the mistake we find in ethics. According to Rush Rhees, who took some of the notes from which these lectures have been

³⁶ L. Wittgenstein, 'Lectures on Aesthetics', in: L. Wittgenstein, *Lectures and Conversations on Aesthetics, Psychology and Religious Belief*, p. 1. The lectures date from the summer of 1938.

reconstructed, we find Wittgenstein speaking in that very sentence of both 'beautiful' and 'good'.³⁷ There is every reason to believe that Wittgenstein would have included ethics in this context.

Whatever Wittgenstein actually said, the very nature of the problem certainly suggests that there is a close connection. This is particularly striking in the context of the ethical philosophy of G. E. Moore, which was well known to Wittgenstein. Thus, in his 'Lecture on Ethics' of 1929, Wittgenstein refers to *Principia Ethica* and its definition of ethics as "the general inquiry into what is good".³⁸

Moore famously thought that this enquiry will take the form of an enquiry which actions or states of affairs have a certain quality, namely that of being good. So it is tempting to adapt the passage from the lectures on aesthetics: 'Good is an adjective, so you are inclined to say: "This has a certain quality, that of being good"'. This, Wittgenstein seems to say, is a mistake, and that is what he would have said earlier, although he would have said it for a different reason. At the time of the *Tractatus*, Wittgenstein thought that when we try to say of something that it is good or bad or beautiful, we are not only likely to be making a mistake, but uttering nonsense. Still, a close connection between ethics and aesthetics is already present in the *Tractatus*:

³⁷ L. Wittgenstein, 'Lectures on Aesthetics', p. 1.

³⁸ L. Wittgenstein, 'A Lecture on Ethics', p. 4; cf. G. E. Moore, *Principia Ethica*, p. 3: "This, then, is our first question: What is good? and What is bad? and to the discussion of this question (or these questions) I give the name of Ethics, since that science must, at all events, include it".

The Truth in Moral Realism

It is clear that ethics cannot be put into words.

Ethics is transcendental.

(Ethics and aesthetics are one and the same.)³⁹

In the 'Lecture on Ethics', the association is no longer quite as close, but Wittgenstein still thinks that there is a connection:

Now I am going to use the term Ethics in a slightly wider sense, in a sense in fact which includes what I believe to be the most essential part of what is generally called Aesthetics.⁴⁰

The mistake as Wittgenstein describes it in his later years is to construe words such as 'beautiful' or 'good' as names of qualities. But if we are not to construe them as qualities, what is the alternative?

In the 'Lectures on Aesthetics', we find a significant methodical remark: "An intelligent way of dividing up a book on philosophy", Wittgenstein says, "would be into parts of speech, kinds of words". In such a book, there could be a chapter on 'seeing' and 'feeling' and other verbs describing 'personal experience':

We get a peculiar kind of confusion or confusions which comes up with all these words. You would have another chapter on numerals – here there would be another kind of confusion: a chapter on 'all', 'any', 'some', etc. – another kind of confusion: a chapter on 'you', 'I', etc. – another

³⁹ L. Wittgenstein, *Tractatus Logico-Philosophicus*, 6.421; cf. L. Wittgenstein, *Notebooks 1914-1916*, pp. 72-91, esp. p. 83: "The work of art is the object seen *sub specie aeternitatis*; and the good life is the world seen *sub specie aeternitatis*. This is the connexion between art and ethics. The usual way of looking at things sees objects as it were from the midst of them, the view *sub specie aeternitatis* from outside". For a discussion of this point, see J. C. Edwards, *Ethics Without Philosophy*, ch. 2.

⁴⁰ L. Wittgenstein, 'A Lecture on Ethics', p. 4.

kind: a chapter on 'beautiful', 'good' – another kind. We get into a new group of confusions; language plays us entirely new tricks.⁴¹

Once again, we find Wittgenstein treating the philosophical confusions that surround both ethics and aesthetics explicitly in tandem. The cure is, here as everywhere, to realize that language has a range of different uses, and to pay attention to the differences between the ways in which different words are used. Language is like a tool chest. Different tools have different uses.⁴²

The question is, therefore, what kind of tool 'beautiful' and 'good' really are. This calls for a grammatical investigation, which can be conducted by asking how these words are taught. Doing this destroys a variety of misconceptions, and gives us a primitive language in which the word is used:

Although this language is not what you talk when you are twenty, you get a rough approximation to what kind of language game is going to be played.⁴³

Here we find that words like 'beautiful' and 'good' have uses that are very different from the uses that the realist imagines. Wittgenstein considers 'beautiful' and 'good':

If you ask yourself how a child learns 'beautiful', 'fine', etc., you find it learns them roughly as interjections. ('Beautiful' is an odd word to talk

⁴¹ L. Wittgenstein, 'Lectures on Aesthetics', p. 1.

⁴² *Ibid.*; cf. L. Wittgenstein, *Philosophical Investigations*, §§ 11, 14.

⁴³ L. Wittgenstein, 'Lectures on Aesthetics', pp. 1 f.

about because it's hardly ever used.) A child generally applies a word like 'good' first to food. One thing that is immensely important in teaching is exaggerated gestures and facial expression. The word is taught as a substitute for a facial expression or a gesture. The gestures, tones of voice, etc., in this case are expressions of approval. What *makes* the word an interjection of approval? It is the game it appears in, not the form of words.⁴⁴

Still, we will ask if this is the correct analysis of the language game we play 'when we are twenty'. Could it not be that this quite basic language game becomes much more sophisticated than the expressivist analysis implies?

In one sense, the answer to that question must be 'yes'. There is a point at which I could no longer replace the words with an exaggerated gesture or facial expression. This is only to be expected, given that my attitudes become much more complex as my abilities, including my abilities of using language, continue to develop. There is no *natural* expression, now, for my wish that we will have good weather for our honeymoon in Italy this spring. More importantly, there is a point at which the language game becomes complex enough to make it artificial if not inappropriate to say that we are always dealing with 'expressions of approval'.

This is important insofar as the expressivist analysis itself can easily fall into those abstractions or idealizations that obscure rather than illuminate the use of words. Wittgenstein keeps emphasizing differences:

⁴⁴ *Ibid.*, p. 2.

Words such as 'lovely' are first used as interjections. Later they are used on very few occasions. We might say of a piece of music that it is lovely, by this not praising it but giving it a character. (A lot of people, of course, who can't express themselves properly use the word very frequently. As they use it, it is used as an interjection.)⁴⁵

Similarly:

If I say of a piece of Schubert's that it is melancholy, that is like giving it a face (I don't express approval or disapproval).⁴⁶

A similar point emerges in a discussion with Rhees about 'deterioration':

You use 'deterioration' on the one hand to describe a particular kind of development, on the other hand to express disapproval. I may join it up with the things I like; you with the things you dislike. But the word may be used without the affective element; you use it to describe a particular kind of thing that happened. It was more like using a technical term – possibly, though not at all necessarily, with a derogatory element in it.⁴⁷

Again:

"What similarity has my admiring this person with my eating vanilla ice cream and liking it?" To compare them seems almost disgusting. (But you can connect them by intermediate cases.)⁴⁸

What this grammatical investigation can help us see is that there is no one thing, not even expressing approval or disapproval, that forms the common core of all aesthetic judgments. The use of language is much

⁴⁵ *Ibid.*, p. 3.

⁴⁶ *Ibid.*, p. 4.

⁴⁷ *Ibid.*, pp. 10 f.

⁴⁸ *Ibid.*, p. 12.

more complex, and it changes and develops. If we want to understand it, we must look at the complicated setting of activities in which words of adults, as opposed to words of children, have their use: "What belongs to a language game is a whole culture".⁴⁹ The same is true in ethics.

Now none of this discourages a realist, who may waive all aspirations to a uniform style of analysis. But there is still the point that what we are inclined to say about words such as 'good' or 'beautiful' is philosophically confused, and we have reason to expect that the more basic language game exposes that confusion. If Wittgenstein is right in stressing the pragmatic, the expressive, the affective side of ethics and aesthetics, as he clearly seems to do, then we must conclude that moral or aesthetic realism does not provide the adequate interpretation of our attitudes even when considering the language game we play 'when we are twenty'. Realism is the illness, not the cure:

Would it matter if instead of saying "This is lovely", I just said "Ah!" and smiled, or rubbed my stomach? As far as these primitive languages go, problems about what these words are about, what their real subject is, don't come up at all.⁵⁰

More clearly still:

You could regard the rules laid down for the measurement of a coat as an expression of what certain people want. People separated on the point of what a coat should measure: there were some who didn't care if it was broad or narrow etc.; there were others who cared an enormous lot. The rules of harmony, you can say, expressed the way people wanted chords

⁴⁹ *Ibid.*, p. 8.

⁵⁰ *Ibid.*, p. 3.

to follow – their wishes crystallized in these rules (the word ‘wishes’ is much too vague.)⁵¹

This is an interesting observation because it helps us to avoid two different kinds of mistake.

First, there is no suggestion that whenever we say that a coat should be cut in a certain way, this is merely an expression of a personal preference. There is a standard that is independent of a given preference, here and now, which is clear from the fact that one can dislike such standards. Here, we find an interesting parallel with ‘good’. It makes perfect sense to say that hotels have to meet a certain set of standards to be good hotels, but that one simply does not like a good hotel. ‘I simply don’t like staying at good hotels’ is an intelligible thing to say.⁵²

Second, there is no suggestion that there is a standard of correctness for the ways in which coats should be cut that goes beyond the rules that were ‘laid down’. The rules themselves are said to answer, not to some realm of facts about the way coats should *really* be cut, but to our attitudes and expectations. Of course, it is not normally true that there is a clear division, so that first there were the wishes, all articulate and clear, and then the rules were made to fit them. The process is much more involved than that. Wishes change as rules develop. Even the talk of ‘wishes’ can at some point become quite misleading, which is why Wittgenstein enters a

⁵¹ *Ibid.*, pp. 5 f.

⁵² Cf. B. Williams, *Ethics and the Limits of Philosophy*, p. 125; cf. B. Williams, *Morality*, pp. 52-68.

second caveat: "And although we have talked of 'wishes' here, the fact is just that these rules were laid down".⁵³

This is not, I take it, all that realists would want to say about this kind of situation. Few people are realists about the standards for the measurements of coats, but the situation is the same in ethics. Wittgenstein, for one, refers explicitly to Moore in this connection:

If I had to say what is the main mistake made by philosophers of the present generation, including Moore, I would say that it is that when language is looked at, what is looked at is a form of words and not the use made of the form of words.⁵⁴

Here we have the prime example of the moral realist who looks at moral language through the spectacles of some misleading theory. Rightly realizing that 'good' cannot be identified with, for example, 'pleasurable' or 'useful', he concludes that it must stand for an intrinsic, irreducible and very special property of objects. And this is a mistake. Ultimately, it is the expression of the myth that every word stands for a certain *object*, or failing that, for a *quality* of such an object. This is the Augustinian picture, and it is deeply flawed.⁵⁵

If we do not want to be Moorean realists, we might think we have no choice but to become expressivists. But as we have seen, we have a choice between three ways of understanding ethics and aesthetics. On the one hand, we can say that this domain of discourse is not 'really'

⁵³ L. Wittgenstein, 'Lectures on Aesthetics', p. 6 n. 2.

⁵⁴ *Ibid.*, p. 2.

⁵⁵ Cf. L. Wittgenstein, *Philosophical Investigations*, §§ 1 ff.

representational. This is what expressivists will say, and what Wittgenstein says can be called on to support that kind of answer.

Alternatively, we can say that language in ethics and aesthetics is representational, and that what is represented are features like goodness or beauty that are features of objects or actions that exist independently of us. This is the view that Wittgenstein emphatically rejects. There can be no such thing as a science of aesthetics, as the realist construes that term:

You might think Aesthetics is a science telling us what's beautiful – almost too ridiculous for words.⁵⁶

Again, the same is true in ethics. At this point, it is important to remember that there is a third way of interpreting this puzzling area of language. According to the interpretation I suggested, language in ethics or aesthetics is frequently representational. But it does not represent features of objects that exist independently of us. This is the subjectivist version of the secondary quality analogy, and it seems to me to be the most promising account.

Does it give the moral realist everything he wants? Surely not – but what the realist wants is difficult to get, and the thoughts and pictures that motivate moral realism are much too complex to allow any simple refutation. At bottom, I suspect that they need therapy, not argument.⁵⁷

⁵⁶ L. Wittgenstein, 'Lectures on Aesthetics', p. 11.

⁵⁷ Cf. L. Wittgenstein, *Philosophical Investigations*, §§ 133, 255: "There is not a philosophical method, though there are indeed methods, like different therapies"; "The philosopher treats a question; like an illness".

Nevertheless, we have seen that there is a set of arguments that strongly support this subjectivist position. What remains to be seen is how well the subjectivist hypothesis can make sense of other central aspects of ethical life and ethical experience.

3

Are There External Reasons?

One feature of our ethical experience that naturally suggests a realist interpretation is our ability to act for reasons. In this chapter and the one that follows, I will consider this ability in more detail. I believe that the correct account of our ability to act for reasons is given by a view known as internalism.

The argument proceeds as follows. I begin with a discussion of internalism as defined by Bernard Williams. After that, I consider his objection to externalism, and conclude that it is inconclusive as it stands. With this result established, I will try to motivate internalism in a new and in my view more convincing, way. It will emerge in the course of that discussion that internalism itself must undergo important changes.

1 The Structure of Internalism

In a famous paper called 'Internal and External Reasons', Williams contrasts two interpretations to which claims about reasons for action are susceptible.¹ On the first, the presence of a reason is conditional upon the presence of a motivation. On the second, there is no such condition for the presence of a reason. Call the first the 'internal', and the second the 'external' interpretation of statements about reasons for action.² The question raised by Williams is how to interpret reason statements where the external interpretation is called for, and whether, so interpreted, these sentences are ever true.

What is involved in the internal interpretation? By definition, internalism claims that reasons for actions are based on a set of motivations. Call it the subjective motivational set, or *S*. Let the variable 'A' stand for a given agent, and ' ϕ ' for a given course of action. The idea that drives internalism is that for it to be true that A has a reason to ϕ , there has to be what Williams calls a *sound deliberative route* from A's present motivations to A's ϕ -ing. Accordingly, his preferred formula for the internalist position is this:

¹ B. Williams, 'Internal and External Reasons', in: B. Williams, *Moral Luck*, essay 8. Williams comes back to the distinction at various points in his publications. The most important of these are the following: 'Internal Reasons and the Obscurity of Blame', in: B. Williams, *Making Sense of Humanity*, essay 3; 'Replies', in: J. E. J. Altham and R. Harrison (eds.), *World, Mind, and Ethics*, ch. 10; and 'Postscript: Some Further Notes on Internal and External Reasons', in: E. Millgram (ed.), *Varieties of Practical Reasoning*, pp. 91-7.

² B. Williams, 'Internal and External Reasons', p. 101.

Are There External Reasons?

A has reason to ϕ only if there is a sound deliberative route from A's subjective motivational set S to A's ϕ -ing.³

This formula calls for further comment.

1. *Is the Condition Necessary or Sufficient?* As it stands, the formula states only a necessary condition for the truth of reason statements. This raises the question of whether the presence of a motivation in S is also a sufficient condition for the truth of a reason statement. This is a complicated question, and it is one that Williams himself tends to leave to one side. Williams does say at one point that he thinks that the presence of a motivation is also sufficient for the presence of a reason. But he does not say why he thinks this.⁴

I believe there is good reason to resist the claim that every motivation that can lead by a sound deliberative route to an appropriate action provides a reason for action. I will come back to that point. For the moment, I will focus on the claim that the presence of a motivation is a necessary condition for the truth of a reason statement.

2. *Mistakes in Reasoning and False Belief.* One important feature of internalism is that it allows, at the very least, for the correction of false reasoning and false belief. A may think he has a reason to ϕ , where that belief is based on mistakes in reasoning, or false beliefs about the facts. In such cases, it may be that the agent has no reason to ϕ , but merely thinks that he does. Sometimes we say that an agent *would* have a reason if he

³ B. Williams, 'Postscript', p. 91.

⁴ B. Williams, 'Internal Reasons and the Obscurity of Blame', pp. 35 f.

knew the facts, as opposed to saying that he *has* a reason but does not *know* it. But in general, reasons for action respond to the facts, and not to our beliefs about the facts.

3. *Normative Force.* It follows that what an agent thinks he has reason to do and what he has reason to do can come apart. It is a conceptual requirement that reasons for action respond to the facts, not to our beliefs about the facts. This is one important reason why 'A has reason to ϕ ' has 'normative force'.⁵ A conception of practical reason that leaves no room for this normative dimension is certainly too narrow. It reduces normative reasons to explanatory ones, and this would eliminate the normative dimension that comes with the notion of a reason.

4. *Rationality and Moral Reasons.* According to Williams, the reason why the correction of false reasoning and false belief can be incorporated into the internalist account is that the correction of such errors is already part of the *S* of every rational agent. By contrast, there is no reason to think that moral reasons are already part of every rational agent's *S*, even where agents are fully rational.

Kantians believe that they are, but that needs an argument, and it is by no means obvious what that argument could be. The generalization test would be a criterion of moral right and wrong that works with basic notions of consistency, but it is well known that this test faces serious

⁵ Cf. B. Williams, 'Internal Reasons and the Obscurity of Blame', p. 36; 'Postscript', pp. 92 f.; 'Internal and External Reasons', pp. 102 f.

problems.⁶ If it is said that there are basic normative facts, such as the fact that rational beings exist as ends in themselves, this begs the question against the internalist. Moreover, it seems that what is needed is a normative fact that can support claims to ethical objectivity, and there is no reason to think that there are any such facts.

5. *The Indeterminacy of S.* Moral motivations can, but need not be present in a rational agent. This observation suggests that the contents of *S* can vary, and that its scope is, to a certain extent, indeterminate. On any adequate account, *S* must contain such things as “patterns of evaluation, patterns of emotional reaction, personal loyalties, and various projects, as they may be abstractly called”.⁷ There is no suggestion that *S* could not include altruistic or moral patterns of motivation.

6. *Reasons and Causes.* Davidson argues in ‘Actions, Reasons, and Causes’ that reasons for action are based on

desires, wantings, urges, promptings, and a great variety of moral and aesthetic principles, economic prejudices, social conventions, and public and private goals and values insofar as these can be interpreted as attitudes of an agent directed toward actions of a certain kind.⁸

There is a close parallel with internalism insofar as the list of types of motivation is quite open. But there is also a difference, because Davidson

⁶ Cf. B. Williams, ‘Internal Reasons and the Obscurity of Blame’, p. 37 and ‘Postscript’, p. 92. Williams refers to a suggestion by C. Korsgaard to the effect that the internalist requirement could be met by a Kantian conception of morality; cf. C. Korsgaard, ‘Skepticism about Practical Reason’, in: C. Korsgaard, *Creating the Kingdom of Ends*, essay 11.

⁷ B. Williams, ‘Internal and External Reasons’, p. 105.

⁸ D. Davidson, ‘Actions, Reasons, and Causes’, in: D. Davidson, *Essays on Actions and Events*, pp. 3 f.

is primarily interested in action explanation. Williams offers an account of normative reasons. Action explanation features in his argument in favour of internalism, but it enters at a different level. There is no suggestion that reasons are causes, and where normative reasons are concerned, it is not even clear what it could mean to say that reasons are causes. This kind of claim seems to involve a category mistake.

7. *Varieties of Practical Deliberation.* Internalism does not construe all instances of practical reasoning in terms of given attitudes combining with beliefs about the means to given ends. That would be a gross oversimplification, and there is no reason to think that the domain of rationality is exhausted by adjusting the relations between means and ends.

For one thing, practical reasoning can lead to the conclusion that one has reason to ϕ in a large variety of ways, and there is no suggestion that what constitutes a sound deliberative route is fully determinate.⁹ For another, there is deliberation about ends themselves.¹⁰ Agents may deliberate about how best to combine various elements in S , either by time ordering, or with a view to what goal or project they think most important. There is also the phenomenon of creativity in practical deliberation.

⁹ B. Williams, 'Internal and External Reasons', p. 104; cf. 'Internal Reasons and the Obscurity of Blame', p. 38.

¹⁰ Williams refers to A. Kolnai, 'Deliberation Is of Ends', in: A. Kolnai, *Ethics, Value, and Reality*, essay 3.

8. *Effects of Rational Deliberation.* It is obvious that the process of deliberation can add to, or subtract from, or modify the elements of *S*. Imagination can show what it would really be like if some course of action were to be pursued, and as a result, prevent the agent from pursuing it. Conversely, deliberation can reveal that the agent has reason to do something he did not think he had a reason to do.¹¹ So we should not think of *S* as static, like a list of preferences to work off, or desires waiting *en bloc* to be satisfied. 'A has reason to ϕ ' goes significantly beyond both 'A thinks he has reason to ϕ ' and 'A is presently motivated to ϕ '.¹²

Equipped with this sketch of the structure of internalism, we can now ask what externalists find missing. One thing that is immediately clear is that the existence of moral or altruistic reasons cannot by itself provide a reason to reject internalism. There are no special constraints on the contents of *S* on the internalist conception, so there is no reason why *S* could not include altruistic and moral reasons.

Moreover, it is clear that internalists need not consider altruistic and moral reasons as 'one kind of reason among others'. They are, of course, one kind of reason among others in the trivial sense that there are others, as even the most austere externalist is likely to admit. What is meant must be that on the internalist conception, moral or altruistic reasons are no more important or authoritative or compelling than some other type of reason. But in order to show that, we would have to show

¹¹ B. Williams, 'Internal and External Reasons', pp. 104 f.

¹² Cf. B. Williams, 'Internal Reasons and the Obscurity of Blame', p. 36.

that the elements in S that give rise to them are no more important or authoritative or compelling *for the agent* compared to other elements in S , and there is no reason to accept that.

This suggests that the internalist conception can account not only for the situation where an agent *has* reason to ϕ , but equally for situations where he has *most* reason to ϕ , or where he has *more*, or *less*, or *equally* good reason to ϕ than to ψ . What an agent has more or most reason to do may not always be fully determinate, just as it may not always be determinate whether he *has* a reason to ϕ . But to the extent that it is determinate, we should expect the hierarchy or ordering of reasons to be given by the agent's S if the internalist account of reasons is correct.

However, there is still a problem. As we have seen, there is no suggestion that moral or altruistic reasons *cannot* feature in our S . But there is also no suggestion that they *have to* feature in our S , or even that they always *do*. This shows that internalism makes a moral difference. Consider a sexist man who insists that he has no reason to be nicer to his wife:

I say, 'You have reason to be nicer to her'. He says, 'What reason?' I say, 'Because she is your wife'. He says - and he is a very hard case - 'I don't care. Don't you understand? I really do not care'. I try various things on him, and try to involve him in this business; and I find that he really is a hard case: there is *nothing* in his motivational set that gives him a reason to be nicer to his wife as things are.¹³

What are we to say to, or about, this man?

¹³ *Ibid.*, p. 39.

Are There External Reasons?

As Williams points out, we can say very many things to or about this man, including that he is ungrateful, inconsiderate, hard, sexist, nasty, selfish, brutal, and many other disadvantageous things.¹⁴ Of course we will say, whatever else we say, that it would be better if he were nicer to his wife. But there is one specific thing the externalist wants us to say, and that we may refuse to say, namely that the man has reason to be nicer to his wife.

We *may* refuse to say this – but we do not have to. When we say that the man in the example has reason to be nicer to his wife, we may be using this expression as a further affirmation of our moral stance towards him.¹⁵ But externalists will think there is a different sense attached to that expression, and internalists insist that insofar as there is a different sense attached to this expression, it is bound to be false.

This conclusion is in line with a diagnostic remark from ‘Internal and External Reasons’ where Williams says that external reason statements, when definitely isolated as such, “are false, or incoherent, or really something else misleadingly expressed”.¹⁶ Sometimes such a statement will amount to little more than saying that things would be better if the agent acted differently. This, however, will all too obviously be *our* assessment of the situation, but not *his*. We are led to the conclusion that those who use this kind of expression frequently issue optimistic internal reason statements: “We launch them”, Williams writes, “and hope

¹⁴ *Ibid.*

¹⁵ S. Blackburn, *Ruling Passions*, pp. 264-6; cf. A. Gibbard, *Wise Choices, Apt Feelings*, p. 163.

¹⁶ B. Williams, ‘Internal and External Reasons’, p. 111.

that somewhere in the agent is some motivation that by some deliberative route might issue in the action we seek".¹⁷ If no appropriate motivation is forthcoming, we must withdraw our claim that *A* has reason to ϕ . If we insist that there is such a reason, and insist that saying this in some way goes beyond expressing our assessment of the situation, we are *bluffing*.¹⁸

A similar conclusion is suggested in the case where the man thinks that he has reason to be nicer to his wife, but not enough of a reason to make him change his ways. If the man prefers the life he now has to the one he could have if he were nicer to his wife, we may well face the conclusion that he has more, or most, reason to carry on the sexist life he is presently living. Everything will depend on the structure and profile of his motivations, and no external authority could provide a higher court of law to which we could appeal.

Why should we believe this?

2 A Humean Challenge

The question is whether what the external theory asks for can really be had. 'Internal and External Reasons' contains a complex argument that aims to establish that the external theory is false. We must now ask what that argument is.

The argument of 'Internal and External Reasons' is frequently discussed, but rarely understood. It begins with an observation about the

¹⁷ B. Williams, 'Internal Reasons and the Obscurity of Blame', p. 40.

¹⁸ B. Williams, 'Internal and External Reasons', p. 111; cf. 'Postscript', pp. 95 f.

Are There External Reasons?

conceptual connection between having a reason for action and coming to be motivated by it. This is a source of a significant misunderstanding.

It seems clear that, as a rule, something can only be a reason to do something if that reason can contribute to an explanation of an action. Agents must be capable of acting on the reasons that a theory of reasons says they have. This may seem to be easy to explain for the internalist. After all, on his view, reasons are based on the motivations the agent already has. So it appears that there is no deeper difficulty involved in the claim that reasons must generally be capable of motivating.

But this is not quite correct. After all, the belief that one has a reason for action goes significantly beyond the presence of a given motivation. If so, the internalist must invoke a belief that one has such a reason, and that belief must be what motivates the action. It may be asked whether there may be reasons for an action that cannot be a source of motivation. I am inclined to doubt that this is so, but we can leave this question to one side. Given that the internalist must account for motivation in terms of belief, he will be in the same position with respect to action explanation as the externalist. Of course, on the externalist account, there will be no requirement that my beliefs about reasons for action are based on my given motivations. But there will be the same connection between my belief that I have a reason for action *and action* as there is on the internalist account. If that connection is loose, and I may have a reason for action even if I cannot act on it or even be motivated to

act on it, then the externalist will face the same problem, so we can safely set this point aside.

Now the reason itself, or as we may put it less misleadingly, the truth of a reason statement, cannot contribute to the explanation of anything. It is a truth, or fact, and as such it is neither here nor there. The only item that can possibly be called upon in explanations of an action must be the belief, true or false, that there is such a reason. If there are external reasons, then it must, at least as a rule, be possible for the agent whose reason it is to come to believe that he has such a reason. Moreover, his believing that he has that reason must provide him with a motive.

There is a question how this can be so, forcefully articulated by Hume. Hume believes that reason is incapable of giving rise to motivation. "Morals", Hume observes, "excite passions, and produce or prevent actions". But

reason of itself is utterly impotent in this particular. The rules of morality, therefore, are not conclusions of our reason.¹⁹

Given Hume's account of reason, we can paraphrase his view by saying that belief can never motivate an action.

The key premise in the argument is that reason, by itself, cannot 'excite passions, and produce or prevent actions'. This, in turn, is the conclusion of an argument Hume puts forward earlier in the *Treatise*.

¹⁹ D. Hume, *A Treatise of Human Nature*, bk. III, pt. I, sec. I, p. 457.

Are There External Reasons?

There, he argues “*first*, that reason alone can never be a motive to any action of the will; and *secondly*, that it can never oppose passion in the direction of the will”.²⁰ Hume triumphantly concludes:

Reason is, and ought only to be, the slave of the passions, and can never pretend to any other office than to serve and obey them.²¹

But as it stands, his argument is at least inconclusive, and a large number of critics have cast doubt on the claim that Hume seeks to establish.²² In my view, there is no reason to believe that reason, or belief, can never give rise to a motivation. Hume himself does not deny that reason or belief can form part of action explanation. Thus, Hume thinks that reason can tell us, among other things, what means are required to achieve a given end. But if that is possible, what reason is there to think that reason cannot motivate? Even within the Humean framework, it is difficult to argue that reason is, and must be, motivationally inert.

It is not helpful to point out, as Hume does, that there are some rational conclusions that do not provide a motivation. It is true that my belief that $2 + 2 = 4$ does not, as it stands, provide me with a motive. But that only shows that rational processes of *certain kinds* do not provide a motivation *in certain contexts*. That falls short of the general claim that has

²⁰ *Ibid.*, bk. II, pt. III, sec. III, p. 413.

²¹ *Ibid.*, p. 415.

²² For some critical queries, see e.g. B. Stroud, *Hume*, chs. 7 and 8; T. Nagel, *The Possibility of Altruism*, chs. 2, 3 and 5; M. Platts, *Ways of Meaning*, ch. 10; see also J. L. Mackie, *Hume's Moral Theory*, chs. 3 and 4.

to be established. Reason is not the mere slave of the passions. Hume, for one, has given us no reason to think otherwise.

That reason, or belief, can motivate is what the externalist contends, although there are some notable exceptions.²³ It is important that it is also what the internalist contends, and that neither internalism, nor the argument in favour of internalism that Williams puts forward in 'Internal and External Reasons' rely on a Humean theory of motivation. Williams explicitly grants that believing that a particular consideration is a reason to act in a certain way "provides, or indeed constitutes, a motivation to act".²⁴ Moreover, he takes this to be quite obvious: "Of course reason, that is to say, rational processes, can give rise to new motivations".²⁵ The internalist account of practical deliberation relies on that being so. Internalism as Williams describes it not only does not presuppose the Humean theory, but is in fact incompatible with it.

Smith has argued for the claim that "desires are indeed the only states that can constitute our motivating reasons".²⁶ He thinks that some form of the Humean theory of motivation is correct. If it is, then the argument in favour of internalism, and internalism itself will come under pressure. So it is important to ask what reason we have to accept the Humean theory.

²³ Cf. M. Smith, *The Moral Problem*.

²⁴ B. Williams, 'Internal and External Reasons', p. 107.

²⁵ *Ibid.*, p. 108.

²⁶ M. Smith, *The Moral Problem*, p. 125.

Are There External Reasons?

A basic question that arises with respect to the Humean theory is what kind of action explanation it suggests. It is one thing to say that motivation always involves an element of desire. It is quite another to say that desire is the source, or the place where action explanation ultimately comes to rest. Smith remains neutral on the matter. In effect, Smith appears to restrict himself to the claim that motivation always *involves* desire.

So far, this is quite compatible with the view that Thomas Nagel endorses in *The Possibility of Altruism*, which contains important criticisms of the Humean theory. According to Nagel, it may well be true that all motivation involves desire. But Nagel insists that the relevant desire may, in turn, be subject to rational explanation.²⁷ In particular, it may be subject to explanation in terms of belief. But if it is subject to explanation in terms of belief, it is susceptible to explanation in terms of beliefs about what would be good or bad or right or wrong, or what there is reason to do. But if this is the situation, the Humean theory is not so much implausible as quite uninteresting, because it does not show that belief is not involved at a deeper level of action explanation. Moreover, it does not show that belief is not the most important, guiding or directing factor in the aetiology of action. On the present proposal, belief can explain the presence of desire, but the relevant beliefs need not, in turn, be subject to a further explanation in terms of desire. So we are no further on.

²⁷ Cf. T. Nagel, *The Possibility of Altruism*, ch. 5.

Are There External Reasons?

To come up with a more interesting opponent, some have described the Humean theory as claiming not merely that all action explanation must *involve* the presence of desire, but that explanations of desire will have to be derived from the presence of some other desire or desires. Thus, Jay Wallace has described the Humean theory in terms of what he calls a 'desire-out, desire-in' principle. The idea is that processes of thought that give rise to a desire as 'output' can always be traced back to a further desire as 'input'. According to Wallace, the Humean theory claims that desire, not belief "fixes the basic evaluative principles from which the rational explanation of motivation begins".²⁸

As we shall see, there is a truth in this, but it is not one that supports a Humean theory of motivation. The problem with the proposal as a version of the Humean theory is that while it is interesting and certainly says something substantial about action explanation, it is not clear why we should accept it. Even if we are inclined towards a Humean theory of motivation, we would not have to accept the Humean theory as a consequence of accepting some generic principle to the effect that only desires *could* explain the presence of desires of the kind that feature in a theory of action explanation. Lacking such a principle, the claim that desires are always to be explained by more basic desires as input, but not by beliefs, would have to be a piece of *a priori* psychology, or else some rash empirical generalization. Either way, it does not look very plausible. If this is the Humean theory of motivation, it does not look very

²⁸ R. J. Wallace, 'How to Argue About Practical Reason', p. 370.

promising. Smith seems to think that a theory that traces motivation back to some desire has the advantage of simplicity. But it is quite unclear why that should be so. Surely desires and beliefs must contribute to the explanation of an action, and if that is so, the simpler theory seems to be the one that leaves it open whether a desire or belief provides the basic motivation, as opposed to the one that requires a desire in addition to belief.

In fact, there is a deeper problem. The distinction between states that we refer to as desire and states that we refer to as beliefs is less clear and stable than it may initially appear to be. The distinction will presumably involve a difference in 'direction of fit'.²⁹ The idea is that a belief aims at truth, and that for it to be true, it has to 'fit the world', and not *vice versa*. Desires, by contrast, aim to change the world so that the world 'fits' the desire. But who is to say that states that aim to change the world to fit them cannot also aim at truth? Once again, Smith seems to think that any theory of motivation that relies on belief, even where the relevant kind of belief shares some features with a desire, will be inferior to a theory that involves a combination of desire and belief. But it is unclear why he thinks this.

I suspect the basic problem with his Humean theory of motivation is that Smith thinks of beliefs and desires as 'distinct existences'.³⁰ Smith is careful to point out that we should not think of desires on the model of

²⁹ Cf. M. Smith, *The Moral Problem*, ch. 4; the distinction can be traced back via Platts to G. E. M. Anscombe, *Intention*.

³⁰ *Ibid.*, p. 119.

empiricism, as 'perceptions in the mind', or states with a distinct phenomenological content. Nevertheless, it seems to me that there is something of an empiricist remnant even in his refurbished version of the Humean theory. Consider the idea that to have a desire is to be in a 'mental state', which goes with the idea that mental terms refer to mental items.

The problem is not, or not in the first instance, that Smith refers to desires as 'dispositions'. This term is useful as a technical expression in philosophical psychology. But given the expression, the temptation is now even stronger to think of dispositions as 'states' of the agent that have a characteristic explanatory role. An explanation of an action in terms of beliefs and motivations, and of motivations in terms of desires or beliefs, now looks like an interesting, substantial and informative psychological theory. As such, it seems to reveal something deep and interesting about the workings of the mind. The obvious suspicion is that what presents itself as a theory of motivation does not in fact do any such thing, but rather reflects the grammar of a certain range of our psychological vocabulary.

Belief involves a willingness to rely on what is believed in reasoning and action. If we ascribe a belief, we expect a willingness to express it in assertion, absence of doubt, contrary belief and hesitation, and so on. 'Desire' has a different, but equally complicated grammar. But if we take seriously the thought that the ascription of belief and desire is governed by a complicated and not necessarily well circumscribed set of

Are There External Reasons?

criteria, then it may turn out that being disposed to ϕ , or desiring to ϕ , turn out to be among the criteria for the ascription of the belief that there is a reason to ϕ , or that it would be good to ϕ , or something like it.

Such a connection would not be surprising, precisely because the criteria for the ascription of belief often include characteristic dispositions or desires. If you believe that the sun is more than 90 million miles away, this belief has no obvious relevance for any of your actions. But if you believe it, you will be disposed to act in certain ways. You will be disposed to use the proposition as a premise in reasoning; you will be disposed to answer 'More than 90 million miles' when someone asks you how far away you think the sun is; you will look surprised when you hear someone say that the sun is less than a million miles away; and so on and so on.

Now if these are some of the criteria for holding this belief, why should we not think that among the criteria for the ascription of beliefs about reasons for action, or about what would be worth doing, or about what would be a good or a bad thing to do, we find characteristic dispositions on the part of the agent to pursue or to avoid, to recommend or to discourage, to act or to keep oneself from acting, in characteristic ways associated with that action? Of course the link with action should be loose enough to make room for the phenomenon of weakness of will, where some of the characteristic links with motivation and action are severed. But this is quite possible on the present theory. Just as it is

possible to ascribe to you the belief (a) that the sun is more than 90 million miles away, (b) that light travels at roughly 300.000.000 meters per second, and (c) that the light that travels from the sun takes a few seconds to arrive on earth, it is possible to ascribe to you the belief that you should now make an appointment with the dentist, while you lack all corresponding motivation to. Light that travels from the sun takes more than 8 minutes to arrive on earth, as an easy calculation shows, so we are presented with a case of theoretical irrationality. All we have to say about your lack of motivation to make an appointment with the dentist is what we want to say anyway, namely that it is a case of practical irrationality.

If this account is correct, it looks as if the Humean theory of motivation does not give us deeper explanations of the psychological phenomenon of motivation, but rather expresses the criteria that govern the ascription of beliefs about what we have reason to do, should do, and so on. This is not to say that instead of appealing to a combination of desires and beliefs, our motivations are to be explained in terms of a single state with characteristics of both belief and desire, such as a 'besire'.³¹ The point is precisely that they should not be explained in terms of any mental 'states' at all.

³¹ Cf. *ibid.*, ch. 4, esp. pp. 116-25; for an earlier statement, see M. Smith, 'The Humean Theory of Motivation', esp. pp. 55-61.

3 Requirements of Rational Deliberation

According to Williams, the problem with externalism is not that beliefs about reasons for action issue in motivations. The problem is that they must do so in a specific way. The motivation to ϕ and the belief that I have reason to ϕ could be the product of prejudice, rhetorical influence, or pressure. This, however, would not give the externalist what he wants, which is that the belief be acquired as the result of rational reflection. More specifically, Williams thinks

that the agent should acquire the motivation *because* he comes to believe the reason statement, and that he should do the latter, moreover, because, in some way, he is considering the matter aright.³²

This seems a convincing account of the situation, and is fully in line with what externalists will want to say. The agent must acquire the belief and corresponding motivation because he deliberated *correctly*.

At the same time, it is a requirement of the external theory that the belief and corresponding motivation must not always, or even normally, be the outcome of a rational extension of already existing motivations. Once again, this point is obvious, for in that case, all reasons for action would be internal after all. So the idea must be that an agent who deliberates correctly can come to believe the true reason statement and be motivated by it independently of the beliefs and motivations he originally had.

³² B. Williams, 'Internal and External Reasons', p. 109.

Are There External Reasons?

Williams concludes that whatever his motivations are before he acquires the new belief and corresponding motivation, it must be possible for the agent to reach, by rational reflection, the conclusion that he has a reason to ϕ . This conclusion then gives rise to new beliefs and motivations, and may feature in the explanation of an action. And this suggests an argument to the effect that external reasons are impossible. On the one hand, beliefs and motivations must be rationally arrived at. At the same time, they do not have rational support deriving from already existing motivations. And

if this is correct, there does indeed seem great force in Hume's basic point, and it is very plausible to suppose that all external reason statements are false.³³

Ex hypothesi, there is no motivation for the agent to deliberate *from* to reach his new motivation. But the new motivation must be *rationally* arrived at, granted the earlier motivations. At the same time, it must not bear to those earlier motivations the kind of rational relation that would yield internal reasons. Williams concludes: "I see no reason to suppose that these conditions could possibly be met".³⁴

But has that really been shown by the argument? Moreover, is what Williams thinks impossible something the externalist has to assert?

Let us retrace the steps of the argument:

³³ *Ibid.*

³⁴ *Ibid.*

Are There External Reasons?

- (1) If there are any external reasons, it must, in general, be possible to act on them.
- (2) In order to be able to act on an external reason, it must be possible for the agent to form the belief that he has such a reason and be motivated by it.
- (3) Moreover, the agent must form the belief and corresponding motivation because he comes, in some sense, to consider the matter aright.
- (4) So the belief and corresponding motivation must be arrived at by correct deliberation.
- (5) At the same time, the new belief and motivation must not be arrived at as the result of rational deliberation on the basis of the given motivations.
- (6) But they must be arrived at by rational deliberation.
- (7) There is no reason to think that they could be arrived at in such a way.

- (8) Therefore, it is not possible to acquire the belief and motivation in the way required by the external theory.

Is this a sound argument?

4 The Unsupported Premise

Given its conclusion, it is unsurprising that the argument received the critical reception that it did. I will not attempt anything like a review of the literature.³⁵ Instead, I will focus on what I think is the most telling objection to the argument Williams sets out in 'Internal and External Reasons'.

Premises (1)-(4) seem unproblematic. Let us grant, for the sake of argument, that they are true. Let us grant, moreover, that both the internalist and the externalist accept that they are true. Premise (5) expresses a condition that only the externalist theory has to fulfil, so here internalists and externalists part company. But given the definition of externalism, premise (5) is not inherently problematic.

The trouble starts with premise (6). For all that has been shown so far, the agent who has an external reason must be able to arrive at the conclusion that he has that reason because he comes, in some sense, to see

³⁵ Some helpful critical discussions are the following: R. Cohon, 'Are External Reasons Impossible?'; B. Hooker, 'Williams' Argument Against External Reasons'; C. Korsgaard, 'Skepticism about Practical Reason', reprinted in: C. Korsgaard, *Creating the Kingdom of Ends*, essay 11; J. McDowell, 'Might There Be External Reasons?', reprinted in: J. McDowell, *Mind, Value, and Reality*, essay 5; E. Millgram, 'Williams' Argument Against External Reasons'; T. M. Scanlon, 'Appendix: Williams on Internal and External Reasons', in: *What We Owe To Each Other*, pp. 363-73; M. P. Jenkins, *Bernard Williams*, ch. 5.

Are There External Reasons?

the matter aright. That is what the externalist theory has to say, for obvious reasons. If it did not, it would run a risk of losing the idea that there is an external truth to be appreciated, in which case we would be driven back to the internalist account. Alternatively, the external theory could say that there is such a truth, but that the agent does not come to appreciate it in deliberation, in which case the idea of acting for a reason would be lost.

So we can say, given (4), that the belief and corresponding motivation must be arrived at by correct deliberation, and given (5), that it must be possible for the agent to arrive at the new belief and corresponding motivation not by rational reflection on the basis of his given motivations, which merely spells out what the externalist seeks to establish. But why should we assume that (6), the new belief and corresponding motivation must be arrived at by *rational* reflection?

This claim is correct if we read '*rational* reflection' as a mere variant of '*correct* deliberation'. But this reading would need further explanation, and in Williams, it does not get such an explanation. More importantly, this is not the way in which the premise functions. If we interpret the requirement of *rational* deliberation as the requirement of *correct* deliberation, premise (6) comes out true, but the argument is no longer conclusive.

The requirement that the new belief and corresponding motivation must be arrived at by rational deliberation clearly goes beyond the requirement that the belief must be true. If the new belief and motivation

are arrived at by rational deliberation, this must mean that the belief must be arrived at by a process analogous to that involved in the process of rational deliberation on the basis of the given motivations, which Williams describes.

If we read the requirement in this way, premise (6) adds something important, and given the premise, the argument is valid. The problem is that, so understood, we have not been given any reason to accept premise (6). The external reasons theorist needs the claim that agents can arrive at new beliefs and motivations by *correct* deliberation. Why should he also accept that agents must arrive at those beliefs by *rational* deliberation? Perhaps what is needed is something like conversion.³⁶ As McDowell points out, all that the external reasons theorist needs at that point in the argument is that *in* coming to believe the reason statement, the agent is coming to see the matter aright. This leaves it quite open how the transition is effected.³⁷

There is an instructive parallel here with religious belief, where the notion of conversion is at home. Wittgenstein remarks:

“Convincing someone of God’s existence” is something you might do by means of a certain upbringing, shaping his life in such & such a way.³⁸

³⁶ This is brought out clearly by McDowell; cf. J. McDowell, ‘Might There Be External Reasons?’, in: J. McDowell, *Mind, Value, and Reality*, essay 5, esp. p. 102.

³⁷ *Ibid.*, p. 100.

³⁸ L. Wittgenstein, *Culture and Value*, p. 97.

Are There External Reasons?

But note that so far, this does not say anything about the truth of that belief. Conversion can give rise to true belief, to false belief, and to belief to which nothing corresponds to even make it true or false.

On the other hand, there is the possibility that conversion could not be replaced by giving arguments, because the belief to which a person is converted is so fundamental that it lies beyond the reach of arguments. So there is a sense in which conversion might be necessary to bring someone, or oneself, to adopt a new belief, while we would not think that this would compromise the truth of that belief in any way.

In *On Certainty*, Wittgenstein sketches an instructive dialogue that shows that there is that possibility. On the one hand, we have Moore, who like most of us is certain and could not be brought to give up his belief that the earth has existed for a very long time. On the other hand, we have a king who was brought up to think that the earth has not existed for a long time, but came into existence with him. This kind of upbringing, Wittgenstein thinks, can put the king in a position where he has no grounds, and will accept no grounds, for thinking that the earth has existed for a long time. In other words, there is no rational route from beliefs the king already has to the correct conclusion, which is that the earth has existed for billions of years:

And if Moore and this king were to meet and discuss, could Moore really prove his belief to be the right one? I do not say that Moore could not

Are There External Reasons?

convert the king to his view, but it would be a conversion of a special kind; the king would be brought to look at the world in a different way.³⁹

So Moore may well be unable to establish that the earth has existed for a long time, probably billions of years, by producing evidence the king has to accept. It is unlikely that Moore could convince him of the truth of that statement on the basis of what the king already believes, or can be brought to believe by rational reflection. Nevertheless, we will stick to our belief that the world is billions of years old. But to believe that the world is billions of years old is to believe it to be *true* that the world is billions of years old. There may indeed be overwhelming *reason* to believe that the world is billions of years old, even though the king may not be able to appreciate that truth *by reasoning*.

On the external theory, someone who has not been brought up so as to be able to appreciate the truth about their reasons will not necessarily come to appreciate that truth merely on the basis of rational reflection. In fact, it may well be impossible to convince such a person by giving reasons. What may be required is 'to look at the world in an entirely different way'. All the same, it may be true that the person has external *reasons*, even if the person cannot arrive at them, here and now, *by reasoning*. I conclude that the argument for internalism fails.

³⁹ L. Wittgenstein, *On Certainty*, § 92.

5 Back to Basics

Where does this leave us? Not, it is important to remember, with reasons externalism. For all that has been said so far, externalism may be false, while internalism might provide the true account of all reasons for action. What has been shown so far is merely that what looked like a good case against externalism is not as powerful a challenge as it may at first appear.

It is interesting to note in this connection that Williams himself may have come to doubt the cogency of his first argument. In 'Internal Reasons and the Obscurity of Blame', it is conspicuously absent, and Williams merely says that he "does not believe" that the sense of external reason statements is in the least clear, and that he "very much doubts" that any of them are true.⁴⁰ In the 'Postscript', Williams refers to the original argument, but does not seem to commit himself to it.⁴¹ With time, his strategy seems to have become more defensive.

As they stand, none of his arguments give us good grounds to call the external theory into question. But that insight gives us no better grounds, or any grounds, for accepting the external theory. Consequently, nothing in the discussion has so far given us any reason to reconsider internalism. For all that has been said so far, either theory may be correct.

This suggests that the question that we should address fairly and squarely is what reasons we have to believe the external theory. Here, I think, we may be able to make progress if we ask a question, not in the

⁴⁰ B. Williams, 'Internal Reasons and the Obscurity of Blame', p. 40.

⁴¹ B. Williams, 'Postscript', p. 93.

Are There External Reasons?

first instance about any particular statement about reasons, or statements about a particular kind of reasons, but rather about what it could be for a reason statement to be true.

One thing that is clear, both on the internal and on the external theory, is that reason statements may be called true or false, and that a person may well be mistaken in her assessment of her reasons. Just as a person may well be mistaken in her assessment of her own reasons, she may be mistaken in her assessment of the reasons of another person. So there is a sense in which we link the idea of a reason with the idea of truth and falsity, and with it, with the notion of the facts.

The question is how far this thought can take us. Here, I think, it is a helpful strategy to take a wider view and to consider the role of the notion of a reason in the context of our language. This, I think, will help us understand what the status of statements about reasons, including statements about reasons for action, really is.

Given that we speak of truth or falsity when we speak of reasons, it is natural to think of reasons as objective. Reasons seem to be in some sense *there*, making demands on our reason. This picture could be called a form of realism about reasons. I believe that realism about reasons is untenable, and that this fact undermines the basic motivation of externalism.

6 The Problem of Induction

Consider the following remark from *Philosophical Investigations*:

“Why do you believe that you will burn yourself on the hot-plate”? – Do you have reasons for this belief; and do you need reasons?⁴²

This remark concerns, at least in the first instance, the familiar problem of induction. More precisely, the remark concerns the question how a truth or set of truths about the past can give us reasons to form some belief about the future. The problem is evidently not solved by observing that normally, what we think will happen does happen. The question is what reason we had previously to think that it would.

The problem was clearly identified by Hume. As Hume famously argues, no statement about the future logically follows from a statement about the past, unless we have already built into the argument a premise to the effect that nature is uniform. But this is the very thing in question, hence the problem of induction.⁴³ Thus, we are said to infer from

(P) Flames have always burnt me

to

⁴² L. Wittgenstein, *Philosophical Investigations*, § 477.

⁴³ Hume does not use the term ‘induction’. He does discuss ‘constant conjunction’ and refers, for example, to the principle that “instances, of which we have had no experience, must resemble those, of which we have had experience, and that the course of nature continues always uniformly the same” (*A Treatise of Human Nature*, bk. I, pt. III, sec. VI, p. 89). There is a parallel discussion in the *Enquiry Concerning Human Understanding*, secs. IV and V.

(F) This flame will burn me.

But the proposition (F) evidently does not follow logically from (P) alone. It is possible that (P) is true while (F) is false. This could be ruled out, so that (F) would follow logically, only if instead of (P), we had a different premise, namely (U):

(U) Flames always burn me.

But (U) no longer expresses a truth or set of truths about the past. It makes a claim that ranges over past, present and future, and the basic problem is that we have been given no independent reason to accept (U).

Why not? For one thing, (U) is not logically necessary. It is perfectly *possible* that this flame will not burn me. If we say that this is not possible, the impossibility we have in mind here is not logical. If the justification of (U) is not logical, it must be empirical. But if it is empirical, the evidence for (U) can be no better than the evidence for (F) with which we started. If this flame, unlike all previous ones, will not burn me, (U) is simply false. We have as much, that is to say as little, reason to believe that (U) is true as we have reason to believe in (F), our original belief about the future.

7 A Sceptical Conclusion

What does this suggest about our reasons for forming beliefs like (F) about the future? Hume famously draws a sceptical conclusion. He does not express his point in terms of reasons, just as he does not express his observation as one that concerns induction or inductive reasoning. But there is at least a close conceptual connection between 'reason' and 'reasons', and 'reason' is the proper target of the sceptical attacks that Hume mounts here and elsewhere. Given this connection, it is legitimate to reconstruct the argument in terms of reasons.

Let us assume, if only for the sake of argument, that the Humean position were properly expressed like this. There is no sound deductive argument that could show on independent grounds that our beliefs about 'unobserved instances' are justified. Given that beliefs about the future are beliefs about unobserved instances, it follows that there is no sound deductive argument that could show on independent grounds that our beliefs about the future are logically justified. At the same time, there is no other source of reasons that could be appealed to. So we are led to conclude that there is *no reason at all* to form beliefs about the future. Properly speaking, it appears, we have no reason whatsoever to believe that certain things will happen in the future on the basis of our past experience, which is the only basis to provide such reasons we have got.

This does not mean, and Hume himself certainly did not take it to mean, that we will now happily put our hand into the fire. We will do

anything but that, as he is well aware. Indeed, it is a point Hume emphasizes, because it goes to show, as he believes, that reason only seems to be our guide in forming our beliefs about the future. Nature, and not reason, guides our inference. Our confidence and certainty must be explained, as one might say, not by our rational nature, but by our nature as *animals*, which is fundamental. This conclusion is shown by the naturalist Hume to be less destructive than we may take it to be, and take it to be partly because we overestimate the power of reason. But the sceptical conclusion remains in place and curbs at least our philosophical pretensions about reasons.

8 Wittgenstein on Reasons

On the face of it, what Wittgenstein says in § 477 of *Philosophical Investigations* sounds very similar. The emphasis is on the fact that we have no clear conception of our reasons when we form beliefs about the future, and that we may not even need a reason to form a belief such as that this hot-plate, like the previous ones I touched, will burn me. This is just what Hume would have said, if he had expressed his point in terms of reasons.

But a closer look shows that there is an important difference.

Wittgenstein continues:

What reason do I have to think that my finger will feel a resistance when it touches the table? What reason to believe that it will be painful if this pencil pierces my hand? — When I ask this, a hundred reasons present

Are There External Reasons?

themselves, each silencing the voice of the others. "I have experienced it myself innumerable times; and as often heard of similar experiences; if it were not so, it would ...; etc."⁴⁴

One striking fact about this passage is that Wittgenstein does not deny that when we ask ourselves what reason we have to form a belief about the future, we can give an answer that cites *reasons*. Past experience *is* a source of reasons, and Wittgenstein does not wish to deny that this way of looking at it is correct. Consider the proposition 'water boils at 100°C'. All instances that were hitherto observed were in accordance with the proposition. One might think that this amounts to *proof*:

But that the same thing has happened again is not a proof of it; though we do say that it gives us a right to assume it.

This is what we *call* an "empirical foundation" for our assumptions.⁴⁵

Past experience gives us a *right* to assume that water will continue to boil at 100°C. The assumption that it will has an *empirical foundation*, and in the light of the previous remarks about reasons, we can add that past experience with boiling water combined with our knowledge of chemistry and physics gives us *reason* to believe that water will continue to boil at 100°C.

⁴⁴ L. Wittgenstein, *Philosophical Investigations*, § 478.

⁴⁵ L. Wittgenstein, *On Certainty*, §§ 295-6. In the original German, the connection with reasons is quite explicit. What is rendered as 'empirical foundation' is 'erfahrungsmäßige Begründung'. This makes it quite clear that we are dealing with reasons ('Gründe' and 'Begründung').

Are There External Reasons?

This is all very well, you may now say, but in the light of Humean scepticism about reason, or reasons, this can only reinforce the need to say something more informative about those reasons. How *can* past experience provide reasons to form beliefs about the future? Wittgenstein is well aware of the problem, and replies to it as follows:

If it were now asked: But how *can* previous experience be a reason for assuming that such and such will happen later on? – then the answer is: What general concept of a reason for such an assumption have we got? This sort of statement about the past is simply what we call a reason for assuming that this will happen in the future. – And if one is surprised that we play such a game then I refer to the *effect* of a past experience (to the fact that a burnt child fears the fire).⁴⁶

Where the emphasis falls is clear from the previous passage from *On Certainty*. The fact that the same thing has happened every time is what we *call* an empirical foundation for our assumption that the same thing will continue to happen.

The strategy at work is that of paying attention to the way in which a word is used. More specifically, the strategy is to curb all pretensions that we either have, or need, a deeper rational foundation for our language game of giving reasons. In this way, Wittgenstein can be read as questioning what the sceptic and the person trying to refute the sceptic take for granted. Both agree that we must look for something deeper or more fundamental than our language game of giving reasons, something that provides external leverage for it. While the sceptic thinks this search

⁴⁶ L. Wittgenstein, *Philosophical Investigations*, § 480.

is likely to be unsuccessful, those who try to answer scepticism hope that there is something to be found.

Wittgenstein thinks that we should not look for something deeper or more fundamental than the language game of giving reasons: "A good reason is one that looks like *this*".⁴⁷ On the one hand, Wittgenstein refers us to a common *practice*: 'This sort of statement about the past is simply what we call a reason for assuming that this will happen in the future'. On the other hand, Wittgenstein points out that the practice has a *point*: 'And if one is surprised that we play such a game then I refer to the *effect* of a past experience (to the fact that a burnt child fears the fire)'.

This shows, first, that the practice is not arbitrary. It latches on to something given, namely, natural reactions, and the facts of nature as they are and always have been. But it also shows that, second, natural facts and natural reactions do not have the function of *replacing* reasons. Rather, they provide the basis for a linguistic practice, and the linguistic practice is that of exchanging reasons.

So what Wittgenstein says about the role of reasons in the context of induction is quite different from the basic sceptical conclusion. But perhaps it is not as different from what Hume says than it may at first appear. In fact, one might ask whether Hume could not straightforwardly accept what Wittgenstein says about reasons. Conversely, one might ask

⁴⁷ L. Wittgenstein, *Philosophical Investigations*, § 483.

whether the Wittgensteinian view is not close enough to that of Hume to count as a version of it.⁴⁸

9 The Fundamental Facts of Human Nature

Wittgenstein himself tends to appeal to fundamental facts of human nature where justifications have run out. Especially in *On Certainty*, Wittgenstein repeatedly stresses the importance of the fact that reason and belief both have and need deeper foundations, but that they do not have and do not need a deeper foundation *in reason*. Here is a characteristic statement:

I want to regard man here as an animal; as a primitive creature to which one grants instinct but not ratiocination. As a creature in a primitive state. For any logic that serves as a primitive means of communication we need not be ashamed of. Language did not emerge from ratiocination.⁴⁹

One interesting suggestion in the passage is that such a primitive 'logic' is not to be thought of as *opposed* to 'instinct'. Our inferential habits are not *added to*, but are themselves an *aspect* or *expression of*, our being 'creatures in a primitive state'. Language, and the complex abilities that are bound up with it, can bridge the gap between reason and passion. The fact

⁴⁸ Cf. P. F. Strawson, *Skepticism and Naturalism*, ch. 1; see also R. Fogelin, *Wittgenstein*, ch. 15. Fogelin argues, on the strength of remarks such as this one from *On Certainty*, that the philosophical tradition that Wittgenstein's later work most closely resembles is pyrrhonian scepticism. For Strawson's own treatment of the problem of induction, which is richer and in many ways resembles Wittgenstein's, see P. F. Strawson, *Introduction to Logical Theory*, ch. 9. For a constructive sympathetic discussion of Hume on induction, cf. B. Stroud, *Hume*, ch. 3.

⁴⁹ L. Wittgenstein, *On Certainty*, § 475.

remains, however, that beyond a certain point it no longer makes sense to ask for reasons:

Why do I not satisfy myself that I have two feet when I want to get up from a chair? There is no why. I simply don't. This is how I act.⁵⁰

And finally:

Now I would like to regard this certainty, not as something akin to hastiness or superficiality, but as a form of life. (That is very badly expressed and probably badly thought as well.)

But that means I want to conceive it as something that lies beyond being justified or unjustified; as it were, something animal.⁵¹

This is a profoundly naturalistic way of looking at the human intellect. Indeed, the point applies, and was thought by Wittgenstein to apply, to the problem of induction just as much as to the problem of extending a series, or performing calculations, or believing that there is a physical world that is structured by sequences of cause and effect.⁵² Wherever we look for foundations, we find them in forms of life, provided we dig deep enough: "What has to be accepted, the given, is – so one could say – *forms of life*".⁵³

⁵⁰ *Ibid.*, § 148.

⁵¹ *Ibid.*, § 358-9.

⁵² Cf. L. Wittgenstein, *On Certainty*, §§ 133-5, 287, 499 f., 618-20; cf. L. Wittgenstein, 'Cause and Effect: Intuitive Awareness', in: J. Klagge and A. Nordmann (eds.), *Ludwig Wittgenstein: Philosophical Occasions 1912-1951*, pp. 371-426.

⁵³ L. Wittgenstein, *Philosophical Investigations*, pt. II, p. 192.

Interestingly, Wittgenstein seems to think that moral thought is equally based on such 'primitive reactions'. Consider a series of remarks from *Zettel*:

It is a help here to remember that it is a primitive reaction to tend, to treat, the part that hurts when someone else is in pain; and not merely when oneself is – and so to pay attention to other people's pain-behaviour, as one does *not* pay attention to one's own pain-behaviour.

But what is the word "primitive" meant to say here? Presumably that this sort of behaviour is *pre-linguistic*: that a language-game is based *on it*, that it is the prototype of a way of thinking and not the result of thought.⁵⁴

In the first instance, this remark concerns the use of words like 'pain', and the question whether one can know that someone is in pain. But it suggests a natural extension to the moral case. Here, too, we want to say that the language game is an extension of our natural reactions. Moreover, it can have the status of a natural reaction, its immediacy and independence of reflection. Linguistic reactions can be second nature:

Being sure that someone is in pain, doubting whether he is, and so on, are so many natural, instinctive, kinds of behaviour towards other human beings, and our language is merely an auxiliary to, and further extension of, this relation. Our language-game is an extension of primitive behaviour. (For our language-game is behaviour.) (Instinct).⁵⁵

Moral thought itself is an extension, and expression, of our primitive reactions towards other people and ourselves. It is connected with our natural reactions, but not necessarily identical with them. Moral thought

⁵⁴ L. Wittgenstein, *Zettel*, §§ 540-1.

⁵⁵ L. Wittgenstein, *Zettel*, § 545.

and primitive reactions are *not* the same thing. But the language game itself can be as basic, as immediate and natural as a primitive reaction.

The same is true when we form our beliefs about the regularities of nature:

The nature of the belief in the uniformity of events can perhaps be seen most clearly in the case in which we fear what we expect. Nothing could induce me to put my hand into the flame – despite the fact that I *only* got burnt *in the past*.

The belief that the fire will burn me is of the same kind as the fear that it will burn me.

That I shall get burnt if I put my hand in the fire: that is certainty.

I. e. that is where we see what certainty means. (Not just what the word “certainty” means, but also what it amounts to.)⁵⁶

Again, there is no suggestion that our certain expectation to get burnt is *the very same thing* as the fear that we get burnt. It is not the same, and it is important that it is not the same, but the two are similar in nature. Both let us act with certainty in the absence of a deeper reason. Both show that we need no reasons to support our language game.

10 A Misleading Analogy

What does this suggest about our reasons? Hume makes it appear, as sceptics usually make it appear, that we are lacking something that, for some reason or other, it would be quite desirable to have. More importantly, Hume insists, as sceptics usually insist, that this something

⁵⁶ L. Wittgenstein, *Philosophical Investigations*, §§ 472-4.

Are There External Reasons?

that we lack is required for a certain form of words to be appropriate. This is most obviously so in the case of sceptical doubts about the applicability of 'knowledge'. Here, sceptical scenarios are taken to establish that the conditions that would justify us in classifying our beliefs as knowledge, or, more revealingly, 'true knowledge', are not satisfied.

Perhaps the argument even establishes the stronger claim that these conditions could not possibly be satisfied, although if that is the conclusion, we may begin to wonder what kind of knowledge this 'true knowledge' that we lack could possibly be. Is it more than a conceptual fantasy? If the claim that knowledge is impossible is to have sense, it must be shown both what would have to be true for knowledge to be real, and why, for us, no such thing is in fact attainable. Moreover, it has to be shown that this is what we mean by 'knowledge'. That is usually the form the sceptical conclusion takes. Knowledge is a case in point, but 'having reasons to believe something about the future' is just as relevant.

The Humean doubts about induction seem to have that very structure. The problem seems to be that we fall short, and must fall short, of some set standard, and that concepts like 'reason' or 'rational' appear to be tied to that ideal or standard. What the standard tells us is what would have to be true for us to have a satisfying reason, and that is why the standard is tied to, or embodies, a particular ideal.

This ideal is easy to identify. It is logical inference. Given the standard set by the ideal, a simple argument now shows that this standard could not possibly be satisfied by the kind of evidence we now have at our

Are There External Reasons?

disposal. *So there really is no reason*, we are led to conclude, to expect one thing rather than another on the basis of what happened in the past.

If we turn to Wittgenstein, we see that he describes the situation very differently. It is not, of course, that he can provide the argument that Hume was looking for and failed to find. *There is no argument* that would establish that we are able to infer, by standards set by deductive logic, that certain things will happen in the future. This cannot be done, and Hume has said everything there is to say about the reasons why. In this respect, there is a close resemblance between Hume and Wittgenstein. But there is at least a shift of emphasis. Consider the following passage:

The squirrel does not infer by induction that it is also going to need stores next winter. And no more do we need a law of induction to justify our actions and predictions.⁵⁷

One way of reading this remark would be the following. We simply act, and we do not need to justify ourselves by referring to a law of induction or anything else in order to act and make predictions as we do. We act in the absence of such reasons and justifications. Who is to stop us anyway?

This is of course a point that Hume concedes, so on this reading, this remark would merely affirm the Humean predicament. But what the mere fact of our acting and predicting without justification does not show, as Hume rightly insisted, is that we are justified in doing so. We can act without justification, and if Hume and Wittgenstein are right, this is what

⁵⁷ L. Wittgenstein, *On Certainty*, § 287.

Are There External Reasons?

we do sometimes, and when we dig deep enough, it is what we do all the time. But this does not amount to justification for our actions and predictions. Of course, the mere fact of our acting and predicting also does not show that we are *not* justified in doing so. But it is a real question where any alternative justification could be coming from. We may feel driven back towards the sceptical conclusion that we are, quite simply, not justified in forming those beliefs and making those predictions in the way we do.

As both Hume and Wittgenstein make clear, that would not necessarily be something we need to be worried about. A reason or justification is certainly not needed for us to be able to act in the ways we do. However, what Wittgenstein says is not just that we *act* without justification. He says, indeed, that we do not need a law of induction to *justify* our actions and predictions. This seems to suggest something more than that we need no reasons or justifications. It suggests that we *are* justified, and that means *justified as things are*, in making the inferences and predictions that we do. It is not that justification is in principle never available, but needed. Rather, we must realize that our beliefs about the past are all that is needed for our being justified in what we do, and that some elusive, 'real' justification is not needed in the first place.

At first sight, this may seem quite specious, but I think that Wittgenstein has an important point. To repeat, the point is not that our beliefs about the past and all our past experiences provide reasons of the kind Hume was expecting. That, Wittgenstein certainly agrees, cannot be

Are There External Reasons?

done. His reply turns on the point that these expectations are not realistic. In this way, the reply is like the one we might give to the sceptic who thinks that knowledge requires certainty. Such a sceptic may win his entire argument once we grant his assumptions. But the fact is that among those assumptions, we find expectations about knowledge that are not realistic. He is asking for too much. But he is not only asking for too much in the sense that he asks for more than we can give, which would lead to scepticism. He is asking for too much because he asks for more than is required to establish claims to knowledge.

The same, it seems to me, is true of reasons. Someone who asks for 'real' reasons to believe that certain future events will occur is asking for more than we can give. But he is not only asking for more than we can give. This granted, we would once more be led to scepticism. Someone who asks for 'real' reasons is asking for more than we need to give to be able to say, and to be justified in saying, that we have reasons, often good reasons, for our beliefs about the future.

If so, then it appears that the real target of the Wittgensteinian discussion is the assimilation of 'having reasons' to 'having logically compelling reasons'. That this is the hinge on which the matter turns comes out clearly in the following remark from *Philosophical Investigations*:

If anyone said that he is not to be convinced by information about the past that a certain thing would happen in the future, — I should not understand him. One might ask him: What, then, do you want to hear? What kind of information do you call reasons for believing this? What, then, do you call "to convince"? How do you expect to be convinced? — If *these* are not reasons, then what are reasons? — If you say that these are

Are There External Reasons?

not reasons, then surely you must be able to say what would have to be the case for us to have the right to say that there are reasons for our assumption.

For note: Here reasons are not propositions from which what is believed logically follows.⁵⁸

It is not, that is to say, that we have to acknowledge that reasons must be logically compelling reasons for them to be 'real' reasons. Rather, it is precisely because it is, at least on reflection, so very obvious that no beliefs about the past *could* logically license conclusions about the future that we should reject this model as a relevant ideal or standard of comparison for our talk of reasons. If that model is rejected, it is no longer obvious that we need something better than what we undoubtedly already have, namely beliefs about the past that *count* as reasons for beliefs about the future. Giving reasons must not be construed as a version of, or an approximation to, logical inference. It is a language game that works with different standards than the language game of logic, which constantly tempts us to look up to it as the ideal. The language game of giving reasons is autonomous, and it must be acknowledged as such to see the nature of our reasons clearly.⁵⁹

If this is the situation with respect to reasons and we now look back at Hume, we can see that in one sense nothing has changed. We are still in the Humean predicament, if that means no more than that there are no reasons of the kind Hume seems to have thought we need in order to be

⁵⁸ L. Wittgenstein, *Philosophical Investigations*, § 481.

⁵⁹ Cf. P. F. Strawson, *Introduction to Logical Theory*, p. 238: "The perfection of support is proof, but not deductive proof; it is conclusive evidence".

justified in our beliefs about the future. Indeed, it is unclear where any such reasons could be coming from. At the same time, it is clear that Hume was right to insist that no such reasons need be given to enable us to form beliefs about the future and to make predictions in the way we do. In this respect, we are once more confined to the Humean predicament.

In another sense, however, everything has changed and, to put it paradoxically, once *that* has changed, it turns out that even less has changed than Hume had thought. Once we free ourselves from misconceptions of the nature of a reason, we no longer need to withdraw the claim that we often have good reason to form a belief about the future in the way we do. It is simply not true that we have no reason to form a belief about the future on the basis of past experience, provided that we do not expect the notion of a reason to do work it cannot do.

However, there is more to the autonomy of reasons than the mere comparison with logic and deductive reasoning may suggest. The emphasis is, on the one hand, on the independence of the language game of giving reasons from that of the special case of giving logically compelling reasons. But the language game of giving reasons is equally autonomous with respect to nature and the facts of physical reality. This is equally important, because it could seem that the language game of giving reasons could only describe or express our inferential habits, or those regularities of nature to which Wittgenstein refers as providing the basis for our language game of giving reasons. But this is not so, and it is no

part of the analysis to claim that reason statements have merely descriptive meaning.

The analysis is also neutral with respect to what we should say if the facts of nature were to change quite fundamentally. There is a certain independence of the facts, as Wittgenstein points out in *On Certainty*:

In particular circumstances one says “you can rely on this”; and this assurance may be justified or unjustified in everyday language, and it may also count as justified when what was foretold does not occur. *There is a language game* in which this assurance is used.⁶⁰

It is unclear, and it is not at all surprising that it is unclear, where we would move from saying that ‘You can rely on this’ to saying that you cannot, but this lack of clarity in our reactions matches a lack of clarity in the phenomena.

Suppose that by sheer cosmic coincidence, this flame does not burn. Does it follow that you had no reason to believe it would? This would follow if claims about justification functioned like predictions, which are directly susceptible to refutation by experience. But there is no reason to accept that. You would have been justified in saying that you had a reason, just as you would have been justified in saying that you can rely on this even in the absence of good reasons. This is how the language game of giving reasons works.

All the same, if our predictions came out wrong most of the time, this part of the language game of giving reasons would have lost its point

⁶⁰ L. Wittgenstein, *On Certainty*, § 620.

Are There External Reasons?

or rationale. But this is compatible with the fact that I had a reason to expect this flame to burn me even if, contrary to my expectation, in fact it does not burn. Given our language game of giving reasons, the move from 'Flames have always burnt me' to 'This flame will burn me' is easily licensed as correct. One sparrow does not make a summer, and one irregularity does not overthrow the language game of giving reasons.

4

A Fresh Approach

In the previous chapter, I suggested that there is a sense in which the language game of giving reasons is autonomous, and that its autonomy manifests itself in two quite different ways. On the one hand, talk of reasons goes beyond giving descriptions either of natural regularities, or of ways in which humans behave. On the other hand, the language game of giving reasons is autonomous in a very different sense. It is placed in contexts in which the game has a point or purpose, and it is essentially so placed, but this does not mean that what counts as a correct move in the language game is determined by the facts of nature such that, if the facts of nature changed, what counts as a correct move in the language game would always, or directly, have to change accordingly.

These are remarks about the grammar of the language game of giving reasons, and they are remarks about that grammar in a fairly narrow sense. They concern the question how the facts about our reasons are related to some other kinds of fact, in particular, to facts of nature, and they have shown that reasons are not independent of the facts of nature, but not fully, or immediately, dependent on them either. What I now want to suggest is that the language game of giving reasons is autonomous in yet another sense. The facts about our reasons, I believe, are themselves to be explained by reference to the rules that guide the language game of giving reasons.

This suggests an important conclusion, namely that the language game of giving reasons neither answers to, nor incorporates, a relation to an independent realm of reasons. There is no normative reality, if we construe that notion on the model of a form of realism about reasons. As we shall see, this result suggests a form of internalism.

1 Physics, Fools and Heretics

Why should we think that Wittgenstein is opposed to realism about reasons, as opposed to overly ambitious versions of that view? Would a form of realism about reasons not be equally compatible with what he says? Is there any suggestion that Wittgenstein himself does not wish to remain neutral on the matter?

Whatever the answers to these questions about Wittgenstein himself may be, I think that a realist construal of the language game of giving reasons would involve a mistake, and that this mistake is shown for what it is where Wittgenstein calls our attention to what we *call*, or *count as*, reasons. As he puts it in the passage quoted earlier: 'This sort of statement about the past is simply what we call a reason for assuming that this will happen in the future'. One way of accounting for the 'simply' in the sentence is to say that when we call past experience a source of reasons for assuming certain things about the future, we need no further reason or justification for our practice of calling our past experience a source of such reasons. One might go further and say that no such justification is or could be available, so that our practice has to function in the absence of such reasons.

There is no doubt that this is correct as an interpretation of Wittgenstein. The point is that this interpretation would be equally compatible with a view that says that while there is no, and need be no, deeper justification for our counting certain facts as reasons, our calling them or counting them as such may still be thought of as a practice that involves a reference to objective facts about our reasons. These facts, in turn, can be construed along the lines of moral realism, so that it is possible to ask, with respect to every judgment about reasons, whether it is in accordance or in conflict with the facts about our reasons.

That there is a tension in this view, at least as an interpretation of Wittgenstein, can be seen when we go back to the point that we need no further justification for our practices of counting past experiences as a source of reasons for certain expectations. Wittgenstein says that 'we simply call' this kind of consideration a reason for assuming that certain events will take place in the future. But if it is correct that we need no further justification, then it seems that we must already be entitled to assume that *further*, external questions concerning the truth about our reasons simply do not arise. The question whether our reason statements are objectively true is such a question, so there is at least a suspicion that this is one of the questions that simply does not arise. The reason why it simply does not arise is easy to find, and I have mentioned it already. There are no such 'external' or 'objective' facts about our reasons, and facts of that kind have to be presupposed where we ask the external question. What Wittgenstein says about the language game of giving reasons has an important explanatory dimension.

Consider a series of remarks towards the end of *On Certainty*. First, a question is raised and implicitly answered:

Is it wrong for me to be guided in my actions by the propositions of physics? Am I to say I have no good reason for doing so? Isn't precisely this what we call a 'good reason'?¹

¹ L. Wittgenstein, *On Certainty*, § 608.

It certainly is, and question and answer are reminiscent of inductive reasoning. But now the discussion takes a new turn. We can imagine a real alternative to being guided by propositions of physics. That is to say, we can imagine a real alternative to our methods of reasoning, and that means that we can imagine someone who thinks that our methods of forming beliefs are *wrong*:

Suppose we met people who did not regard that as a telling reason. Now, how do we imagine this? Instead of the physicist they consult, say, an oracle. (And for that reason we consider them primitive.) Is it wrong for them to consult an oracle and be guided by it? – If we call this “wrong” aren’t we already using our language-game as a base from which to *combat* theirs?

And are we right or wrong to combat it? Of course there are all sorts of slogans which will be used to support our proceedings.

Where there really are two principles meeting which cannot be reconciled with one another, then each man declares the other a fool and a heretic.

I said I would ‘combat’ the other man, – but wouldn’t I give him *reasons*? Certainly; but how far do they reach? At the end of reasons comes *persuasion*. (Think what happens when missionaries convert natives.)²

The passage helps us to prevent two fundamental misconceptions. First, and most importantly, it reminds us that what is and what is not a good reason is not something that is fixed independently of our linguistic practices. This theme is present in the first remark that sets up the problem and suggests, at least by implication, that the reasons of physics are

² *Ibid.*, §§ 609-12.

paradigmatic instances of what we *call* 'good reasons'. It is also present in the second remark. When we call a different practice of exchanging reasons 'wrong', we are already relying on our language game to combat that of the other party. The fact that we are 'already' relying on our language game shows that we might be thinking of ourselves as doing something more, namely assess the different language game from a point that is independent of all limitations of our practices or points of view. But there is no such standpoint, and our criticism of a different practice of exchanging reasons can only be based on practices that are distinctly ours.

Second, the passage reminds us that when we call something a reason, we cannot, and need not, always give a further reason that can be called upon to justify it as a reason. This theme is present in the last remark, with its insistence that at the end of reasons comes persuasion, which can function like conversion. We are brought to see the world in an entirely different way, but we do not get to see it in that way by reasoning.

These two points do not entail that there is no objective truth about our reasons. But they strongly suggest that conclusion. The emphasis is, from all sides, on the essential relativity of judgments about reasons, and there is no suggestion that we can have what we may have hoped to have, namely a reassurance to the effect that our judgments about our reasons are, by and large, objectively correct. Wittgenstein does not say that we have no right to assume this, but he writes as if this is precisely the idea he is attacking.

It is important that if we accept that, it does not follow that we have no right to call judgments about reasons true or false. Compare a different kind of question that is structurally quite similar. Was it right or wrong for the Greeks to believe that during the winter, birds live on the bottom of the lakes? If we take that as a question about the truth, as opposed to a question about the available evidence, the answer is simple. The Greeks were wrong to think this. If they really believed this, their belief was false. Was it right or wrong for the Greeks to believe that there is good reason to trust the oracle, but not such a good reason to trust the physicists, to find out what will happen? Taken as a question about truth, as opposed to a question about the available evidence, the answer is, once more, that they were wrong to think this. If they really believed this, their belief was false.

Wittgenstein does not dispute this, and this is an important point. The point is *not* to deny that we can distinguish between good and bad reasons, or to deny that we can classify statements about reasons as being 'true' or 'false'. The question is what we can take ourselves to be doing when we distinguish reasons and classify statements in this way. Here, the natural analogy with our beliefs about the birds and lakes soon begins to crack. It is a fact that birds migrate, and it is a fact that we have long come to know. But our judgments about reasons cannot be supported in this way.

At the same time, we have gained a diagnostic insight. It is natural to say that it is *true*, and that we now *know*, that there is more reason to

consult the physicist than to consult the oracle to form a belief about the future. Moreover, we can back up this kind of claim in many ways by citing *reasons*. The fact that we trust physics and not oracles is anything but arbitrary, and the same is true of judgments about reasons. Given the grammatical analogy with facts of nature, given truth and knowledge, and given the reasons that back up our judgments about reasons, we can now find ourselves thinking of our reasons on a realist model. We ask whether a given belief is true or false not relative to some given beliefs or the available evidence, but absolutely, compared to the facts. Questions about reasons are naturally taken in a realistic spirit.

We all want to be realists, and we certainly want to be realists about the facts of bird migration. But should we be realists about our reasons? In my view, the answer is 'no', and if my interpretation is on the right lines, what Wittgenstein says about reasons strongly supports that reply. The evidence that we have gathered from *On Certainty* may not be conclusive. An indirect, but equally instructive treatment of the problem is found in *Remarks on the Foundations of Mathematics*, to which I will now turn.

2 The Expansion of π

It is clear that Wittgenstein builds his philosophy of mathematics around the idea of following a rule, and with it, around the idea of customs or practices. So described, the account is compatible with a platonist interpretation of mathematics, according to which mathematical

statements are true or false objectively. The two interpretations are compatible because to say that mathematics is a practice, in particular a practice of following rules, does not commit us to denying that what we say or do when doing mathematics is true or false by comparison to some objective standard. But the account of mathematics in terms of a set of rules and practices is equally compatible with the claim that there is no such objective standard, so that questions of objective truth in contexts such as those of mathematics do not even arise. The fact is that Wittgenstein radically rejected all forms of platonism, where platonism involves a claim to the effect that there are fundamental logical or mathematical dimensions of reality. Consider the following passage about the expansion of π :

Suppose that people go on and on calculating the expansion of π . So God, who knows everything, knows whether they will have reached '777' by the end of the world. But can his *omniscience* decide whether they *would* have reached it after the end of the world? It cannot. I want to say: Even God can determine something mathematical only by mathematics. Even for him the mere rule of expansion cannot decide anything that it does not decide for us.³

What the passage says about God serves to illustrate a certain platonist temptation, namely to think that mathematical reality is something that lies open to view, capable of being taken in as a whole, by something analogous to sense perception. Our limitations when it comes to calculating π would then be strictly *our* limitations, that is, epistemic

³ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, VII-41; cf. *Philosophical Investigations*, § 516.

limitations of our finite minds. God, being subject to no such limitations, could have some direct access to the mathematical reality that we can only come to know by performing the calculation. All God would have to do, perhaps, is look.

This is clearly a misleading picture. The rule of expansion does not mysteriously contain its application, and even God must *calculate* to have some knowledge of the outcome. But this does not show that there is no objective truth or mathematical reality. A platonist may say that even God must determine something mathematical by doing mathematics, that is, by following the rules. But that leaves open the possibility that following the rules enables him, as it enables us, to discover how things are in the realm of mathematics.

The important point is that Wittgenstein rejects the claim that mathematical reality is something 'there to be discovered', something that our practices of doing mathematics enable us to 'see'. Consider the thought that '777' will occur somewhere in the expansion of π . It seems to be the kind of thought that must be either true or false. After all, what would be the alternative? Wittgenstein, however, now goes on to ask:

What if someone were to reply to a question: 'So far there is no such thing as an answer to this question'?

So, e.g., the poet might reply when asked whether the hero of his poem has a sister or not - when, that is, he has not yet decided anything about it.⁴

⁴ *Ibid.*, V-9.

It is very important to be aware that this is *not* the view Wittgenstein thinks we should adopt about the expansion of π . That '777' *either* appears *or* does not appear in its extension, Wittgenstein says, is itself the expression of a rule:

It seems to me that in saying this you are yourself setting up a rule or postulate.⁵

The passage continues:

The question - I want to say - changes its status, when it becomes decidable. For a connexion is made then, which formerly *was not there*.

Notice that Wittgenstein does not say: '... is decided'. It would be misleading to suggest that mathematical questions have no answer until we have given one. This is a form of finitism, and it is associated with behaviourism in the philosophy of psychology:

Finitism and behaviourism are quite similar trends. Both say, but surely, all we have here is Both deny the existence of something, both with a view to escaping from a confusion.⁶

But in doing so, both get involved in more confusion. While behaviourism denies the relevance, or even the reality of experience, finitism denies that

⁵ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, V-9.

⁶ *Ibid.*, II-61.

rules of expansion, for example, reach beyond the point up to which a series has been developed, and is deemed by participants in the practice to have been developed correctly.

The question about the expansion of π should therefore not be modeled on the question whether the hero in the poem has a sister. This question has no answer until the poet has *decided* whether his hero has a sister. But the mathematical question does not change its status only when it is *decided*. As Wittgenstein makes clear, the point of change is when the question becomes *decidable*, and that does not entail that we have to have performed the calculation, or that no result can be called 'correct' or 'incorrect' unless we have decided that it is correct or incorrect.

The point is that many things have to be in place before such a question becomes so much as *decidable*. What has to be in place is, trivially, but importantly, a structure that allows us to decide such questions, and that structure is a practice of following a set of rules. But the structure that allows us to decide such questions is not something that is simply 'there', or something that would trace mysterious realms of mathematical reality. The question whether the expansion of π contains '777' can only arise within a system, a practice of following rules. There is an answer to it given the system, that is, given our accepted ways of calculating the expansion. But this system is not 'found'. As Wittgenstein insists, the system is 'erected':

Only within a mathematical structure which has yet to be erected does the question allow of a *mathematical* decision, and at the same time become a demand for such a decision.⁷

This is a clear rejection of the platonist idea that our practices of doing mathematics latch on to some given order that is simply 'there'. Nothing mathematical is ever found in *that* sense. Many passages to that effect are found throughout *Remarks on the Foundations of Mathematics*:

I want to say: it *looks* as if a ground for the decision were already there; and it has yet to be invented.

If you want to know more about the series, you have, so to speak, to get into another dimension (as it were from the line into a surrounding plane). – But then isn't the plane *there*, just like the line, and merely something to be *explored*, if one wants to know what the facts are?

No, the mathematics of this further dimension has to be invented just as much as any mathematics.

The mathematician is an inventor, not a discoverer.⁸

This is not the place for a detailed discussion of the Wittgensteinian views on mathematics. One thing, however, seems to me clear, or at any rate as clear as anything concerning the interpretation of this difficult material can be: Wittgenstein rejects all forms of platonism.

The truth in platonism lies in its entirely reasonable opposition to all radically reductive views. Some may hold that there is no such thing as a right way of extending a series or of following a rule *at all*. This cannot

⁷ *Ibid.*, V-20.

⁸ *Ibid.*, V-9, V-11, I-168 respectively; *cf.* Appendix II-2, II-38, V-5, V-9, V-11, VII-5 and VII-18, where the contrast between the discoverer and the inventor is employed again.

be correct. There are certainly rules, and if there are rules, there must be ways of acting and of failing to act in accordance with them. Others may hold that unless we have performed a calculation and explicitly decided on the outcome, there is no sense in which there is an answer to the question what the outcome of the calculation is. This cannot be correct either, because rules extend beyond our actual performance. But there is a right way and a wrong way of extending a series or of doing a calculation only once a practice of following a rule is established. As a result, our understanding of the rule, and our resources for giving the right answers, extend as far as that practice extends, and no further. This is not an epistemic limitation on our part. Where we have no resources for giving the right answer, this is not due to the fact that we have not yet found the right answer. It is due to the fact that as yet, there is no right answer to be found.

Some will ask: 'But is the standard set within the practice the *correct* one?' On a Wittgensteinian view of mathematics and of reasons, this question does not make sense. Standards are set within a practice, but there is no further sense in which the standard that is set within the practice can be, or can fail to be, correct. The situation is like that described by Wittgenstein in ethics: While there may be a correct answer to the question what the husband should do *given* the framework of Christian ethics, it makes no sense to ask whether the standard set by that framework is *itself* correct or incorrect, where this question calls for an

answer that transcends the Christian framework.⁹ But it does not follow that we must not call a Christian or a Nietzschean ethical outlook correct, as long as we remember that to call it correct is to adopt it.

Similarly, if we reject platonism, this does not mean that we must no longer say that it is *true* that $12 \times 12 = 144$, much less that we should say that 12×12 is *not really* 144, or something like that. Similarly, it does not follow from a Wittgensteinian view of reasons that we must no longer say that we have *good* reason to trust our physics, but no *good* reason to trust an oracle. The point is that if we say this, we must remember that what we do is to articulate, or affirm or reaffirm, a certain standard. In the first example, the standard affirmed is that expressed by the rule ' $12 \times 12 = 144$ '. In the second, the standard affirmed is that physics is a source of what we count as good reasons. The mistake is not to say that physics provides good reasons, but to think that when we say it, we do something more than reaffirm our standard. If we are inclined to say that we are right about our reasons, while the primitives are wrong, there is a sense in which this would be just another slogan to support our proceedings.

3 Reasons, Bedrock and Foundations

It seems to me that such a view of mathematics and of reasons may justly be called subjectivist, as long as 'subjective' is distinguished, for example, from 'arbitrary', 'a matter of personal preference', or 'private, as opposed

⁹ Cf. R. Rhees, 'Some Developments in Wittgenstein's View of Ethics', p. 23.

to something shared by others'. 'Subjective' captures the essential feature of the view, which consists in the denial that the facts of mathematics or the facts about our reasons are objective.

As the earlier discussion in *Philosophical Investigations* makes clear, what counts as a reason is something that is grounded in a language game, a linguistic practice of exchanging reasons. That is what is meant, or at the very least suggested, when Wittgenstein says: "A good reason is one that looks *like this*".¹⁰ In the context of a linguistic practice, criteria are established for what does and what does not count as a good reason, but the judgments made within that practice do not answer to a normative reality. That is one aspect, at least one natural extension, of the claim that for the standard of what is a good reason, there is ultimately no further reason. In particular, there is no further reason to be found in some supposed fact such as that what we think is a good reason *really is* a good reason, where that 'is' receives a realist interpretation. In the light of our discussion of ethical value, we can say that the way in which a proposition can correspond or fail to correspond to physical reality has no counterpart here.¹¹

The upshot is that practices of giving reasons are contingent, and neither have nor need some further rational foundation, just as the practices of counting and of calculating are contingent, but neither have nor need further foundations in reason, intuition, or objective fact. To

¹⁰ L. Wittgenstein, *Philosophical Investigations*, § 483.

¹¹ Cf. R. Rhees, 'Some Developments in Wittgenstein's View of Ethics', p. 24.

think of these practices as answering to a normative or mathematical reality is to indulge in philosophical mythology:

The difficult thing here is not, to dig down to the ground; no, it is to recognize the ground that lies before us as the ground.¹²

This, I think, captures the essence of the Wittgensteinian claim that we should stop looking for a deeper explanation of our practices – the sense in which “our disease is one of wanting to explain”.¹³

It does not follow, and it is important that it does not follow, that our practices cannot be explained. They can certainly be *causally* explained: “For mathematics is after all an anthropological phenomenon”; “Following a rule is a human activity”.¹⁴ So there is a sense in which “what we are supplying are really remarks on the natural history of man”:

not curiosities however, but rather observations on facts which no one has doubted and which have only gone unremarked because they are always before our eyes.¹⁵

What is of interest to philosophers are not the causes, but what we may call the ‘rational’ connections between concepts and the facts of nature.

“Our interest certainly includes the correspondence between concepts and very general facts of nature”, but our interest does not fall back on the

¹² L. Wittgenstein, *Remarks on the Foundations of Mathematics*, VI-31.

¹³ *Ibid.*

¹⁴ *Ibid.*, VII-33; VI-29.

¹⁵ *Ibid.*, I-142; cf. VI-49, VII-33; *Philosophical Investigations* pt. I, §§ 25, 415; pt. II, sec. xii, p. 195.

“possible causes of the formation of concepts”.¹⁶ Philosophy does not have to engage in science. All the same, what I called the ‘rational’ connections between facts of nature and our practices of forming concepts can help us to understand those practices.

What does this entail for our talk about reasons? In a first step, we can extend the application from the specific case of mathematics to practices of taking measurements generally. Consider the proposition ‘12 inches = 1 foot’:

No one will ordinarily see this last proposition as an empirical proposition. It is said to express a convention. But measuring would entirely lose *its ordinary character* if, for example, putting 12 bits each one inch long end to end didn’t ordinarily yield a length which can in its turn be preserved in a special way.¹⁷

Again, we have a rule, and there is such a thing as a ground or a framework or foundation for the rule in our ordinary practices of measuring. These practices, in turn, are grounded in some highly general facts of nature. Does this mean that the proposition ‘12 inches = 1 foot’ says something *about* these practices and facts of nature?

No. The proposition *is grounded in* a technique. And, if you like, also in the physical and psychological facts that make the technique *possible*. But it doesn’t follow that its sense is to express these conditions.¹⁸

¹⁶ L. Wittgenstein, *Philosophical Investigations*, pt. II, sec. xii, p. 195.

¹⁷ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, VII-1.

¹⁸ *Ibid.*, VII-1; cf. IV-53, VI-9 and VII-18.

After all, to deny that '12 inches = 1 foot' is not to say that rulers are not rigid or that we do not take measurements in the same way. The fact that the rule is grounded in a technique does not affect its status, which is, and remains, that of a rule. A rule *qua* rule is "detached". It stands "as it were alone in its glory", although "the facts of daily experience" are what gives it its importance.¹⁹ Compare the task of describing the office of a king:

What I have to do is something like describing the office of a king; – in doing which I must never fall into the error of explaining the kingly dignity by the king's usefulness, but I must leave neither his usefulness nor his dignity out of account.²⁰

The office of a king is much more than a merely 'useful' institution. All the same, its usefulness must not be left out of the account. More importantly, the fact that it is useful does not necessarily conflict with saying that the king has, first and foremost, great authority and dignity.

So it appears that highly general facts of nature have something important to contribute to our understanding of our practices, and that this is true even in the case of mathematics, where we are not dealing with conventions of the kind expressed by statements like '12 inches = 1 foot'. Mathematics is not 'arbitrary'. The rules of mathematics are much more than mere 'conveniences', and it does not follow from the fact that every practice is contingent that we can adopt any set of rules we like, while continuing to think of ourselves as doing mathematics. One reason for that

¹⁹ *Ibid.*, VII-3.

²⁰ *Ibid.*

is that not many of the things I could do would still count as doing mathematics. If it is replied that these conceptual restrictions are themselves purely arbitrary, the reply is that this is precisely what they are not, and the explanation of that fact once more appeals to fundamental facts of nature:

“Then according to you everybody could continue the series as he likes; and so infer *anyhow!*” In that case we won’t call it “continuing the series” and also presumably not “inference”. And thinking and inferring (like counting) are of course bounded for us not by an arbitrary definition, but by natural limits corresponding to the body of what can be called the role of thinking and inferring in our life.²¹

Again:

We should presumably not call it “counting” if everyone said the numbers one after the other *anyhow*; but of course it is not simply a question of a name. For what we call “counting” is an important part of our life’s activities. Counting and calculating are not—e.g.—simply a pastime. Counting (and that means: counting like *this*) is a technique that is employed daily in the most various operations of our lives. And that is why we learn to count as we do: with endless practice, with merciless exactitude; that is why it is inexorably insisted that we shall all say “two” after “one”, “three” after “two” and so on. — “But is this counting only a *use*, then; isn’t there also some truth corresponding to this sequence?” The *truth* is that counting has proved to pay.²²

Language relates to a way of living.²³ Rules are bound up with certain expectations, and such expectations are essential for our lives. That is why mathematics is no mere ‘convenience’:

²¹ *Ibid.*, I-116.

²² *Ibid.*, I-4.

²³ *Ibid.*, VI-34.

A Fresh Approach

Now can it be said that the concepts which mathematics produces are a convenience, that essentially we could do without them?

First and foremost the adoption of these concepts expresses the *sure* expectation of certain experiences.

We do not accept e.g. a multiplication's not yielding the same result every time.

And what we expect with certainty is essential to our whole life.²⁴

If we choose to describe mathematics as a 'mere convention', we must bear in mind that there may be a deep need for the convention.²⁵

Still, it may be asked whether we have any clear conception of a different practice, which seems to be presupposed in the idea that concepts are formed in the context of contingent linguistic practices. I think that the answer to this question is a 'yes' with some qualifications. The qualifications are due to the fact that the changes we would have to imagine would indeed be changes in facts that are fundamental.

On the one hand, these could be changes in our customs and practices. This can be illustrated by means of the Wittgensteinian wood sellers. (It is interesting to note that this example follows the remark about the natural history of man.) Here, we are asked to imagine a group that has some different ways of measuring and paying for the wood with which they trade. Say they sell their timber by cubic measure, rather than by weight, or some other standard. Are these people *right* in doing so?

²⁴ *Ibid.*, IV-52.

²⁵ *Ibid.*, I-74.

Well, are we *right* in using other measures? Neither question makes much sense if we read the question 'Who is right?' as asking for 'absolute correctness'.

This reply may well seem too quick. Imagine that the timber is piled up in heaps of varying heights and sold at a price proportionate to the area that is covered by the piles. We point this out and get the following reply: "Of course, if you buy more timber, you must pay more":

How could I shew them that - as I should say - you don't really buy more wood if you buy a pile covering a bigger area? - I should, for instance, take a pile which was small by their ideas and, by laying the logs around, change it into a 'big' one. This *might* convince them - but perhaps they would say: "Yes, now it's a *lot* of wood and costs more" - and that would be the end of the matter. - We should presumably say in this case: they simply do not mean the same by "a lot of wood" and "a little wood" as we do; and they have a quite different system of payment from us.²⁶

The point is that the difference between us and this community would have to be much deeper than it may initially appear. All the same, there is the genuine possibility that the people who sell wood in this way can consistently apply their standard. When confronted with ours, this *might* convince them that they are involved in some deep inconsistency - but it does not *have to*. The depth of our commitments and our practices does not stand in the way of the Wittgensteinian account. If anything, it provides it with additional support.

²⁶ *Ibid.*, I-150.

Conversely, imagine our practices as they now are, and imagine certain facts of nature to be radically different. Wittgenstein contemplates this kind of scenario in the following example:

Put two apples on a bare table, see that no one comes near them and nothing shakes the table; now put another two apples on the table; now count the apples that are there. You have made an experiment; the result of the counting is probably 4. (We should present the result like this: when, in such-and-such circumstances, one puts first 2 apples and then another 2 on a table, mostly none disappear and none get added.) And analogous experiments can be carried out, with the same result, with all kinds of solid bodies. – This is how our children learn sums; for one makes them put down three beans and then another three beans and then count what is there. If the result at one time were 5, at another 7 (say because, *as we should now say*, one sometimes got added, and one sometimes vanished of itself), then the first thing we said would be that beans were no good for teaching sums. But if the same thing happened with sticks, fingers, lines and most other things, that would be the end of all sums.

“But shouldn’t we then still have $2 + 2 = 4$?” – This sentence would have become unusable.²⁷

Once more, the changes in the facts of nature would have to be very radical indeed to support entirely different practices. But once more, this does not show that mathematics is not based upon contingent linguistic practices. If anything, it strengthens the view that it is:

When, in teaching the child to calculate, one asks, “How do $3 + 2$ make 5”? – what is he supposed to show? Well, obviously he is supposed to move three beads up to 2 beads and then count the beads (or something like that).

And if the child now shows how 3 and 2 make 5, then he shows a procedure that can be regarded as a ground for the rule “ $2 + 3 = 5$ ”.²⁸

²⁷ *Ibid.*, I-37.

²⁸ *Ibid.*, VI-9.

All the same, ' $2 + 3 = 5$ ' is not empirical, and this marks the leap the pupil needs to make.

In the light of these examples, I think we should accept that there is a general connection between the facts of nature and the rules of mathematics. The point is that this general connection is not as direct and as immediate as it may be thought to be. If the world in which we lived were such that objects spontaneously came into and went out of existence, and if this happened on a scale that would render the practice of counting things useless, then the practice would have lost its point. It does not follow that the rules of counting, or addition or subtraction, are themselves empirical in nature, or susceptible to refutation by the facts. But the practice needs a certain context, and as such, it needs a foundation in the facts.

On the view of reasons that emerges, judgments about reasons do not answer to, or incorporate, a reference to a normative dimension of reality. This, I think, is what Wittgenstein suggests when he calls our attention to what we call, or count as, reasons. Some of the evidence for this account was found in the *Remarks on the Foundations of Mathematics*, which presents a powerful alternative to platonist accounts. Similar conclusions can be drawn, and are drawn by Wittgenstein, about all platonist accounts of logic. Here, one focus of interest is the hardness of the logical 'must'. That there are more general applications of the point, I argued, is suggested by the parallel discussion of our practices of

measuring. In the remainder of this chapter, I consider what this view suggests about the metaphysics of our reasons.

4 A Realistic View of Reasons

The language game of giving reasons is not the same as that of calculating or inferring or measuring objects. In fact, it is obviously not the same, and that is why Wittgenstein not only points out, but emphasizes that the language game of giving reasons should not be construed on the model of logical inference. Nevertheless, there are some striking similarities. (1) The language game of giving reasons is autonomous in the same sense as the language game of mathematics. (2) The language game of giving reasons is not empirical in nature, and in this respect, it is again strikingly similar to mathematics. (3) Nevertheless, and again like the language game of mathematics, the language game of giving reasons is grounded in, or based upon, empirical, contingent facts. The facts in question are both facts of nature that are highly general, such as facts about the laws and regularities of nature. They include, just as importantly, facts about our form of life and our psychologies.

In what follows, I will try to explain these parallels and draw some general conclusions that strike me as important for the discussion of reasons for action. The conclusion that I aim to establish is that a subjectivist account of reasons for action strongly supports a form of internalism.

The most important difference between a subjectivist account of mathematics or of logic and a subjectivist account of reasons will be that such an account will portray the former, but not the latter, as rules. Propositions of logic and of mathematics do, whereas propositions about reasons do not, normally present themselves as *a priori* necessary truths. They cannot be refuted by experience and concern not merely what is actual, but extend to what is possible or necessary. So their modal range is wider.

This is not so in the case of reasons, which are closer to, and more closely interwoven with, the facts of our experience. For one thing, claims about reasons may include a range of modal notions, but the mode in which these claims present themselves is factual. It is the case, or true, that someone has a reason, and this does not yet say anything about the modal status of that claim. For another, claims about our reasons are more closely interwoven with the way the world happens to be. If we suspect someone to be denying that $2 + 2 = 4$, we will not quite understand him. If someone rejects our claims about our reasons, our reactions can be like that, and that is an important fact. But they need not, and typically do not, take that form. When someone thinks that a certain belief about reasons is false or unjustified, we often find what such a person says perfectly intelligible, and it may even convince us. '2 + 2 = 4' is timeless, necessary, and secured from refutation by experience. 'Past experience gives me reason to believe that this flame will burn me' is, in one sense, also timeless. But what that

means is that the laws and regularities of nature that form the foundation of such claims are thought of as unchanging. Given a radical change in the facts of nature, the claim that we have reason to expect certain events may well become *false*, and not just useless. Becoming useless is, by contrast, all that can possibly happen to a proposition that is mathematical or logical.

At the same time, claims about reasons are not subject to refutation by experience in the way predictions are. Even if it turns out that this flame does not burn me, it may be true that I had a reason to believe it would. This shows, I think, sufficiently that the language game of giving reasons is autonomous in at least three different respects. First, it is not empirical, if that means that the language game of giving reasons has the function of giving descriptions or making predictions. If anything, it helps us to arrive at adequate descriptions or to make predictions, but that does not mean that it has the function of describing or predicting. Second, it is not empirical in the different sense of having the function of giving an account of some supposed normative reality. In this respect, it is analogous to mathematics and logic, which are not to be thought of as describing some supposed logical or mathematical reality. Third, the language game of reasons is autonomous in the sense that it is not immediately susceptible to refutation by experience. The claim that I had reason to expect that something would happen may not need to be withdrawn even if what I expected did not happen.

Still, it is clear that the language game will have a point only if we already take the fundamental facts of nature and of human life for granted. Compared to the language games of logic and of mathematics, its dependence on contingencies and context is indeed much more immediate and firm. Reasons are often *given* by the facts. What we call a reason to expect that this flame will burn me, for example, is *that this is what has always happened in the past*. Conversely, what that reason is a reason for is the belief, or expectation, *that this flame will burn me if I touch it now*. But there is a more complex dependence also. What counts as a reason for what is something that is responsive to the facts, which is one more unsurprising implication of the view that reasons are not there to be appreciated, but grounded in our linguistic practices.

What does all this have to do with ethics? In particular, what does it have to do with our reasons for action? The basic connection lies in this. Reasons for action should be understood as being grounded in a set of linguistic practices. In particular, they should be understood as being grounded in a practice of exchanging reasons. This practice will not be the same as that for giving reasons for beliefs about the future. After all, we are dealing with reasons for action. These are reasons for doing certain things, and for doing some things rather than others. But this is a difference in the subject matter, not a difference in the metaphysics of our reasons. The grammatical relation '*... is a reason for ...*', or '*...speaks in*

favour of...’ remains the same, even if the right hand side contains a verb of action.

If so, this suggests a first important observation about reasons. It is true that, as externalists tend to insist, reasons are given by the facts. They are not, or at any rate not normally, given by our attitudes. Thus, the reason why I should not jump off the tree is that it will hurt me, not, or not immediately, that I do not want to be hurt. The reason for doing my friend a favour is that it will please her, not that I want her to be pleased. Even in the case where reasons are contingent upon some present desire or a preference, they are not a mere reflection of that preference. My reason for buying a bar of chocolate is not, or not immediately, my desire to eat chocolate. Properly speaking, what gives me a reason to buy a bar of chocolate is *the fact that I want to have chocolate*. This, I think, is the important truth in reasons externalism. But as we shall see, this idea does not yet take us to externalism.

If this account of reasons is correct, it also does not yet amount to some form of expressivism. There is no reason to believe, for example, that statements about reasons for action all share a single function, such as that of serving to express some attitude.²⁹ In particular, there is no reason to think that such statements can never be true or false. Of course such statements may function as mere expressions of some attitude. But they do not have to. Even if they essentially involved expressions of some attitude,

²⁹ Cf. S. Blackburn, *Ruling Passions*, p. 264-6.

it would not follow from that fact alone that these expressions cannot be true or false. If we are to believe that reason statements can be neither true nor false, this needs independent argument.

In fact, there is a deeper problem with expressivism. Expressivists may think of reason statements as expressive tools. If so, they can freely accept that reasons statements can be true or false. This would be a familiar move in keeping with the tenets of expressivism. Problems arise when this analysis is thought to provide the *only* sense, or the only *relevant* sense, in which a reason statement may be true or false. For in that case, a reason statement would always be used to express some attitude towards an action, or at best acceptance of a set of norms that speak in favour of that action.³⁰

Here, a genuine dimension of our talk about reasons for action seems to have been lost, namely the possibility that someone may have a reason to do something that we think is undesirable or wrong. We may well wish to say, like Williams does in his example, that while *we* think that the sexist man ought to be nicer to his wife, it may still be *false* that he has any such reason. It is not obvious to me how Blackburn or Gibbard would account for this possibility. Given that it is a possibility, I think we should be wary of endorsing expressivist analyses across the board to

³⁰ The latter formulation is preferred by A. Gibbard; cf. *Wise Choices, Apt Feelings*, p. 163.

make sense of our statements about reasons. There seems to me to be a truth in the idea that reason statements can be bluff, as Williams puts it.³¹

At the same time, there is an important truth contained in the expressivist account. As expressivists insist, reason statements should not be interpreted as involving reference to a realm of normative facts that is, in some sense, simply there. This is how the realist construes them, who portrays the facts about our reasons as the moral realist portrays the moral facts. As we have seen, there is every reason to resist this picture. The point is that we do not have to accept it even if we think that we should remain neutral on expressivist analyses. Reasons are based on distinctive moves in a linguistic practice, and that practice does not answer to, or itself incorporate, a normative dimension of reality. For all that, statements about reasons may be true or false.

5 From Subjectivism to Internalism

If this account of reasons is correct, we should abandon platonism about reasons. But why should a subjectivist account of reasons lead us to a form of internalism? In fact, it is not obvious that the two views are even compatible. After all, subjectivists believe that when it comes to deciding what is and what is not a reason for action, there are no objective facts of the matter at all. But one way of construing internalism would be to say that internalism tells us which claims about reasons are objectively true

³¹ Cf. B. Williams, 'Internal and External Reasons', p. 111; cf. B. Williams, 'Postscript', p. 95.

and which ones false. According to internalists, no statement about reasons for action can be correct unless it is backed up, in the appropriate way, by some present motivation. That looks like a claim that tells us how things stand with respect to our reasons. Moreover, I have said that a subjectivist account of reasons can respect one of the insights of externalists, namely that reasons are given by the facts, not, or not immediately, by preference or desire. Is it even possible to be an internalist once we accept this version of subjectivism about reasons?

The point of fundamental importance is that it is a mistake to construe the question what reasons there are on the model of an empirical question. Facts about our reasons are not like the facts of nature in that they are somehow already there, waiting to be discovered, and making demands on our reason. Externalists think of reasons in this way, and to the extent to which internalists are also guided by this picture, their view will equally conflict with a subjectivist account of reasons. Williams remains neutral on the matter, but only because he never clearly addresses the question. There is some evidence, however, that he tended towards some form of subjectivism about reasons.

Consider the following objection to internalism. It starts from the thought that what counts as a reason for an action is a social or institutional question, and not an individual or psychological question as internalism seems to suggest. It is then pointed out that in addition to the reasons that are based on the psychologies of agents, prudential or moral

reasons, for example, are also recognized as reasons. So it would seem that a complete account of reasons must go well beyond the limits set by the internalist.

Here Williams thinks there is a truth, but he insists that it does not help the externalist:

What can rationalize or render intelligible various kinds of action is certainly a social, and in some part an institutional, matter. It can be a question of historical, anthropological, and philosophical interpretation, how far these various practices are, also, variable and local. But whether they are local or more widely spread, it will equally be true that they need a basis in individual psychologies.³²

Whatever the merits of his argument, Williams seems to agree that what counts as a reason may be variable and local, and so relative to social context.

Williams seems to think that the externalist who pursues such a strategy still finds himself caught up in a dilemma. Reason claims need to be given a basis in individual psychologies, and *ex hypothesi*, this is precisely what they do not have on the externalist account. Reason claims, Williams seems to be saying, need to get a grip on the agent, but they do not get a grip on those agents who have not internalized the practice. Where they do not get a grip, we are left with nothing more than a blank assertion that there are such reasons. There is no problem of this kind

³² B. Williams, 'Postscript', p. 94.

against the background of a practice in which certain reasons are accepted, and this is of course precisely what internalists predict.

Is this a good reply to the objection? One thing that is clear, and rightly emphasized by Williams, is that we have no reason to expect that those who are unaware of the fact that they have some reason for action will be convinced by the mere claim, directed at them by others, that they have such a reason. This is the point about the limits of justification again, and the externalist is well advised to admit it. All justification comes to an end somewhere, and the same is true of all claims about reasons. Deliberation, including taking account of the relevant reasons, is here and now, and here and now it may well be too late for an agent to be able to appreciate his reasons, or to be able to appreciate them fully and correctly. All the same, it may be true that he has those reasons. There is no inconsistency involved in the idea that an agent has a reason for action of which he is, and may remain, unaware.

So it seems that this reply is not enough to undermine externalism, if it is not coupled with overly ambitious views about the epistemology of reasons. Its metaphysics could remain intact. Williams could reply that there is a specific way in which a reason statement must have a basis in the individual psychology of an agent. But if that is more than a blank assertion that internalism is correct, this is likely to lead back to the original argument that invokes a requirement of rational deliberation, and we have seen that there is no such requirement.

All the same, I think that Williams has touched on an important issue when he draws attention to the fact that what counts as a reason for an action is a social and, to some extent, an institutional matter. This suggests an argument to the effect that internalism is, after all, the right account of practical reasons.

The crucial point is that on the subjectivist view of reasons as I have described it, the language game of giving reasons is basic, and neither answers to, nor incorporates, a normative reality. Suppose we accept this view of reasons, but still wish to maintain some form of externalism. We will now have to say that it can be true that agents have reasons to do certain things even though they are not at all disposed towards doing those things in any of the ways recognized by internalism – after all, if this turns out to be impossible, we are stuck with the internalist account. To imagine this, we have to imagine an agent who is not disposed towards acting in a particular way, but still has a reason to act in that way. And how can that be? If it is true, as subjectivism says it is, that there are no reasons independently of our practices of taking certain considerations as reasons, and if the person does not take the considerations that we take as reasons for action *as* reasons for action, or takes them as reasons for acting *differently*, then how could it be true, of him, that he has such a reason? What exactly could the agent be said to be missing if he refuses to accept our claim that he has such a reason?

The situation will be one in which the agent does not share, or only shares in part, our normative perspective. He will, that is to say, not be involved in the same language game. At the limit, he will cease to play the game of giving reasons, because he might be disposed to act in ways that we could no longer recognize as moves in such a game. Here, the situation would be that of someone who persists in counting differently. We might say that the person does not play the language game we play, that he uses a different standard of measurement, perhaps like those who pay for timber in a different way. But that is all – there is no deeper sense in which that person could be said to get things wrong. Perhaps we would feel very alien from such a person. But there is no sense in which we could be said to know something about his reasons that he refuses to accept, and this amounts to the denial of externalism. If we now say that our practice of taking certain considerations as reasons is *correct*, this is harmless if it merely reaffirms our standard. If, in saying this, we aspire to do something more, then we are bluffing.

6 Reconsidering Reality

Given that subjectivism about reasons naturally leads to a form of internalism, there is one obvious line of reply available to the externalist. The argument in favour of internalism rests on the assumption that our practice of taking certain considerations as reasons does not answer to, and does not itself incorporate, some normative reality. But what reason is

there to accept that? I have relied on a subjectivist account of mathematics, and of measuring more generally, but it is clear that this account is open to dispute. Moreover, it would be desirable to have independent arguments in favour of subjectivism. I have suggested that there is no such thing as a normative reality, a standard against which our judgments about reasons could be measured. What reason is there to accept this?

To come to terms with this objection, I will once more consider the parallel objection to subjectivist accounts of mathematics and of logic. The basic point is that platonist accounts will always raise more questions than they answer. First, there are substantial metaphysical and epistemological concerns about a normative dimension of reality. If there is such a reality, how are we to think of it, and how can we come to know it? More importantly, given that we can have knowledge of it, how is it that so many people are, and have been, ignorant of it? It is not easy to say.

Even if these problems can be overcome, there is still a substantial problem. Other things being equal, the simpler and more economical account of the phenomena is preferable. This is subjectivism, so there is some independent reason to prefer subjectivism to any platonist contender. What I now want to suggest is that what Wittgenstein says about the role of 'true' and 'fact' and 'correspondence to reality' strongly supports this line of argument. This analysis completes the case against all forms of platonism.

It might be thought that when it comes to questions of truth and correspondence, Wittgenstein does little more than to appeal to minimalist or deflationist accounts of truth. But in fact, there is much more to his analysis than that. Consider the following exchange:

“But isn’t there a truth corresponding to logical inference? Isn’t it *true* that this follows from that?” – The proposition: “It is true that this follows from that” means simply: this follows from that. And how do we use this proposition? – What would happen if we made a different inference – *how* should we get into conflict with the truth?³³

It appears that Wittgenstein admits that the deflationary manoeuvre by itself does not settle much, if anything, of theoretical interest. The question is *how to account* for the idea that different inferences would be ‘false’ or ‘incorrect’. What is the standard compared to which the truth or falsity of inferences can be established? We have a fairly clear conception of the way in which a person gets into conflict with the truth if that person denies, for example, that water boils at 100°C or that Great Britain is a monarchy. If someone does not accept some rule of modal logic or a mathematical proof that we think of as incontrovertible, we have no clear conception of *what it would be* to get into conflict with the truth.

It is not, perhaps, that we have *no* conception of such a conflict, or that the very idea of some objective fact in logic or in mathematics could be set aside as incoherent. The point is that we have no clear conception of what such a conflict *looks like*, and that is one more way in which the way a

³³ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, I-5.

physical reality corresponds or fails to correspond to a proposition has no counterpart here.³⁴

This is not to say that we cannot express ourselves by saying 'It is a fact that $2 \times 2 = 4$ ', and the same is true of talk of reasons. But this is not evidence for the truth of some form of realism, but part of a diagnosis of a philosophical mistake. What looks like a close kinship is in fact only a superficial relationship:

We are used to saying "2 times 2 is 4", and the verb "is" makes this into a proposition, and apparently establishes a close kinship with everything that we call a 'proposition'. Whereas it is a matter only of a very superficial relationship.³⁵

Again:

I am trying to say something like this: even if the proved mathematical proposition seems to point to a reality outside itself, still it is only the expression of acceptance of a new measure (of reality).³⁶

The mathematical proposition is a rule, and that means that it belongs with the measurements, not with the things measured. The same is true of reasons. Propositions about reasons do not describe reality. They are instruments for dealing with reality. There is no question of a correspondence here.

³⁴ Cf. R. Rhees, 'Some Developments in Wittgenstein's View of Ethics', p. 24.

³⁵ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, Appendix III-4.

³⁶ *Ibid.*, III-27.

Reason statements, as we saw, are not themselves expressions of rules. To the extent that they are, they are clearly not expressions of rules like the ones that we find in logic or in mathematics. Propositions about reasons are, for example, neither necessary nor *a priori*. But despite these differences, there is a deeper parallel. If someone disagrees with us about our reasons, and if that disagreement persists on reflection, then all that is left to say in good conscience is that we have parted company. There is nothing objective that one party could be said to understand or know, while the other party remains ignorant of it. In that respect, the situation is once more like that in logic or in mathematics:

"This follows inexorably from that." – True, in this demonstration this issues from that.

This is a demonstration for whoever acknowledges it as a demonstration. If anyone *doesn't* acknowledge it, doesn't go by it as a demonstration, then he has parted company with us even before anything is said.³⁷

'He has parted company with us' *does not mean* 'He is wrong, and we are right', if that is to be more than some emphatic affirmation of our practices. The metaphor of 'parting company' suggests a view from an external standpoint. When we take up that standpoint, it is clear that while we follow one path, another group follows a different path – and that is all there is to it.

³⁷ *Ibid.*, I-61.

This point is reinforced by an appeal to practices and shared responses. What ‘keeps us on the rails’ when doing mathematics is a set of shared responses, and as we have seen, these responses are not arbitrary. Indeed, this is part of the reason why we have no need for the platonic picture. We are *certain* that in the series ‘add 2’, 2006 comes after 2004. This makes it plain that we have *no need* for the platonic picture:

“But do you mean to say that the expression ‘+ 2’ leaves you in doubt what you are to do e.g. after 2004?” – No; I answer “2006” without hesitation. But just for that reason it is superfluous to suppose that this was determined earlier on. My having no doubt in face of the question does *not* mean that it has been answered in advance.³⁸

It is superfluous – but could the platonic picture not be *true*? Wittgenstein replies:

The *truth* is that counting has proved to pay. – “Then do you want to say that ‘being true’ means: being usable (or useful)?” – No, not that; but that it can’t be said of the series of natural numbers – any more than of our language – that it is true, but: that it is usable, and, above all, *it is used*.³⁹

Again, the point can be applied to reasons. Our practice of exchanging reasons has no foundation in normative facts, but that does not mean that it has no foundation. It has a foundation in our practices. Ultimately, there is nothing more available to keep us on the rails. But nothing more is

³⁸ *Ibid.*, I-3.

³⁹ *Ibid.*, I-4.

needed. The ultimate supports of reasons are found neither in reason, nor in objective fact, but in the human form of life that lies at their foundation.

It is noteworthy that the Wittgensteinian view does not amount to the claim that mathematical and logical statements are 'true in virtue of meaning' or 'true in virtue of convention'. This kind of account evidently fails to solve the real problem, because it merely pushes it one step further back. If we are trying to understand the nature of logical inference, for example, it certainly does not help to give an account of it that says that logical propositions are true in virtue of certain conventions about the use of signs from which all truths of logic follow. The nature of logical inference is precisely what needs to be explained, and it is not much of an explanation to be told that logical inference is the product of convention and subsequent logical inference. The fact that we *accept* a rule is fundamental, and this acceptance is what Wittgenstein invokes to bridge the gap between empirical, contingent propositions and the *a priori* necessary ones that we find in logic and in mathematics.

Others will feel that this kind of account already concedes too much to the empiricist. Quine, for example, suggests that mathematical truth is not in principle exempt from alteration or revision in the light of experience. To the extent that it is, Quine believes, this is due to a policy of doing whatever needs to be done to keep the entire system stable.⁴⁰ If Wittgenstein is right, however, propositions of experience and rules of

⁴⁰ Cf. W. V. Quine, *Pursuit of Truth*, p. 14.

representation are not on the same plane. Wittgenstein admits that rules may change, that language games may lose their point because the world in which they have their role is changing, and that new language games arise. After all, counting – and that means: counting *like this* – has proved to pay. But it does not follow that the rules of mathematics and of logic have a status similar to that of empirical propositions. The office of a king is not adequately described as a useful institution. The office has a special dignity. Rules in mathematics or in logic are useful. But they are more than that, because they are completely isolated from experience and not in principle susceptible to confirmation or refutation by experience. That is what *their* special dignity consists in.

Still, one may feel that there is something paradoxical in the idea of trying to account for truths of logic and of mathematics in terms of contingent linguistic practices. If linguistic practices form the foundation of rules, and if these practices are to be thought of as contingent, then different practices must at least be possible. But if different practices are *possible*, how can the truths of mathematics and logic be *necessary*?

I think that the correct reply to this objection is to ask whether what we are said to be looking for is something that we can have *on any account* of logical and mathematical necessity, and I suspect that the answer is ‘no’. This strongly suggests that those who think that we need more than what Wittgenstein gives us are in the grip of a picture.

Imagine, for the sake of argument, that there is a platonic realm of abstract objects, or that there are necessary truths or facts. How would this help? After all, what we would have is an arrangement of platonic objects, and all such arrangements could presumably be otherwise. Alternatively, all we have is a set of necessary facts or truths. But *if it is a fact* that there are such truths or facts, then it must be possible for these facts to be otherwise. And if that is so, how can such a model solve our modal paradox?

Is there a way to escape this dilemma? It is not easy to see what it could be. The problem is that for whatever we say to be so, it must be possible to say that it is otherwise. So the problem seems to be inherent in the very thought of abstract objects or of necessary truths or facts. This, in turn, is not surprising, because such talk is evidently modeled on familiar talk about contingent facts and empirical objects. As long as there is no way out of that dilemma, *any* account of logic or of mathematics will eventually face the same problem. And if we eventually have to face the same problem, should we not face it straight away?

7 Is This Still Internalism?

This line of thought completes the case against all forms of platonism, and with it, against externalism about reasons. If this is the situation, it suggests the somewhat surprising conclusion that the key to understanding the case for reasons internalism lies neither in questions of

motivation, nor in questions of rational deliberation. It lies in the metaphysics of reasons itself. The mistake is to believe that reasons are features of a normative reality that is somehow 'there' and makes demands on us in virtue of our sensitivity to reasons.

The picture that emerges, then, is that of internalism, but with an important difference. The dispositions on which reasons are based are, primarily, the dispositions to take certain considerations as reasons, and not, or not primarily, the dispositions to act that an agent presently has. If that is so, then the basis of all reasons for action is indeed a set of dispositions on the part of the agent. In this respect, subjectivists are in agreement with internalists. But they can also capture one important truth in externalism, which is that desires are not the only kind, and perhaps not even the usual kind, of input for a practical reason statement. Realists tend to say things like 'Reasons are given by the facts, not by our desires', and it is difficult to deny that this is so. The point is that this does not yet commit us to externalism. The question is *how* the facts provide the reasons, and here the internalist has something more to say.

There is a further difference between my version of internalism and the version that Williams originally described. Even if we grant that all reasons for action must be based on dispositions, we should not conclude from this that every disposition gives rise to reasons for action. Intuitively, we will probably think this is false.

Williams himself has a reply to this problem. He admits that some elements in the subjective motivational set do not necessarily give rise to reasons for action. But he thinks that this phenomenon will be explicable in terms of false belief or mistakes in reasoning.⁴¹ Where no such mistakes are present, Williams thinks that the presence of a desire is not only a necessary, but also a sufficient condition for the truth of a reason statement.⁴²

I will not try to dispute this claim directly, but I think there is good reason to be sceptical that desires of all kinds give rise to reasons for action. Perhaps nothing much is lost if we accept that they do. But I think that it is worth emphasizing that desires and reasons *can* come apart, and that the present account provides the conceptual space for that to happen.

Consider the example of desiring money. This is something that most of us do, for obvious reasons. Money is required for the basic necessities of life, and much else besides. However, it is clear that we do not accept that even a strong and quite basic desire, like the desire for money, will justify *anything*. To take an extreme example, money gives us no good reason to kill another person and take theirs. Indeed, we may find that this thought should not even occur to us, and if it were to, that would be a moral failing. For a person that is ethically disposed as she should be, this kind of consideration no longer provides a reason for killing another person *at all*.

⁴¹ B. Williams, 'Internal and External Reasons', p. 103.

⁴² B. Williams, 'Internal Reasons and the Obscurity of Blame', pp. 35 f.

It is unclear how internalism as described by Williams can make sense of this possibility. After all, even ethically disposed people desire money, and it is all too obvious that stolen money serves just as well as legitimate income to buy us what we want. So it is evidently not true that either the desire or the reasons for desiring money have now disappeared. There are no obvious mistakes in belief or reasoning. All the same, we want to say that the desire for money does not give us a reason, not even an immediately overruled reason, to kill another person. How can that be?

The answer to this is that the internalist can leave it open which desires give us reasons to do certain things, and which desires give us no reason for action whatsoever. Similarly, internalists can leave it open which desires give us *some* reason for action, and which ones give us *less* or *more* reason for action compared to that desire. The point is that our judgments about reasons will be a reflection of our linguistic dispositions, in particular, our disposition to take certain considerations - which may involve a reference to the presence of desires - *as* reasons for action. This seems to me to be another truth in the externalist account that is well worth preserving.

But is this view still close enough to that of Williams to count as a version of it? What I have said so far has emphasized the differences. But it is clear that these are differences against the background of a fundamental agreement. On either view, we should not think of our reasons for action as being radically independent of our desires and

dispositions. In short, we should not think of them as independent of our subjective motivational set. This makes it appropriate, I think, to describe my account as a version of internalism.

A crucial point to remember is that Williams thinks of the basis for reason statements in intentionally broad and open terms. Thus, he writes that on any adequate account, the subjective motivational set must contain “patterns of evaluation, patterns of emotional reaction, personal loyalties, and various projects, as they may be abstractly called”.⁴³ This feature of internalism certainly allows us to include the special set of dispositions that I called the dispositions for taking certain considerations as reasons. So there is what may be called a *formal* parallel between the two accounts. But there is a *material* parallel, too. The parallel consists in this, that our dispositions for taking certain considerations as reasons for action, will not, in general, be independent of our given desires and dispositions.

First, there is good reason to expect a general convergence of desires and reasons because both respond to fundamental facts of nature, and are shaped by facts of our psychologies and the human form of life. But there will also be demand for psychological coherence at the individual level. This demand will typically ensure, and at the very least encourage us to work towards, a state in which our judgments about our own reasons and our dispositions are, to a certain extent, brought into line. So we should not think of internalism as I have described it and

⁴³ B. Williams, ‘Internal and External Reasons’, p. 105.

internalism as Williams originally described it as two radically different theories. They do, however, take different perspectives on the nature of our reasons, and with it, on the nature of internalism.

Perhaps what Williams says in a paper called 'The Primacy of Dispositions' comes rather closer to the view that I envisage. Here, Williams argues that all value ultimately rests in dispositions. Dispositions, Williams explains, are unique among the objects that have value, because they make all evaluation possible:

In a certain sense, they give the value of those objects, even though the value of those objects cannot be reduced to theirs.⁴⁴

If I am right, the primacy of dispositions has even wider repercussions. It reaches beyond the realm of ethical evaluation into the realm of reasons and rationality.

8 Transcendental and Empirical Idealism

If the dispositional view of reasons is correct, it entails a significant shift of perspective in our understanding of practical reason. In conclusion, I will now compare this shift with a different, but instructive change of perspective in the history of philosophy.

⁴⁴ B. Williams, 'The Primacy of Dispositions', in: B. Williams, *Philosophy as a Humanistic Discipline*, essay 6, p. 75.

In the *Critique of Pure Reason*, Kant faces the charge of maintaining a particularly radical form of idealism. This is not surprising, given that he argues that the framework in which talk of objects as being in space and time, as being causally connected and so on, has a subjective origin. As he puts it, space and time have 'no reality' independently of the perceiving subject. But as Kant points out, once we confine ourselves to the framework set up by our *a priori* intuitions and refrain from advancing a claim about the nature of things in themselves, it will once again be true to say that physical objects exist independently of us. But this is so not despite, but rather in virtue of, the truth of transcendental idealism. Moreover, Kant does not deny that there are 'things in themselves' and that they exist independently of us. So Kant is not an idealist at the transcendental level.⁴⁵

Compare this with a view like that of Bishop Berkeley, who is famous – or perhaps infamous – for his version of idealism. Berkeley denied that anything exists over and above space and time and spiritual beings. As a result, we are faced with the conclusion that matter is unreal, a mere appearance.

As Kant points out, or rather insists, this is not the kind of idealism he defends. He does not deny that there is something that exists independently of us, and he denies it at neither the transcendental nor the empirical level.

⁴⁵ This point comes out particularly clearly in the *Prolegomena*; cf. I. Kant, *Prolegomena*, § 13, esp. Note III.

This, then, is the transition from what Kant calls 'empirical' or 'material' to 'transcendental' or 'formal' idealism.⁴⁶ With respect to everyday purposes, this kind of idealism is meant to leave everything as it is. It should be clear, however, that transcendental idealism *really is idealism*. In fact, it is an extremely radical form of idealism. The truth is that what Kant suggests is much more radical than anything that had been contemplated by 'empirical' idealists like Berkeley.⁴⁷

What does all this have to do with practical reasons? Consider what may be called 'empirical' or 'material' internalism. This is the view that all reasons for action are given by present motivations or desires. As we have seen, this view of reasons for action is not very plausible. Reasons are typically given by facts, not by desires. It would be a mistake, and a mistake analogous to the mistake made by empirical idealists as Kant describes them, to tie our reasons to desires in such an immediate, 'material', or 'empirical' way.

The fact is that the framework of exchanging reasons as a whole is a social institution. It is based on dispositions of agents, so it is, in the sense I described, 'merely subjective'. Like space and time and physical reality on the Kantian conception, reasons have 'no reality' independently of us. This might be called the transition from an 'empirical' or 'material' internalism to a 'transcendental' or 'formal' one. According to formal

⁴⁶ Cf., e.g., I. Kant, *Critique of Pure Reason*, A 368-81, B xxxiv, 52, 274-9, 519-25.

⁴⁷ In this respect, I fully agree with Strawson: "Kant, as transcendental idealist, is closer to Berkeley than he acknowledges" (P. F. Strawson, *The Bounds of Sense*, pp. 21 f.).

internalism, it is not true that reasons for action are always, or even normally, given by present desires or dispositions. But that is true only as long as we describe the language game of giving reasons. If we ask about the metaphysics of our reasons, we will say that the language game of giving reasons does not involve a reference to a normative reality. The resulting view is, in a way, much more realistic than any view that claims that reasons are directly given by desires. But at the same time, it is radically and openly idealist, and in this way fundamentally opposed to all forms of externalism about reasons.

The general conclusion is that despite appearances, subjectivists can make good sense of our ability to act for reasons. They are neither forced to question that we act for reasons, nor do they have to think of this ability as driven by desire. This is as it should be, because it touches on, and preserves, the fundamental truth that human agency has a distinctive character, and that this distinctive character is bound up closely with our ability to reason and to speak a language.

Ever since antiquity, man has been described as the creature that has both the ability to reason and to speak a language. One way of describing the result of this investigation is to say that these abilities are connected much more closely than they may at first appear to be.

5

Frameworks and Foundations

The line of thought that I have been developing suggests that in at least two different respects, we should abandon the idea that ethical language answers to, or incorporates, such a thing as an objective ethical or normative reality. I suggested that this is the situation with respect to moral and evaluative judgments, and statements about reasons for action. This connection is, of course, not arbitrary. Language should be understood as our tool for constructing, or creating, reasons and values. This process does not take place in a vacuum, and it is anything but arbitrary. Our practices are shaped by fundamental facts of human nature, and they evolve in response to a reality that exists objectively. But reality contains no such thing as an ethical or normative order. This is subjectivism.

In this chapter, I begin to weave these strands together. Given the materials that I have assembled, the goal is to give an accurate, and in the relevant sense naturalistic, account of the foundations of our ethical outlook. I suggest that these foundations can be found in unsupported practices – in our *accepting* certain values and ideals, and in our *accepting* certain considerations as reasons for action. These practices provide the framework, or scaffolding, of ethical thought. In this way, they set the limits within which rational argument and justification in ethics are possible. So interpreted, ethics has a foundation neither in reason, nor in sentiment. Its foundation lies in forms of life.

The text on which this view is modeled is one of the last written by Wittgenstein, posthumously published as *On Certainty*. It has long been recognized as one of his masterpieces, and ethics is a field where its insights can be usefully applied.

1 An Elusive Piece of Writing

A point that has been stressed throughout the present work is that the later Wittgenstein provides a number of important insights and ideas that, while originally applied in other contexts, can be put to work in ethics. We have already seen how his views on logic and on mathematics and on reasons, facts and the autonomy of grammar suggest a number of surprising applications in moral philosophy. What I now want to show is that the same is true, in yet another way, of the ideas expressed in *On*

Certainty. It seems to me that there are striking parallels between our epistemic situation and our practical predicament, given the description of the former that we find in Wittgenstein. This parallel shows that there is a link between theoretical and practical reason that runs deep, but has gone largely unnoticed. This leaves us with the task of making sense of this elusive piece of writing.

I begin by identifying the theme of *On Certainty* that strikes me as most important for the purposes of this enquiry. The book deals with a number of related concepts, all of which are familiar topics in epistemology. There are detailed discussions of 'belief', 'doubt', 'knowledge', 'reason', 'evidence', 'experience', 'ground', 'mistake' and many others. So in a sense, the grammatical investigation of 'certainty' and 'certain' is only one of several interlocking topics. Still, certainty is one of the main themes. The picture that emerges is one where *certainty*, not knowledge, argument, or evidence lies at the foundation of all our believing. On the one hand, beliefs are 'held fast' by beliefs that lie around them. But the system as a whole has no further foundation, either in experience, or proof, or fact. On the other hand, it needs no such foundation. The reason is that our beliefs are indeed held fast:

Knowledge is in the end based on acknowledgment.

At the foundation of well-founded belief lies belief that is not founded.¹

¹ L. Wittgenstein, *On Certainty*, §§ 378, 253.

There certainly is such a thing as giving reasons and justifying the evidence. But:

giving reasons, justifying the evidence, comes to an end; – but the end is not that certain propositions immediately strike us as true, i.e. it is not a kind of *seeing* on our part; it is our *acting* which lies at the bottom of the language-game.²

Ultimately, the foundations of belief lie in *our believing*, that is, in *accepting* that which is believed. This is the sense in which our *acting* lies at the foundation of the language game.

The thought is that all exchanging of arguments and giving justifications takes place against a background of assumptions that goes largely unquestioned, and must go largely unquestioned if giving arguments and justifications is to be so much as possible. This background forms a ‘scaffolding’ or ‘framework’ of our thought, and its essential feature is that as a whole, it lies beyond justification.

Elsewhere, Wittgenstein portrays this scaffolding or framework as the ‘river-bed’ of thought, which provides the moving waters that run through it with support, firm boundaries and a direction.³ River-beds can change with time, and change under the influence of water that flows through it. Even so, there has to be a bed if there is to be a river, just as in thought there remains a distinction between those things that are

² *Ibid.*, § 204.

³ *Ibid.*, §§ 96-9.

moveable and those things that stand fast. These ideas suggest that there might be a parallel to ethics.

2 Knowledge or Unshakeable Conviction?

What are the fundamental propositions? Consider the following list given by G. E. Moore in 'A Defence of Common Sense':

There exists at present a living human body, which is *my* body. This body was born at a certain time in the past, and has existed continuously ever since, though not without undergoing changes; it was, for instance, much smaller when it was born, and for some time afterwards, than it is now. Ever since it was born, it has been either in contact with or not far from the surface of the earth; and, at every moment since it was born, there have also existed many other things, having shape and size in three dimensions (in the same familiar sense in which it has), from which it has been *at various distances* (in the familiar sense in which it is now at a distance both from that mantelpiece and from that bookcase, and at a greater distance from the bookcase than it is from the mantelpiece; also there have (very often, at all events) existed some other things of this kind with which it was *in contact* (in the familiar sense in which it is now in contact with the pen I am holding in my right hand and with some of the clothes I am wearing)).⁴

Other propositions that Moore mentions as the list continues are the following: 'The earth had existed for many years before my body was born'; 'I am a human being'; 'I have had dreams'; 'Other human beings, during their lifetime, had many different experiences'. Moore claims that this list contains propositions "every one of which (in my own opinion) I *know*, with certainty, to be true".⁵

⁴ G. E. Moore, 'A Defence of Common Sense', in: G. E. Moore, *Philosophical Papers*, ch. 2, p. 33; cf. 'Proof of an External World', *ibid.*, ch. 7.

⁵ G. E. Moore, 'A Defence of Common Sense', pp. 32 f.

Wittgenstein asks at the very outset of *On Certainty* whether that is really so:

If you know that *here is one hand*, then we'll grant you all the rest.⁶

And again, adding more explanation:

If Moore says he knows the earth existed etc., most of us will grant him that it has existed all that time, and also believe him when he says he is convinced of it. But has he also got the right *reason* for his conviction? For if not, then after all he doesn't *know* (Russell).⁷

There is a crucial difference between *believing* something with certainty and having the right *reason* for that belief, and it is important that the two can come apart. As Wittgenstein portrays the situation, Moore has every right to say that he *believes* these propositions, with certainty, to be true. But this does not yet amount to knowledge: "That he does know remains to be shown".⁸

Clearly 'I know it' does not, by itself, establish that I know it, and in this respect 'I know' belongs to a fundamentally different grammatical category than 'I believe'. There is nothing wrong with saying 'I know that there is a tree in front of me', either in an ordinary or a philosophical context. But 'I know' does not guarantee that I know, and while it is true

⁶ L. Wittgenstein, *On Certainty*, § 1.

⁷ *Ibid.*, § 91; cf. §§ 18, 19, 59.

⁸ *Ibid.*, § 14.

that what I know cannot be false for grammatical reasons, I may only think I know.

At the same time, Wittgenstein acknowledges that Moore has a valid point, and this provides the second, more constructive half of the diagnosis. Propositions like those on the list may or may not amount to knowledge, but they are clearly, and quite independently of that, things that we believe with certainty. We *cannot but* believe them, act on them, and take them for granted in our lives. This mode of belief is fundamental, and its hallmark is a form of certainty that is *not* a form of knowledge on the Wittgensteinian account.

Suppose that we replace 'I know' with 'I am of the unshakeable conviction'.⁹ Now what Moore says is no longer presumptuous or unjustified. If I say that I am of the unshakeable conviction that this is a tree, there is no need for me to answer 'How?' or 'Why?'. If I say that I know that there is a tree in front of me, by contrast, questions such as 'How do you know?' or 'Can you justify this claim to knowledge?' will always be right around the corner. When I truthfully say that I am of the unshakeable conviction, then setting self deception aside, this is sufficient to justify the claim that I am: "'Knowledge' and 'certainty' belong to different *categories*".¹⁰ If anything, the form of certainty that Wittgenstein describes provides *foundations* for our claims to knowledge.

⁹ Cf. *ibid.*, § 86.

¹⁰ *Ibid.*, § 308.

This might make it appear as if we needed something more than certainty, when certainty is taken to amount to nothing more, or deeper, than unshakeable conviction. But the fact is *that there is* that kind of unshakeable conviction, and the fact that this conviction is unshakeable is enough to refute the idea that in the absence of a reason, our beliefs might as well be abandoned. The truth is *that they cannot* be abandoned, and this reveals a deeper sense in which what we believe with certainty is not subject to the will. There is already something paradoxical in the idea of adopting a belief simply because one wants to.¹¹ Here, the paradox would be that we call something into question and try to support it with more argument, while what is called into question is part of the framework of all doubts and arguments. On the one hand, there are no reasons that might be given to support the framework set up by what we believe with certainty. On the other hand, there are no reasons that could be given so that, in response to them, those fundamental beliefs would have to be abandoned.

Moreover, it does not follow that we have to 'accept the existence of things outside us merely on faith', as Kant complains in the *Critique of Pure Reason*, and thinks of as a 'scandal to philosophy'.¹² For Moore, this is untrue because he thinks that we have certain knowledge of the world around us. Where Kant complains that proof was lacking, Moore claims that proof is easily available. For Wittgenstein, by contrast, the claim that

¹¹ Cf. B. Williams, 'Deciding to Believe', in: B. Williams, *Problems of the Self*, ch. 9.

¹² I. Kant, *Critique of Pure Reason*, B xxxix, note.

we must take the existence of things outside us 'merely on faith' is untrue because it misdescribes the nature of our unshakeable conviction. That something stands fast for me, he writes, "is not grounded in my stupidity or credulity".¹³ Our conviction is not based on reasons – not even the reason that we should believe what we have no alternative to believing. We *cannot but* believe what is expressed by framework propositions, and this is not normally true of those things that we 'merely accept on faith'. Believing something without justification does not amount to believing it without right.¹⁴

3 A Peculiar Logical Role

The picture of our epistemic situation that emerges, then, is one in which some propositions have a peculiar logical role. They are believed with certainty. This certainty is not a form of knowledge, but it does not, for that reason, turn into blind faith.

One of the central tasks Wittgenstein sets himself as the remarks that constitute *On Certainty* unfold is to illuminate the nature of this certainty, and to give a more detailed description of the propositions we believe with unshakeable conviction. What is their peculiar logical role?

When Moore says he *knows* such and such, he is really enumerating a lot of empirical propositions which we affirm without special testing; that is

¹³ L. Wittgenstein, *On Certainty*, § 235.

¹⁴ Cf. L. Wittgenstein, *Philosophical Investigations*, § 289.

propositions which have a peculiar logical role in the system of our empirical propositions.¹⁵

For many of those fundamental propositions, we do not even have a clear idea of what such a 'special test' would look like. That is one respect in which they have a peculiar logical role.

This might suggest that Wittgenstein thinks that all fundamental propositions are completely isolated from experience. But this would be a mistake. It is not true, for example, that none of them could turn out to be false. Nor is it true that their falsity could not be established by experience. The obvious example is 'No one has ever been to the moon'. Far from being something we believe with certainty to be true, this proposition is one that we no longer even believe. But in 1950, one could not simply *assume* that someone had been to the moon:

For this demands answers to the questions "How did he overcome the force of gravity?" "How could he live without an atmosphere?" and a thousand others which could not be answered.¹⁶

Now they can be answered, but that does not refute, but rather confirms the fundamental and unshakeable conviction concerning the fundamental facts of nature. If someone were to challenge that system of evidence, we would no longer know quite what to say. Suppose we get the reply:

¹⁵ L. Wittgenstein, *On Certainty*, § 136.

¹⁶ *Ibid.*, § 108.

“We don’t know *how* one gets to the moon, but those who get there know at once that they are there; and even you can’t explain everything.” We should feel intellectually very distant from someone who said this.¹⁷

An equally intriguing case is ‘Books do not simply vanish’. Unlike ‘No one has ever been to the moon’, this is something we still firmly believe. Were someone to challenge us, we would say, once more, that our entire system of evidence forbids us to believe that books just vanish. But at the same time, it is clear that this is a genuine possibility. As Wittgenstein wittily points out, books disappear all the time, so it is not that we have evidence that they do not simply vanish that is derived from constant direct observation. But then, there is no evidence that books evaporate either, and everything in our system of belief speaks for, and nothing against, the proposition that books do not simply vanish. So here we have an example of a proposition of which it is true (a) that we simply accept it, (b) that it is supported by experience and (c) that in response to new experience, it could turn out to be false. Should we find that under particular circumstances books did vanish away, we would alter our beliefs accordingly. No doubt this would be a most radical revision of our system of belief, but it is not clear that we could not do this and having done it, carry on.¹⁸

This provides a useful illustration of the point of the analogy between ‘hardened’ and ‘fluid’ propositions and a river and its river-bed. A river-bed might change under the influence of water that runs through

¹⁷ *Ibid.*

¹⁸ *Cf. ibid.*, § 134.

it, just as 'hardened' propositions can change in response to changes in experience. What stands fast at one time can be washed away:

It might be imagined that some propositions, of the form of empirical propositions, were hardened and functioned as channels for empirical propositions that were not hardened but fluid; and that this relation altered with time, in that fluid propositions hardened, and hard ones became fluid.¹⁹

Note that Wittgenstein does not oppose 'hardened' and 'empirical' propositions. The distinction he draws is better expressed by opposing empirical propositions that are hardened to empirical propositions that are not, but subject to change in the flux of experience. Hardened propositions are not rules. They are *like* rules in a particular respect, namely insofar as they are used, at times, for taking measurements. Thus,

if someone were to say "so logic too is an empirical science" he would be wrong. Yet this is right: the same proposition may get treated at one time as something to be tested by experience, at another as a rule of testing.²⁰

This is why the bank of the river consists partly of 'hard rock', partly of 'sand', which 'now here now there gets washed away or deposited'.²¹ Foundational propositions are not, all at once, subject to confirmation or disconfirmation by experience. In this respect, they are like rules. But unlike rules, they have representational content. It makes sense to ask, of a foundational proposition, but not of a rule, whether it is true or false.

¹⁹ *Ibid.*, § 96.

²⁰ *Ibid.*, § 98.

²¹ *Ibid.*, § 99.

Moreover, what is used at one time like a rule may at another time be open to empirical investigation. The distinction between 'hardened' and 'fluid' propositions is not a sharp one. But it does not follow from that lack of sharpness that there is no distinction to be drawn.

4 Building a World Picture

If this is the situation, then we should admit that even at the fundamental level, changes of belief in response to new experience must take place against a background, or a framework, or a system of beliefs that are not subject to revision all at the same time. This system is constituted by a set, which is not necessarily sharply defined, of propositions that stand fast, and in this way provide the basis for revisions in belief in response to new experience. In a different simile, these beliefs provide the hinges on which doubts and confirmations turn.²² Take the example of 'I have two hands'. If I were to doubt it, I could no longer be sure of *anything*. It is a tempting thought that this belief of mine rests on my sense experience. After all, we want to say, my two hands are here before me! But this is misleading. I may *think* that I constantly confirm my belief by sense experience, but as Wittgenstein points out, this is an illusion. If I were to doubt that I have two hands, I could not, for example, make sure that I have two hands by looking:

²² Cf. *ibid.*, §§ 341-3.

If a blind man were to ask me "Have you got two hands?" I should not make sure by looking. If I were to have any doubt of it, then I don't know why I should trust my eyes. For why shouldn't I test my *eyes* by looking to find out whether I see two hands? *What* is to be tested by *what*?²³

It belongs to the logic of empirical investigation that certain things "are *in deed* not doubted".²⁴ The status of 'in deed' is that of a practical necessity.

Wittgenstein writes: "My *life* consists in being content to accept many things".²⁵ Again:

I really want to say that a language-game is only possible if one trusts something (I did not say "can trust something").²⁶

'Can' indicates a rational entitlement, something that would have to be based on sufficient reasons. It is not the 'can' of 'being able'. It is the 'can' of 'being entitled'; the sense in which 'You really cannot do this' is not relevantly answered by 'Yes, I can', or 'Watch me do it'. To justify ourselves against these charges, we need to cite justifying reasons, and the same is true when we cite evidence for 'fluid' propositions. But at the fundamental level, further justification is neither available nor needed.

This might suggest that sense experience is something separate from belief, but this, too, is an illusion. Experience is already saturated with, and informed by, belief. There is no such thing as justification by experience *independently* of our system of belief:

²³ *Ibid.*, § 125.

²⁴ *Ibid.*, § 342.

²⁵ *Ibid.*, § 344.

²⁶ *Ibid.*, § 509.

One wants to say “*All* my experiences show that it is so”. But how do they do that? For that proposition to which they point itself belongs to a particular interpretation of them.

“That I regard this proposition as certainly true also characterizes my interpretation of experience.”²⁷

Again, now with respect to ‘experience’ and ‘reason’:

But isn’t it experience that teaches us to judge like *this*, that is to say, that it is correct to judge like this? But how does experience *teach* us? We may derive it from experience, but experience does not tell us to derive anything from experience. If it is the reason (and not just the cause) for our judging like this, then there is no further reason for seeing this in turn as a reason.

No, experience is not the ground for our game of judging. Nor is its outstanding success.²⁸

The interesting suggestion here is that experience provides *reasons*, as opposed to causes, for our judgments and beliefs. This suggestion allows us to reject the myth of the given without falling into the other extreme.²⁹ While the myth of the given is built on the idea that sense experience carries its interpretation with it, and so can provide some independent foundation for belief, there is an extreme reaction against the myth of the given that amounts to saying that sense experience can only ever be a cause, and not a reason, for a judgment or belief. Wittgenstein does not deny that sense experience is a source of reasons. The point is that it is a source of reasons because *we treat it as a source of reasons*. Sense experience

²⁷ *Ibid.*, § 145.

²⁸ *Ibid.*, §§ 130-1.

²⁹ The ‘myth of the given’ was named and criticized by W. Sellars; cf. W. Sellars, *Empiricism and the Philosophy of Mind*.

is a source of reasons, but these reasons are not independent of our linguistic practices. In particular, it is not independent of our linguistic practices of *taking* certain experiences as reasons for beliefs and judgments.

This raises the question of whether Wittgenstein commits us to a form of relativism with respect to our most fundamental beliefs. A first point to note in response to this question is that Wittgenstein does not deny that even ‘hardened’ and ‘foundational’ propositions can be true or false. He is quite explicit on this point: “The *truth* of certain empirical propositions belongs to our frame of reference”.³⁰ If we remind ourselves of the grammatical point that beliefs and facts can come apart, there is no need to worry that truth itself, or what is true, must be in every instance relative to our beliefs or our experiences. This point is brought out clearly in the following remark, which stands as a constant reminder over the following discussion:

From its *seeming* to me – or to everyone – to be so, it doesn’t follow that it *is* so.

What we can do is ask whether it makes sense to doubt it.³¹

Truth essentially transcends our system of belief.

Those who are worried by this kind of threat may now point to passages such as the following: “If the true is what is grounded, then the

³⁰ L. Wittgenstein, *On Certainty*, § 83.

³¹ *Ibid.*, § 2.

ground is not *true*, nor yet false".³² But note that Wittgenstein only asserts the implication, not the antecedent of the implication. Thus, he is not committed to accepting that the true *is* what is grounded, so he is not committed to accepting that the ground *is* neither true nor false. This may sound rather specious, but the fact remains that Wittgenstein claims that the *truth* of certain empirical propositions belongs to our frame of reference. If the truth of those empirical propositions belongs to our frame of reference, then even the ground must, as a rule, be either true or false.

Perhaps the remark is really an echo of a point about the epistemic status of the framework:

Everything that I have seen or heard gives me the conviction that no man has ever been far from the earth. Nothing in my picture of the world speaks in favour of the opposite.

But I did not get my picture of the world by satisfying myself of its correctness; nor do I have it because I am satisfied of its correctness. No: it is the inherited background against which I distinguish between true and false.³³

Here, the idea of 'inheritance' reinforces the sense that 'satisfying myself of its correctness' applies to propositions *within* the world picture, but not to the picture as a whole. Even the thought *that I am* satisfied of its correctness would be misleading, if it means more than *that I believe*, and believe with unshakeable conviction. It portrays my epistemic situation as if it were backed at least by a *belief* about my reasons. But there are no

³² *Ibid.*, § 205.

³³ *Ibid.*, §§ 93-4.

reasons, and that means that the belief that I have reasons at the fundamental level is exposed as an illusion.

This is an important result, but it might merely reinforce our sense that we are heading for a view in which certainty is built on relativity. After all, Wittgenstein describes the propositions that make up the riverbed as a kind of 'mythology'.³⁴ Clearly, what stands fast for people changes over time, and it is something different for different people. Rational resources to resolve such disagreements will come to an end. And if I have my world picture neither because I *have* satisfied myself of its correctness, nor because I *am* satisfied of its correctness, but because it is the inherited background against which I distinguish between true and false, it is a genuine question what right I have to call my picture of the world the correct one, or to criticize others. This, in turn, raises the question of how, with respect to the foundations, there can be such a thing as firm conviction and belief.

The answer is that there could not be anything but firm conviction and belief, and that the threat of relativity is not as serious in practice as it may be thought to be. First, the fact that we will at some point run out of reasons does not show that the propositions for which they were meant to serve as reasons are false. It also does not show that they are true, of course, but foundational propositions are such that we cannot but believe them, and to believe them is of course to believe them to be true. Wittgenstein writes:

³⁴ *Ibid.*, § 95.

Really “The proposition is either true or false” only means that it must be possible to decide for or against it. But this does not say what the ground for such a decision is like.³⁵

Nor does it say, one might continue, that there have to be *any* grounds for such a decision. Propositions may be believed to be true in the absence of all reasons.

Moreover, reasons themselves are already informed by, and contribute to the content of, our world picture:

So reasonable mistrust must have a reason?

We might also say: “the reasonable man believes this”.

Thus we would not call anybody reasonable who believed something in spite of scientific evidence.³⁶

These remarks suggest a substantial conception of rationality, and in this way, they help to cure an important philosophical confusion. It is not that we have foundational beliefs on the one hand, and a conception of rationality on the other. This kind of picture encourages us to look for independent justifications for foundational beliefs, and in the absence of such reasons, it will tend to undermine our confidence in those foundational beliefs. The fact is that foundational beliefs *include* conceptions of good reasons, and that our conception of a good reason is already shaped by, and itself a part of, our world picture.

³⁵ *Ibid.*, § 200.

³⁶ *Ibid.*, §§ 323-4.

This might still be worrying, and Wittgenstein himself is well aware that what we think of as reasonable is subject to change and variation:

But what men consider reasonable or unreasonable alters. At certain periods men find reasonable what at other periods they found unreasonable. And vice versa.

But is there no objective characteristic here?

Very intelligent and well-educated people believe in the story of creation in the Bible, while others hold it as proven false, and the grounds of the latter are well known to the former.³⁷

This might sound like defeat. But once again, the passage seems to me best construed as a gentle reminder that even science, and with it, a particular system of evidence and practice of exchanging reasons, has no privileged foundation. Reasons have no recognizable, objective characteristics that would allow us to tell by looking whether or not they are good reasons. That is an illusion, and the comparison with religion brings that out.

I do not think that it follows, or was taken by Wittgenstein to follow, that we should simply accept what religious people say as a different, but equally well founded belief or unshakeable conviction, or that this conviction will always be such that it could not possibly be reached by reason and argument. It is true that we will reach a point at which all that anybody can say is 'Here I stand', 'This I believe with certainty', or 'This is the position I will hold'. But it is a very different question whether religious beliefs can be located at that level. Reasons

³⁷ *Ibid.*, § 336.

must run out eventually, but it is far from clear that they run out that soon.³⁸

Still, it may be felt that there is a deep tension between this picture of our epistemic situation and the ideal of scientific objectivity. Can we accept the Wittgensteinian picture *and* the claim that physics has a hope of representing the world 'as it is anyway'? I do not see why we could not. Wittgenstein criticizes certain epistemic aspirations, but in doing so, he does not challenge, but relies on our beliefs about those things that exist independently of us. At the very least, these beliefs include beliefs about a physical reality, and a belief to the effect that this reality exists quite independently of us. The distinction between the elements of our world picture that are in some sense subjective and the elements that are not is drawn *within* our world picture.³⁹ The claim that we cannot appeal to reasons or experience to *convince* ourselves of its correctness does not entail that that our world picture can only be *our* world picture, and has no better hope of objectivity than any other.

We are *certain* that there is a physical reality of which we are a part. There is nothing in the Wittgensteinian picture that puts this belief in jeopardy. If anything, the Wittgensteinian picture reinforces and supports it. The kind of objectivity that is certainly ruled out by the account is one that would require us to step out of our world view altogether and compare it with a 'pure' reality that is neither experienced nor

³⁸ On this point, *cf.* J. C. A. Gaskin, *Hume's Philosophy of Religion*, chs. 6 and 7.

³⁹ *Cf.* B. Williams, 'Philosophy as a Humanistic Discipline', p. 185, and 'Replies', p. 209.

conceptualized. Similarly, there is no hope that we can support our world picture from within, either by appealing to what might be 'given' in experience, or by *a priori* reasoning. But it is obvious enough, and was obvious enough already, that these hopes for independent reassurance are a fantasy.

5 Theory and Practice

The thought that all belief rests ultimately on belief that is ungrounded, and in that sense, on a certain way of acting, is a most compelling one. In what follows, I apply that insight to the practical domain. The thought is that there is a close analogy between our theoretical and practical commitments, and that the picture that emerges in *On Certainty* can be usefully applied to illuminate the nature of our practical commitments.

One way of describing the analogy would be to say that our beliefs are held fast by beliefs and attitudes that lie around them, where the metaphor of holding fast indicates both a psychological and rational connection. The system as a whole, however, neither has, nor needs, a deeper rationale. Similarly, I believe, our practical beliefs and attitudes, which might include beliefs about what would be right, or good, or worthwhile or desirable, are held fast by beliefs and attitudes that lie around them, and once again, this indicates both a psychological and rational connection. But as before, the system as a whole neither has, nor needs, a deeper rationale. Just as there is justification of belief, there is

justification of attitudes, evaluative judgments, and moral thought. But in both cases, there is a basis from which we must start, and a limit to how much agreement we can hope to achieve by giving reasons.

One of the striking facts about this picture is that it is already present in Hume. Here is a striking passage:

It appears evident that the ultimate ends of human actions can never, in any case, be accounted for by *reason*, but recommend themselves entirely to the sentiments and affections of mankind, without any dependance on the intellectual faculties. Ask a man *why he uses exercise*; he will answer, *because he desires to keep his health*. If you then enquire, *why he desires health*, he will readily reply, *because sickness is painful*. If you push your enquiries farther, and desire a reason *why he hates pain*, it is impossible he can ever give any. This is an ultimate end, and never referred to any other object.⁴⁰

One might ask whether health really is an ultimate end. It appears that there are reasons for desiring health, so this reply is not, as it stands, the point where all justification comes to an end. Hume admits this. He suggests that a man may reply that health is necessary for the exercise of his calling. So here we move into a region where we find a network, or a system, of sometimes mutually supportive, sometimes conflicting, considerations that define a practical point of view. Once we reach the realm of those commitments, reasons will once more tend to run out. Consider the man who desires health because he wants to be engaged in some profession:

If you ask, *why he is anxious on that head*, he will answer, *because he desires to get money*. If you demand *Why?* *It is the instrument of pleasure*, says he.

⁴⁰ D. Hume, *An Enquiry Concerning the Principles of Morals*, Appendix I, p. 293.

And beyond this it is an absurdity to ask for a reason. It is impossible there can be a progress *in infinitum*; and that one thing can always be a reason why another is desired. Something must be desirable on its own account, and because of its immediate accord or agreement with human sentiment and affection.⁴¹

The passage echoes Aristotle, who observes that

we do not choose everything because of something else, for if *that* is the case, the sequence will go on to infinity, making our desire empty and vain.⁴²

The difference is that for Aristotle, this line of thought points upwards. Following this line of thought, Aristotle sets out to identify the highest good for man, which is the good that embraces, or overarches, all others. For Hume, this kind of search is likely to be futile. For him, the line of thought points downwards, to a basic set of human dispositions. It is an empirical question, for Hume, and for us, to what extent our fundamental commitments can cohere and lead to virtue.⁴³

It has been said that this account from the *Enquiry* captures best what Hume had in mind when he argued that reason is, and always must be, subordinate to passions.⁴⁴ I think this is correct. It is, at any rate, a much more promising position than the outright denial that reason can

⁴¹ *Ibid.*

⁴² Aristotle, *Nicomachean Ethics*, bk. I, ch. 2, 1094 a 20-1.

⁴³ Bernard Williams has suggested that the Aristotelian hope for harmony is bound up with a view of human nature that is no longer available (cf. B. Williams, *Ethics and the Limits of Philosophy*, ch. 3, and 'Replies', p. 199).

⁴⁴ Cf. B. Stroud, *Hume*, p. 169. For an instructive treatment of the parallels and differences between Hume and Wittgenstein in the theoretical domain, cf. P. F. Strawson, *Skepticism and Naturalism*, ch. 1. Strawson himself applies the picture to moral attitudes; cf. P. F. Strawson, 'Freedom and Resentment', in: P. F. Strawson, *Freedom and Resentment and Other Essays*, essay 1, and *Skepticism and Naturalism*, ch. 2.

give rise to motivations, which remains associated with the slogan from the *Treatise* that “reason is, and ought only to be, the slave of the passions, and can never pretend to any other office than to serve and obey them”.⁴⁵

The truth is that all rational reflection has to start from somewhere, and that reason itself is not a place where rational justification can come to rest. When reasons run out, we are left with a number of fundamental beliefs and commitments. They form a framework for our theoretical and practical orientation. We act, and have no choice but to act on the basis of these commitments, just as we act, and have no choice but to act, on the basis of our beliefs. Neither of these frameworks can be given further rational foundations. But it does not follow from *that* that we cannot act with certainty and confidence, just as it does not follow that we cannot form beliefs and act on them with certainty and confidence.

It is characteristic of the rationalist temptation to underestimate the resources of belief and the resources of desire, no doubt as a result of overestimating both the need for rational foundations and their availability. If we resist that rationalist temptation, we can see that nothing over and above our frameworks of belief and attitudes is available, and nothing needed. The language game of giving reasons is not based on reasons. It is unpredictable, and neither reasonable nor unreasonable: “It stands there – like our life”.⁴⁶

⁴⁵ D. Hume, *A Treatise of Human Nature*, bk. II, pt. III, sec. III, p. 415.

⁴⁶ L. Wittgenstein, *On Certainty*, § 559.

6 Ethical Thought Without Foundations

What would be the consequence of this for ethics? So far, I have left it open whether we should think of the foundations of our outlook in the practical domain as attitudes that can be true or false. This is important, because in my view it is not at all clear *where* we pass from desires or some other attitude to the cognitive realm of belief. It is indeed an open question whether we can say, of every moral attitude, that it must be *either* a desire *or* belief. But one thing that is clear is that insofar as moral judgments can be true or false, they can be the object of belief, so the analogy with theoretical belief and its foundation is, if anything, still closer than it would be on a Humean account that holds on to a firm distinction between 'reason' and 'passion'.

Similarly, there is no suggestion that the picture of *On Certainty* commits us to a form of realism about value. While our framework of belief includes beliefs about a physical reality that exists independently of us, it also contains the resources to explain the difference between those elements of our outlook that can lay a claim to representing the world as it is independently of us, and those that cannot. The picture that emerges in *On Certainty* puts restrictions on our prospects of *establishing* that science represents the world 'as it is anyway'. But it does not compromise the content and coherence of that thought.

The basic thought that there are fundamental commitments, which may be beliefs or other types of attitudes, and that these commitments

form a framework or a system invites an application to the moral case. When it comes to moral matters, and to reasons for action in general, there are very many things that we believe in the absence of rational reassurance of any kind. We have no justification for them, and we would not be able to give them a justification should the need for that arise. These commitments form a framework, or a system, in which moral thought and argument takes place.

Once more, there will be a distinction between those parts of our outlook that are grounded and those that are not, and this may be no sharp distinction. All the same, we can distinguish between the foundations and the structure that is built on these foundations. But the metaphor must be interpreted with care:

I have arrived at the rock bottom of my convictions.

And of this foundation-wall one might almost say that it is carried by the whole house.⁴⁷

In the practical case, we find a number of values, concerns, ideals, projects and purposes that form the foundations of a practical outlook. But they are not independent of other kinds of judgments. In particular, they are not independent of other moral or evaluative judgments, and these can be of different kinds.

Consider moral thoughts like the following: 'If it is avoidable, causing pain is wrong', 'Autonomy is valuable' or 'One should not lie

⁴⁷ *Ibid.*, § 248.

without sufficient reason'. They already involve a sensitivity to context and call for judgment in their application. When is causing pain avoidable? Are there exceptions to that rule? What follows from the fact that autonomy is valuable? And how valuable is it? What exactly counts as lying, and what are the kinds of circumstance that would count as providing sufficient reason to lie? All the same, thoughts of this kind can have the role of practical foundations. There is no need for these thoughts to be general in form or content. 'This is something that I have to do' or 'That would be best' can, in suitable circumstances, have a foundational function. In that case, they stand fast, and set up rational constraints for action and for other attitudes.⁴⁸

This thought takes us back to themes touched on at the beginning, in particular, to the autonomy of moral language. There is no reason to expect that our evaluations will conform to any simple pattern, or that there is any underlying structure to be found that could provide the rationale for ethical thought. Ethical language is at the heart of the ancient city, with its maze of little streets and squares. There is a system, or a framework of concerns and projects, in which particular judgments about particular cases receive rational support from various sources.

If this is correct, it has important consequences for the aspirations of philosophy to put forward general theories of the ethical. We find a pluralism of value at the surface level of the ethical phenomena, and we

⁴⁸ This is a point where the picture has a certain affinity to the basic tenets of particularism, especially to the holism of reasons; cf. J. Dancy, *Ethics Without Principles*, esp. ch. 5.

have resources, some of them derived from Wittgenstein, to account for this complexity. Given the pluralism of ethical experience, and given the resources we have to describe it, we should resist attempts to find, or to construct, some supposedly more basic structure. I think it is rather unclear what it would mean to have something 'more basic' to rely on than the ethical judgments we make, and make with confidence, in the absence of a theory. Generality is, by itself, no guarantee for truth or even adequacy.

It is one consequence of this that we should not attempt to reduce value judgments to a single type or currency. There is a strong temptation in some quarters to think that a single good or value, such as welfare or good consequences, provides the foundation for all others. But welfare is only one consideration among others. Without a doubt it often is a most important one, and frequently the fact that welfare is likely to be significantly increased by following one course of action rather than another is decisive. But people act for other reasons, and this needs no apology. They care for style, for justice, for some cause close to their hearts. They are concerned about the good for some, but not, or not so much, for that of others. They have different sensibilities of the important, the outrageous, the embarrassing, the brave. They have certain talents, which they may want to pursue for no other reason than that they want to see them come to fruition. People find themselves inspired by various ideals, pursue a range of different projects, and admire people for their creativity - and so on and on. It would be a kind of cosmic accident if this

variety of considerations that provide practical guidance turned out to involve some rational order or system, and it is particularly paradoxical to think of the variety of things that people value as the ultimate source of all value, and at the same time to ignore the variety of value judgments that these people make, and deeply care about, in favour of a single type of judgment. Our concerns do not, and they need not, make rational sense at a deeper level. Their complexity must be respected, not explained.

Of course there is no guarantee that ethical considerations feature prominently in a given system. Consequently, there is no guarantee that ethical considerations feature prominently in the motivational system of a given agent. But it must be remembered that there is no guarantee that they feature prominently in that system *on any realistic view* of our psychologies. The point is that for those who respect ethical reasons and are willing to use ethical words as their own, ethical reasons and judgments do not, as a whole, need any further support. They are at the foundations of the system, and they guide our judgment as effectively, and with as much authority, as any other type of reason.

The foundations will include, in addition to a number of ethical thoughts, a substantial conception of reasons, including an understanding of ethical reasons. Of course what counts as a good reason can vary and change, and it is not something that could be fixed independently of our ethical convictions. All the same, there is significant agreement about what counts as a reason, what makes human lives go better, what would be desirable, what has to be avoided, and so on. This makes it possible to

have rational discussions about ethics. It is important to remember, especially when doing ethics, how far these agreements reach.

For all that, there will be limits to all rational discussions, and there will be a point beyond which no ethical justification will reach. At some point we will *have* to say: 'A good reason is one that looks *like this*', and if the disagreement then persists, that will be the end of argument. But it will be the end of argument on any view of rationality that has any hope of being true. That this is the situation is a plain fact of experience.

The upshot is that all justification, theoretical and practical, must come to an end somewhere. This includes the resources of justification themselves, for what we count as good reasons is not always backed up by further reasons. There is such a thing as giving a reason for something that is a reason, and justifying the evidence:

But giving reasons, justifying the evidence, comes to an end; – but the end is not that certain propositions immediately strike us as true, i.e. it is not a kind of *seeing* on our part; it is our *acting* which lies at the bottom of the language game.⁴⁹

The difficulty is to realize the *groundlessness* of our believing.⁵⁰

⁴⁹ L. Wittgenstein, *On Certainty*, § 204.

⁵⁰ *Ibid.*, § 166.

6

Living With Subjectivism

Why are we so naturally disposed to reject subjectivist accounts of ethics? No doubt the explanation of our intuitions in this area will be a most difficult affair, but one thing seems to be fairly clear. When we ask why we feel that subjectivism cannot possibly be correct, we tend to point to features of our ethical experience.

The previous chapters have given us a sense of the resources of subjectivist accounts of ethics. I have argued that there is good reason to resist all realist accounts of moral facts and moral reasons. Subjectivism has emerged as the more plausible alternative, both with respect to moral facts and moral reasons. As we have seen, subjectivism can draw on a wide range of resources to make its claims both more robust and more compelling.

In this chapter, I will put some of these resources to use in a subjectivist analysis of ethical experience. The features of ethical experience on which I will concentrate are those that seem, at least initially, to be the most resistant to subjectivist analysis. It has often been remarked that ethical thought naturally suggests a realist interpretation, and the thought that ethics rests on a foundation in objective fact is not a philosophical invention. It is based on distinct aspects of our ethical experience.

The complaint is not, perhaps, that subjectivism fails to make *any* sense of ethical experience. This is quite unlikely to be true, unless subjectivists are all completely blind to the phenomena, or choose to ignore them because they are already and irreversibly committed to their theory. Subjectivists have quite a lot to say about the facts of ethical experience. Subjectivism may turn out to be false, but it is certainly not *obviously* false. Ethical thought does not obviously *require* a realist interpretation.

Still, it may be that some form of moral realism makes better sense of the phenomena of ethical experience. The purpose of the present chapter is to argue that subjectivists can make good sense of these phenomena, and that a realist account of them is by no means compelling. Given that subjectivism is already established as the less demanding theory, this completes the case in favour of subjectivism. At the same time, it suggests a diagnosis of some of the salient philosophical mistakes that drive objectivist accounts of ethics.

1 Stability and Independence

Consider a first source of reasons for resistance to subjectivism. One striking feature of our ethical experience is the independence and stability of moral motivation. What I have in mind is a familiar fact of ethical experience. For most of us, most of the time, moral motivation is not only authoritative and important, insofar as we take moral demands quite seriously. Moral motivation is also reliable, effective and autonomous. It has a stable influence on thought and action, and needs no external motivational support. Moreover, moral motivations have a recognizable identity. On the one hand, they must be distinguished from narrowly egoistic motivations, but they must be distinguished from narrowly altruistic ones as well. At the same time, they are prominent and powerful. This seems to show that they have their source not in further motivations, but in something of an entirely different order. It is natural to think that moral motivations have foundations in belief and reason, rather than in things like dispositions, attitudes, psychological and cultural phenomena, or something essentially contingent such as forms of life.

One reply open to subjectivists is to explain away this seemingly compelling intuition. Thus, the most basic forms of expressivism would simply deny that moral motivation involves such things as moral facts, beliefs and reason. Reason, they may say, is restricted to finding means to given ends, or more generally, to weighing evidence, finding out how

things stand, what options there are, and assessing consequences. But there is no sense in which motivation, either in the moral or the non-moral case, is properly explained by reason. What initially appears to be an attitude with representational content and looks like the activity of reason is to be explained in terms of attitudes that are neither rational nor representational.

Sometimes, this will be the correct reply, and it is important that neither the realist nor the subjectivist has to deny that. But there is a different, and in my view more convincing, line of argument open to the subjectivist. It is that even from the subjectivist perspective, there is a sense in which it is quite simply true that moral motivation has its source in facts, beliefs, and reason. The subjectivist insists that moral judgments are not adequately thought of as objective. But as we saw, it does not follow that none of them are ever true. Consequently, there is no general objection to speaking of beliefs and knowledge in this area.

Moreover, there is no reason to think that reason is only peripherally involved in moral motivation. There are two sides to this argument that I will take up in turn.

On the one hand, we should not restrict reason to purely instrumental rationality. We have seen in the discussion of internalism that there is a wide, and to some extent indeterminate, variety of ways in which someone may rationally reach the conclusion that he has a reason to act. That is enough to free us from the misconception that subjectivists

have to rely on a reductive view of rationality, such as the view that reason only ever finds means to given ends. This is clearly not so.

However, there is a subtler version of the same objection. True, one may say, subjectivists do not rely on purely instrumental views of rationality. But they rely on something that may be called a purely formal, or procedural, account of rationality, and this account seems to leave out something important. On a purely formal or procedural account, what it is reasonable for someone to believe or do will depend entirely on the given beliefs and attitudes. In other words, for every p , both believing p and not believing p can in principle be reasonable. It all depends on the given framework of beliefs and attitudes. Similarly, for every action A , both doing A and not doing A can be reasonable. It all depends on the given framework of desires, attitudes and practical commitments. This seems to leave no room for the idea that a certain belief or attitude may be reasonable or unreasonable *in its own right*. In other words, it leaves no room for a substantive view of rationality.

It is easy to see why it looks as though subjectivists endorse the formal, procedural picture. After all, subjectivists stress that justification is always, and essentially, relative to substantial preconceptions, and that these form the foundations for all reasoning about what would be good, or right, or sensible to do. Internalism seems to be directly committed to this view. After all, internalists claim that what a person has reason to do is relative to her present desires, attitudes and practical commitments. This

seems to be simply another way of saying that a formal or procedural account of rationality is correct.

This is not so. That there is something wrong with this description can be seen from the fact that internalists do not claim that all reasons for belief are given by beliefs. They say that in the standard case, reasons are given by the facts. That is why Williams is right to insist that someone who mistakenly believes that there is gin in the glass in front of him when in fact it contains petrol *does not have a reason* to drink what is in the glass. He merely *thinks* he has a reason, but this belief is false.¹ If so, purely procedural accounts of rationality are already ruled out.

Still, this does not yet take us all the way to a substantial view of rationality. This is as it should be, because we want to allow that formal, or procedural accounts of rationality capture something true about our use of 'rational' and 'rationality'. There is a truth in this account, and we can understand how it can help us focus, for example, charges of irrationality. It is at least a part of being rational that one is not involved in inconsistency or incoherence. But it is important to remember that there is a further use of the expressions 'rational' and 'rationality', and that these uses normally involve further commitments. This kind of use is captured well by Wittgenstein, when he says 'the reasonable man believes *this*'.² It is, for us, part of being rational to believe that only a few people have ever been far from the surface of the earth, that cars do not grow out of the

¹ B. Williams, 'Internal and External Reasons', pp. 102-3; cf. B. Williams, 'Internal Reasons and the Obscurity of Blame', p. 36.

² L. Wittgenstein, *On Certainty*, § 323.

earth, or that every human being has parents. It is, for us, part of being rational to think that wanton cruelty is wrong, that other things being equal, causing pain is bad, or that every human being deserves equal consideration. Everything else is not merely an unorthodox view. We find it unreasonable.³ But the subjectivist does not dispute that.

It is, of course, important to remember that charges of irrationality are strong, and that not every kind of criticism of an action or an attitude is adequately put by using those expressions. In a moral dilemma, we will hesitate to call the person opting for one choice unreasonable. There are strong reasons on both sides of the dilemma, and we will certainly not think of a person who chooses one option as unreasonable merely because we would have opted for the other. This is different in the case of wanton cruelty. If someone does not recognize the fact that an action would cause harm or pain even as a reason, or sees nothing wrong with making other people suffer, we would feel ourselves very alien from that person. Here, there is something more at issue than mere disagreement. The disagreement is one about fundamental values, and the basic principles of conduct. Here, the charge of violating a requirement of reason may well be in order.

The mistake is not to think that there is a substantial notion of 'reason', 'rational' or 'rationality' that could support those charges. If there is a mistake here, it is the mistake of thinking that substantial notions of

³ On this point, cf. H. Frankfurt, 'Rationality and the Unthinkable', in: H. Frankfurt, *The Importance of What We Care About*, essay 13.

reason and rationality introduce us to a realm of objectivity, where our practical commitments are backed up by some normative order that is simply there. There is no such order, but it does not follow that we can call an act of wanton cruelty wrong, but not unreasonable. It is true that ethical considerations are already written into our idea of reason and rationality. But that does not mean that they are objectively binding. If they are, that remains to be shown.

So there is room for a substantial notion of 'reason' and 'rationality' on the subjectivist account, and these notions will inform subjectivist accounts of moral motivation. A second aspect of the argument is that on the subjectivist account, 'reason' is something much more complex and more opaque than it at first appears. Nietzsche said of the will that it is "above all something *complicated*, something that is a unity only as a word".⁴ The word, he says, disguises "a plurality of feelings", and Nietzsche claims it is a common prejudice to take what is a mere appearance on the linguistic surface for a glimpse of the reality that lies beneath.

The same is true of 'reason'. Indeed, Nietzsche himself points out that in addition to a range of feelings, thought is an essential element, or as Nietzsche puts it, an "ingredient" of willing, and he issues sensible warnings against those forms of reductionism that tend to conflate willing and feeling. Thus,

⁴ F. Nietzsche, *Beyond Good and Evil*, § 19, p. 48.

in every act of will there is a commanding thought; - and one should definitely not think that this thought can be separated from the 'willing', as if the will would then remain!"⁵

What we refer to as the will *essentially* involves processes of thought, and not just feeling, and this is one way in which the appearance on the linguistic surface is deceptive when it represents as one what is in fact a complicated process.

Reason is a competence, a competence, moreover, that involves a range of other forms of competence. In particular, it involves a competence of giving reasons, which are sometimes, but not always, such as to license inductive or deductive inferences from one proposition or set of propositions to another. Practical reason involves the ability of giving reasons that pertain to action, or to practical affairs more generally. In its moral manifestation, practical reason involves the ability of giving moral reasons, which cover a wide range of relations between propositions, attitudes and actions that are, taken together, constitutive of moral thought and argument.

The point on which subjectivists insist is that abilities like that are rooted in our linguistic practices, and that our linguistic practices explain the competence in question. At the social level, reason is a practice of exchanging reasons. At the individual level, reason is a form of linguistic competence that involves a complex set of dispositions. The basic disposition is, as the subjectivist insists, a disposition to take some

⁵ F. Nietzsche, *Beyond Good and Evil*, § 19.

consideration *as a reason*, and we should not think of that as an ability that needs further objective groundings.

So the subjectivist need not deny that fact, belief and reason are involved in moral motivation, and this makes the subjectivist reply much more compelling from a realist perspective than it might otherwise appear. It is simply not true that there is no such thing as a moral belief, or that a moral argument has no basis in fact or reason. Equally, it would be a mistake to think that reason must be silent when it comes to thinking about ends, as opposed to thinking about means to given ends, where ends are supposedly set by 'passion'. The question the subjectivist will ask is how far that thought, by itself, can take us, and the answer is 'not very far'. It needs to be shown, and cannot simply be taken for granted, that the ideas of fact, belief and reason commit us to more in the way of ethical objectivity than the subjectivist allows.

2 Making Sense of Moral Motivation

What does this show for the analysis of moral motivation? In particular, what does it show for the analysis of its stability and independence? In my view, it shows that subjectivism stays closely in touch with the phenomena, and that it has no need to take recourse to 'passions' or some supposedly more basic psychological phenomena that are alien and fleeting. This would be a misleading characterization even on the terms of the expressivist account, since expressivists will rightly point out that our

basic attitudes are anything but alien and fleeting. This is a characterization of the psychological that owes much more to Kant and Plato and those inspired by them than to those who held that ethics is at bottom a psychological phenomenon.⁶

However, I think that a stronger claim can be established. Whatever objections may be raised on the grounds that there remains a mismatch between what needs explaining, namely moral motivation, and what does the explaining, namely attitudes and dispositions, can be discarded once we get a clear view of the attitudes and dispositions on which moral motivation rests. On the subjectivist view, there is no need to deny that moral attitudes frequently are what they appear to be, namely beliefs and sometimes even knowledge. There is no need to explain appearances away. When realists insist that ethics rests on fact, belief and reason, subjectivists should not try to explain what happens differently. They should simply agree.

This suggests a way of dealing with a seemingly compelling intuition that seems to show that subjectivism must be false. It is sometimes claimed that the subjectivist lacks the resources to make sense of the difference between thinking something valuable, or important, and merely desiring it. This is sometimes expressed by saying that for the subjectivist, 'it is all a matter of mere preference', or that for subjectivist, 'all value ultimately rests on nothing but desire'. This seems to leave out,

⁶ Here it is important to remember Hume, in particular his discussion of the 'calm passions'; cf. D. Hume, *A Treatise of Human Nature*, bk. II, pt. III, sec. III, pp. 417 f.

and does leave out, something important. After all, there is a difference between thinking something valuable and merely desiring it. The two expressions have entirely different uses. The question is what that difference consists in. Realists will tend to say that we must appeal to substantial objectivity to make sense of value, whereas a subjectivist might say that there is a difference between valuing and desiring, but that we can connect them by intermediate cases.⁷ How can we account for the difference between mere desires on the one hand, and deeper evaluative commitments on the other, and still be subjectivists?

One line of response to this question has been explored by Harry Frankfurt. In a series of writings that begins with 'Freedom of the Will and the Concept of a Person', Frankfurt has developed an account that relies on the idea that we have desires, and that these include desires to desire. Frankfurt suggests, in other words, that we have desires of a first and of a second order.⁸

The difference can be illustrated by example. A drug addict desires the drug to which he is addicted, but it does not follow that he desires to desire the drug. He may have a strong, but ineffective desire to break out of his addiction, so that he wants, but does not manage to, overcome his

⁷ Cf. L. Wittgenstein, 'Lectures on Aesthetics', p. 12: "What similarity has my admiring this person with my eating vanilla ice cream and liking it?" To compare them seems almost disgusting. (But you can connect them by intermediate cases)".

⁸ Cf. H. Frankfurt, 'Freedom of the Will and the Concept of a Person', in: H. Frankfurt, *The Importance of What We Care About*, essay 2. The theme is further developed in a number of later writings, including 'The Importance of What We Care About', 'Identification and Wholeheartedness', essays 7 and 12. The discussion continues with 'The Faintest Passion', 'Autonomy, Necessity, and Love', and 'On Caring', in: H. Frankfurt, *Necessity, Volition, and Love*, essays 8, 11 and 14. The most recent statement of his view is found in H. Frankfurt, *The Reasons of Love*.

strong desire for the drug. If so, Frankfurt suggests, the desire for the drug may not really be his. He may, on the other hand, also be a willing addict. In that case, he has a strong desire for the drug, and moreover, a desire to have that desire. He is identified with being an addict, and in that sense, he is a willing addict. Frankfurt thinks that the unwilling addict lacks not freedom of action but freedom of will, while the willing addict is free in both senses, and so liable to blame.

A criticism that suggests itself is that evaluations are not adequately captured by this model.⁹ The objection, pressed by Watson, is that the hierarchical account faces a dilemma. Either it merely ends by *fiat* what would otherwise be an interminable ascent to yet higher orders, in which case the phenomenon of identification with desires or a course of action would be lost. If this ascent is not to be cut off arbitrarily, it must come to rest with some stable commitment. Moreover, it must come to rest with something that shows how, and why, the agent is really identified with what he is doing. But if that is so, it is no longer obvious that the notion of a second order desire is what really does the work. If the notion of identification is not to be empty, it must be explained independently of the hierarchical account.¹⁰

I agree with these criticisms to the extent to which they are focused on the notion of a second order desire. I think that identification and evaluation are not adequately captured by the hierarchical account as it

⁹ Cf. G. Watson, 'Free Agency', in: G. Watson, *Agency and Answerability*, essay 1.

¹⁰ Cf. G. Watson, 'Free Agency', sec. III.

stands. It may of course be true that what a person values is naturally expressed in terms of what the person desires to desire, rather than by just considering what she desires. But this does not show that the second order of desire does the work.

In fact, desires of a second order seem to be peripheral. Deliberation and evaluation are typically first order activities. Those who care for justice or their family will say, very obviously, that they care for justice or their family. This is the most direct expression of their practical commitment. This practical commitment will, at least up to a point, be backed up by reasons. So far, then, Watson is right to emphasize that we respond to reasons, and that we respond to what strikes us as important, or desirable, and this leaves it quite open what to say about the underlying structure of our motivational psychologies. But Watson tends to back up this point about deliberation by some more demanding, and explicitly Platonic claims about practical reason. Reason, Watson believes, must be thought of as a distinctive source of motivation and insight. Otherwise, the relevant distinction between valuing something and desiring it cannot be maintained.

But why should we accept that? Wittgenstein can help us see that there is no need to buttress a structural distinction between more or less foundational beliefs with some explanatory structure, for example by demanding that the propositions that are located at the foundations must be insights into some sort of objective realm of reason. If someone were to ask, 'But how can you believe anything if all you have to go on is belief?',

the answer is that we must not be misled by a false picture of the nature of belief.

We are inclined to think of belief as something flexible, something that stands in need of further reasons, as open to doubt and to dispute. If that were the situation, it would not be possible for us to form a picture of the world and to be satisfied with it. This would be a picture not of certainty, but of the absence of all certainty. But not all belief is like that. There are beliefs that we would not give up for anything, and this shows that these beliefs have a special status in the system. But for all that, it is a system of *belief*. If we recognize that fact, we can see why nothing more is needed, whether or not something more can be had:

The question is really this: "What if you had to change your opinion even about these most fundamental things?" And to that the answer seems to me to be: "You *need* not change it. That is just what their being 'fundamental' is."¹¹

The same is true in the sphere of the practical. Some desires are too flimsy, or inconstant, or too vague to form the basis for evaluations. Such desires could not form the basis of what we would call a moral outlook. But to infer that therefore, something more than mere desire, or evaluative attitude, or belief, is required is to make the same mistake that we met earlier. There are desires and evaluative commitments and beliefs that are so fundamental that they form the basis for evaluations. In their entirety, they constitute our practical identity, and define our point of view.

¹¹ L. Wittgenstein, *On Certainty*, § 512.

Given an account of moral motivation in terms of such basic dispositions and specific kinds of linguistic competence, it is no longer clear why the subjectivist should be thought to be unable to account for the independence and stability of moral motivation. Successful moral education leaves those who receive it in no doubt that moral reasons are important, although they are not necessarily, and perhaps not even generally, our most immediate and most important reasons. But it is a fact that moral reasons often need no further rational support, either from self-interest, some other desire, or altruistic sentiment. Moral thinking becomes second nature. It shapes our practical identities and outlooks, and moral thought that is embedded like this in our practical identities needs no further foundation in objective fact.

The root of the mistake lies in the idea that a subjectivist can only think of moral thought as 'just another type of motivation'. This way of describing it suggests two different kinds of misunderstanding. First, it makes it appear that moral motivation is not rational, but substantially impulsive, or irrational, as opposed to being due to insight or reflection. As a result, the state of moral motivation seems to be cut off from the intellect, which is not how we think of moral thought and motivation. This, however, is not what the subjectivist is saying. What exactly the contrast between rational and a-rational features of our motivational states amounts to is itself unclear, and whatever that contrast may amount to, the subjectivist has ways of interpreting the dispositions on which he relies so as to avoid locating moral motivation on one side of the divide. In

fact, the subjectivist is well advised to leave behind all simple oppositions between 'reason' and 'passion', which will not do justice to the structural complexities of our psychological condition.

A second mistake that naturally suggests itself when we describe moral thought as 'just another kind of motivation' is that moral motivations are ones we could as well do without, because they have no deeper significance for our lives and our sense of our identities. In this way, moral motivation tends to be associated with the most transient of desires, or to those that are experienced as a hindrance, as causing unbidden distractions, or as merely befalling us. Once again, this would be deeply inappropriate as an analysis of moral motivation. But what is inappropriate about it is not so much the fact that moral motivation is explained in terms of dispositions and desires, but the way in which the explanation is spelled out.

To think of moral motivation as on a par with the most transient of desires is to think that we could simply turn our backs on it, because we do not really care about morality and the aims that it promotes. This is certainly not so, and subjectivism can explain the reasons why. Morality is something that we care about, something that matters. It is not the only thing we care about, and there may be things about which we care more. But it is among the things we care about, and it has a prominent place among them.¹² Morality is something with which we identify deeply,

¹² For a sensitive discussion of the phenomena in question, cf. H. Frankfurt, *Necessity, Volition, and Love*, and *The Reasons of Love*.

something that is, now, part of our nature, if only contingently, not essentially, and not always safely. It is not one set of motivations among many, if that entails a judgment of importance, or suggests that we would not experience losing moral motivations as a loss.

Indeed, there is a question as to what standpoint we occupy when we experience some desires as occurrences, as hindrances, distractions, or as alien forces. It would seem that there has to be some operative contrast. Occurrences must occur to *someone*. Hindrances must stand in the way of *something*. Equally, distractions must distract *someone* from *something* that would be more worthwhile. Alien forces can only be alien if they are not part of our practical identities. If so, they must be incompatible with those things that we really care about, and find truly important.

If we want to understand the contrast, then, we need a clear view of the relevant conception of the self. Put quite schematically, desires that are alien, come unbidden or distract us, are those that are incompatible with, or at any rate opposed to something that we care about, find worthwhile or important, or see good reason overall to do. To use a textbook example, the desire to watch television or to go for a walk in the woods can be like that when we think that we really should be doing something else, such as cleaning up or marking papers. Complicated cases call for complicated comment, but I take it that the disagreement with the moral realist does not concern this opposition and its basic structure. Rather, what the disagreement is about is how to understand this opposition, and here the

subjectivist has something more to say than that we find a simple conflict of desires, struggling for predominance over our thoughts and our wills.

First, it is not true that what we think we have good reason overall to do, or care about, or find worthwhile or important, is just a reflection of our momentary motivations. We have different words to mark the differences between a mere desire on the one hand, and caring, thinking important, seeing good reason to do, and so on, on the other. These differences in words mark genuine differences in the phenomena, and it would be a crude form of reductionism to try to account for all these different kinds of state of motivation in terms of a theory that likens the range of motivations to one type, that of a desire or a merely momentary motivation.

Nevertheless, on the subjectivist view even what we think we have reason to do, or think of as worthwhile and important, is based on or expresses our dispositions, and in that sense, there is certainly a truth in the idea that when there is a conflict of what we see good reason overall to do and actual desire, this conflict is one between different dispositions. But if the subjectivist is right about that, this is not to his disadvantage. First, the conflict between what we see good reason to do and what we desire *is* a conflict of our dispositions even on the most radically rationalist account. After all, on that account, reason is a source of motivation. There is a conflict of subjective motivations, a conflict that is felt and often difficult to settle, and this is true on *any* view of our psychology that has a chance of being true. All the subjectivist says is that reason itself is to be

explained in terms of dispositions, not in terms of some mysterious objective insight, and in this sense, subjectivism provides the more powerful explanatory theory.

Rationalists tend to resist that view. In response, they may point out that what we see good reason to do or find important is typically not experienced as being based on a disposition. True enough, but why should it be any different? Think of the idea that our experience of free choices can reassure us of the freedom of the will. The argument that where we experience no causes there are no causes is obviously invalid. As Schopenhauer grimly points out, we cannot translate a lack of knowledge of the causes of our choices into knowledge that these choices have no causes. One might even think of Kant:

Who can prove by experience that a cause is not present? Experience shows only that it is not perceived.¹³

Indeed, Kant himself thinks that this might serve as a reminder that all moral imperatives might be “covertly hypothetical”.¹⁴

3 The Ring of Gyges

A different way of bringing out the difference between moral motivation and mere transient desire is to ask what would happen if we knew that we could pursue our interest narrowly conceived, and ignore demands of

¹³ I. Kant, *Groundwork for the Metaphysics of Morals*, pt. II, AA 4:419.

¹⁴ *Ibid.*

ethics without facing any blame or sanction.¹⁵ In the *Republic*, Plato lets one of his characters describe such a situation. Glaucon mentions the legendary ring of Gyges, which is said to have the power of making its bearer invisible:

Suppose now that there were two such magic rings, and the just put on one of them and the unjust the other; no man can be imagined to be of such an iron nature that he would stand fast in justice. No man would keep his hands off what was not his own when he could safely take what he liked out of the market, or go into houses and lie with any one at his pleasure, or kill or release from prison whom he would, and in all respects be like a God among men.¹⁶

The picture that emerges is one of an egoistic underlying nature, which is held in check by other selfish motivations, in particular, the desire to avoid the social sanctions that are the normal outcome of transgressing what society demands. Men are at bottom selfish, and merely 'diverted into the path of justice by the force of law'. Justice is better than a state of powerlessness, but considerably inferior to the most desirable state of all, which is that of being able to pursue what one wants without having to suffer for it. All justice can ever be is second best.

This challenge is quite serious, and Plato is occupied with answering it for the rest of the *Republic*. For him, an adequate reply to it requires that it can be shown that for any man in any circumstances, it is rational to stick to justice. It is very probable that this challenge cannot be met. But it is also very probable that it need not be met to answer the

¹⁵ Cf. B. Williams, *Morality*, pp. 21-3.

¹⁶ Plato, *Republic*, 360 b-c, p. 201.

challenge as described by Glaucon. It is simply not true that all people are at bottom selfish, or that they would no longer respect moral considerations if they could be sure of getting away with it. Moral dispositions are rooted more stably and more deeply in our characters than that, and they will not lose their influence on us just because social sanctions no longer threaten. No doubt we often tend to underestimate the power and predominance of egoistic motivation. But if this is something we do not wish to admit, least of all to ourselves, there must be more powerful patterns of motivation that will not let us admit it. These must be moral motivations, and the fact that they make themselves known to us in this way shows that self-interest is neither the predominant nor the most authoritative source of motivation.

Moreover, this reply will still respect the plain truth that self-interest is part of the rationale of morality. Institutions such as justice generally benefit those who participate in it, and if it were not so, they would not exist. As Hume said, instead of 'departing' from our own interest or from that of our nearest friends by abstaining from the possessions of others, "we cannot better consult both these interests, than by such a convention".¹⁷ Self-interest remains a powerful incentive to ignore demands of justice, but for most of us, most of the time, it is not powerful enough.

It may still be asked to what extent self-interest is the force that drives compliance with morality, and it is this kind of thought that made

¹⁷ D. Hume, *A Treatise of Human Nature*, bk. III, pt. II, sec. II, p. 489.

Kant wonder whether there is 'any genuine virtue to be encountered in the world':

Out of love for humanity I am willing to grant that most of our actions are in accord with duty; but if we look more closely at our scheming and striving, we everywhere come across the dear self.¹⁸

Kant is being realistic, and if there is something in his view of human motivation that is not realistic, it is his requirement of what would have to be the case for there to be 'genuine virtue'. To be sure, there is more selfishness involved in our scheming and striving than we like to admit, and there is such a thing as moral indifference. But there is no doubt that most people care about their fellows, that there is a basic form of altruistic motivation, and that people often go to great lengths to do what they think they should do, where what they think they should do involves some significant personal loss. Moreover, many people take the view that this is precisely what gives their actions 'moral worth'. They prefer to pull the oars, and do not do so reluctantly. They are motivated in accordance with morality even when there is an opportunity to throw all caution to the wind and get away with acting selfishly. Most people will indeed think better of themselves for that very reason. For them, morality has become second nature, and for it to have a stable anchor in human motivation, it does not have to be more than that.

– 'But is it *rational* to act like that?'

¹⁸ I. Kant, *Groundwork for the Metaphysics of Morals*, pt. II, AA 4:407.

I think that we should not be unduly worried by this kind of question. After all, what reason is there to leave the framework of ethical thought behind? If the answer is that we have more reason to pursue what benefits us, this is simply not true on the view of reasons that I am suggesting. Moral reasons are among our most important reasons. There is no sense in which the pursuit of those things that benefit me, as opposed to those that benefit you or someone else, enjoys a privileged status as a source of reasons.

This reply is likely to attract further criticism. After all, on the view I am suggesting, all our reasons are based on our dispositions, in particular, on our dispositions to take certain considerations as reasons. But these dispositions are, as one might put it, merely *our* dispositions. Reasons, including moral and prudential reasons, are based on a language game of giving reasons, and that language game is in a sense distinctly ours.

Now the language game of giving reasons is certainly not distinctly ours in the sense that other beings *could* not share it, or *could* not have reasons similar to ours. But it is distinctly ours in this sense, that the language game of giving reasons is subject to local and historical and cultural variation, and moreover, in that it makes no sense to ask what reasons there *really* are, independently of our linguistic practice. How can such a structure provide guidance and stay stable on reflection when it lacks all grounding in objective fact?

In reply, I think we need to be clear what difference it would make if some objectivist account of ethics or of reasons or of both were true. If it were true that we have some objective reason to be moral, why should we care about that truth? Similarly, if it could be shown that morality were a requirement of pure practical reason, even a demand grounded in a demand for rational consistency, why should we be impressed by that?

The point is not to deny that dispositions to have true beliefs and to be rational are deep and valued dispositions. This they evidently are. Most of us, most of the time, take a keen interest in knowing the truth, as well as in being consistent and fully informed and rational. The point is, first, that it is doubtful that these disposition could, all by themselves, carry the weight of our moral commitments, and second, that it is unclear why those commitments need that kind of rational support.¹⁹ Clearly, ethical upbringing has, and has to have, the significance subjectivists ascribe to it even on realist accounts. If moral reasons are objective, it remains to be established that they are, and if that is so, our dispositions must take care of themselves in the absence of such theoretical assurance. Ethical dispositions must be second nature if moral life is to continue in the form in which we now have it, and this must be so regardless of the question whether realist accounts of ethics can be sustained.

The truth is that our moral dispositions are stable and enjoy a high deliberative priority because they are concerned with most important

¹⁹ Cf. R. Nozick, *Philosophical Explanations*, p. 408; B. Williams, *Ethics and the Limits of Philosophy*, ch. 2.

matters. As Socrates said, it is not a trivial question with which we are concerned – the question is how one should live.²⁰ Ethical commitments are part of our practical identities. Losing our ethical commitments would be tantamount to losing part of ourselves, part of who we are and who we want to be.

Given the undoubted importance of the ethical, what can be the contribution of our disposition to believe what is true, and to be rational? The realist may say that ethical truths or truths about reasons are especially important truths. But if they are, they are important because they are truths about our reasons and the ethical, not, or not primarily, because they are truths. If our dispositions to act ethically and to respond to reasons were not stable in themselves, our commitment to believing the truth or to doing what is rational would make very little difference. Given that they are stable in themselves, it is unclear why anything beyond a set of stable dispositions is required to sustain ethical life.

One might be tempted to conclude that ethical attitudes are neither true nor rational, because the criticism leveled at someone who acts extremely badly obviously goes beyond the criticism we associate with purely cognitive mistakes. This is true, but it does not show that ethical attitudes can be neither true nor rational. What it does show, by contrast, is that the sense of importance we feel is associated with the ethical comes from the subject matter of the ethical, not from the supposed requirement

²⁰ Cf. Plato, *The Republic*, Bk. I; for discussion, cf. B. Williams, *Ethics and the Limits of Philosophy*, ch. 1, and R. Crisp (ed.), *How Should One Live?*

of truth or rationality. Ethical life continues as long as we have a stable, firm and independent set of ethical commitments, and from a point of view defined by those commitments, that is certainly enough.

4 Deliberation and Experiences of Externality

There is certainly a sense in which ethical experience naturally supports a realist interpretation because it seems to present us with an ethical reality. Ethical thought, especially deliberative thought about what would be good to do in the given situation, seems to put us in touch with normative dimensions of reality, something that is there independently of us.

I suspect that this interpretation most naturally suggests itself in cases where we are in no doubt about the right assessment, so that no reviewing of our reasons or considerations that pertain to other facts is likely to affect our judgment of the situation. Suppose you are presented with an act of deliberate cruelty, where there is no question that such cruelty, in such circumstances, is not morally defensible. Say you witness an attack on a small child on the street. A bully hits and kicks the child merely because he feels like doing it, or to reassert his power. There is no question that the act is cruel, and there is no question that such cruelty might in fact be morally defensible, so there is no question that the act is wrong. Similarly, say you realize that you must keep your word despite the fact that you would like to be released from it. You are aware of your

disinclination, but conclude that for you, here and now, keeping your word is what you should do.

Here, it seems that in deliberating about what to do and arriving at that conclusion, you are presented with a truth, something calling on you from the outside, something that is external to you in the sense that it does not present itself as flowing from your beliefs and attitudes. It is natural to describe judgments of this kind as insights, and to describe them in this way is to bring out their epistemic aspirations and their claim to objectivity.

It is perhaps no accident that this description fits most easily in cases where comparatively little thought and reasoning is called for. Your assessment of the situation simply arrives, and it arrives where you find it quite easy to arrive at firm moral conclusions.

What is happening here, I suggest, is that we are faced with paradigmatic cases in which ethical concepts apply. To arrive at the conclusion that something would be a good or bad or right or wrong thing to do in a given situation can require sensitivity and judgment, and in many cases this assessment must be tentative, open to reassessment and revision, or remain surrounded by substantial doubt. In some cases, however, moral judgment is not like that. Deliberate cruelty is one example, but keeping your word is too, even though it may of course sometimes be right to break it.

The fact is that these cases present paradigmatic examples of ethical instruction. As such, they present paradigmatic cases where ethical

concepts apply. The result of such instruction is that it comes naturally to us to use these concepts and to arrive at these conclusions, and it is no surprise that once we have acquired these concepts and use them as ours, we feel that there is no choice but to apply them. Given that our use of ethical words has a motivational dimension, it is no surprise that following this recognition, we act as we do.

Moreover, when moral concepts apply, this need not be an immediate reflection of present desires or evaluative attitudes. Rather, on the subjectivist view as I conceive of it, we should think of our abilities to form moral judgments as we think of our abilities of forming other kinds of judgment, that is, as an exercise of a distinct conceptual and rational capacity. It is not obviously the same to say that recognizing that something is right or wrong reflects a complex linguistic disposition to form moral judgments, and to say that all moral judgments have one and the same distinctive task and character, namely to express an attitude or disposition. The same is true of other notions that are relevant to reaching practical conclusions, such as those applied when we judge what is good or prudent or desirable.

It may be objected that this kind of account reintroduces the same externality, only at a different level. For now, we have explained the thought that in thinking about right and wrong and related notions, we are guided by rules that determine when to use these words. Thus, it looks as if we have introduced the normative back into the account, and that the

normativity that comes with rules of concept application will cause the same problems, only at a different level.

But will it? It is clear that in one sense at least, the account will have to remain normative. Rules and normative behaviour that surrounds them cannot be reduced to something else that is no longer normative. In this sense, there is no intention to get rid of normativity, or to explain away the thought that our ability to speak a language is an ability to follow rules. With the concept of a rule comes the possibility of getting it right or wrong, and in that sense, rules and normative behaviour are basic.

So the idea must be that if we accept this basic notion of a rule, and accept the idea that there is basic normative behaviour, then we are already committed to the kind of externality that ethical experience seems to present. But this is not necessarily so. To the extent that there are genuine prospects here for explanation, there are explanations of rule following and applying concepts that do not involve a reference to some normative objective order. We can stick to what we know, which is that we use ethical words in certain ways, and that such practices of using words are governed by rules that we internalize and follow.

5 What Is a Free Exercise of Reason?

A further point frequently called on by supporters of a realist account is that deliberation does not seem to be concerned with items in our psychologies, but is directed at the world. To the obvious objection that

practical thought may be directed at the world but may all the same be driven by desires, it is then replied that we do not experience the activity of practical deliberation in this way. Thomas Nagel gives this line of thought articulate expression, when he describes what he calls 'the standpoint of decision':

This standpoint introduces a subtle but profound gap between desire and action, into which the free exercise of reason enters. It forces us to the idea of the difference between doing the right thing and doing the wrong thing (here, without any specifically ethical meaning as yet) - given our total situation, *including* our desires. Once I see myself as the subject of certain desires, as well as the occupant of an objective situation, I still have to decide what to do, and that will include deciding what justificatory weight to give to those desires.²¹

This 'step back', Nagel believes, opens a slight space between inclination and decision, which is a condition that permits the operation of reason with respect to belief as well as with respect to action. It is only when, instead of 'simply being pushed along' by impressions, memories, impulses, desires, or whatever, one stops to ask 'What should I do?' or 'What should I believe?' that reasoning becomes possible - and, having become possible, becomes necessary, as Nagel puts it. This fact is then construed as evidence for the idea that reasoning transcends the point of view of the first person:

One is suddenly in the position of judging what one ought to do, against the background of all one's desires and beliefs, in a way that does not merely flow from those desires and beliefs but *operates* on them - by an

²¹ T. Nagel, *The Last Word*, p. 109.

assessment that should enable anyone else also to see what is the right thing for you to do against that background.²²

It is difficult to deny that this kind of description has a ring of truth to it. But it is also seriously misleading, and I think that it involves a serious misrepresentation of the facts of ethical experience. Subjectivism can preserve the truth in the account without being subject to the charge of misrepresentation. It makes better sense of ethical experience.

On the most primitive view, practical thinking directly reflects the current state of motivation, where the current state of motivation is thought of as constituted by the current set of rationally unmotivated dispositions, and nothing more. Such a view might be ascribed to Hume, but as we have seen, his view is much more subtle. Hume allows, indeed insists, that reasoning about the means to given ends is properly regarded as an exercise of reason, so as Hume himself conceives of it, reason certainly plays an important role in the aetiology of action.²³ Still, Hume seems to think that reason is incapable of generating motivations. But as we saw earlier, this claim is quite baseless.

A much more interesting idea, also found in Hume, is that there is no clear content to the idea that moral facts or truths are something 'in the world', as opposed to something added or imposed by us, or in some other way due to our subjective standpoint. But as we have also seen, it does not follow from this categorical denial of objectivity in ethics that

²² *Ibid.*, p. 110.

²³ Cf. B. Williams, 'Internal and External Reasons', pp. 101-103.

there is no truth or falsity in the ethical domain. So it does not follow that subjective styles of explanation must categorically deny that ethical thought is frequently just what it seems to be, namely capable of being either true or false. But if it is, then by Hume's own account it can be thought of as an exercise of reason. Reason is conceived of as concerned with truth and falsity, so if there can be truth and falsity in ethics, ethics has its distinctive place in the domain of reason.

What follows from this for the standpoint of deliberation? First, and most obviously, that the subjectivist has no need to accept the claim ascribed to him by many realists, including Nagel, that moral thought or evaluative thinking more generally is nothing more than the expression of desires or some other kind of rationally unmotivated motivation. There is certainly a sense in which it is quite true that when we occupy the standpoint of deliberation, we occupy a standpoint that takes present desires and features of the current situation into account. Moreover, it is true that in occupying it and reaching a practical conclusion, we are not directly driven by those motivations. Rather, once the question arises what there is reason to do, or even what our desires might incline or drive us to do, this question calls for an assessment of our reasons. Reasons, in turn, are something others can share in modes of rational participation, and that means that thoughts about reasons cannot simply be reduced to thoughts about the motivations we presently have. The question to be answered is not what we most desire but what we should do or have most reason to

do, and this question is not relevantly answered by invoking either the original desire or some other kind of rationally unmotivated motivation.

The point at issue is that so far, this leaves it entirely open what to say about reason itself. Nagel seems to think of it as an epistemic phenomenon, so that reason puts us in touch with a normative dimension of reality. But as I have argued, we need not accept that. Reason is better thought of as the ability to critically reflect on reasons, and that ability can be explained in terms of a disposition of taking certain considerations as reasons. The practice in which these dispositions have their place is internalized, and it is one in which others can participate, so there is no obstacle to thinking of reasons as public.

It may be objected that on this view of practical reason, psychological phenomena effectively replace the normativity that comes with the notion of a reason. Nagel, for example, claims that the subjectivist would have to show that all purportedly rational judgments about what people have reason to do are 'really' expressions of rationally unmotivated desires or dispositions of the person making the judgment, and he thinks that these would have to be desires or dispositions to which normative assessment has no application. Thus, Nagel claims, the motivational explanation "would have to have the effect of *displacing* the normative one – showing it to be superficial and deceptive".²⁴

This is patently untrue. It is true that normative assessment is a psychological phenomenon, as it must be if it is to feature as an item in

²⁴ T. Nagel, *The Last Word*, p. 111.

moral psychology. Moreover, it is true that normative assessment goes significantly beyond the mere expression of desires, impulses or preferences. Normative assessment is a different language game, and as such it calls for different kinds of thoughts and concepts. It is of course also true that the subjectivist is going to explain this psychological phenomenon in terms of dispositions, in particular the disposition of taking certain considerations as reasons. But it does not follow that the normative drops out of the picture as irrelevant. Nor does it follow that subjectivists unmask it as a deceptive, or that the normative will be simply displaced by the psychological, whatever that may mean. There is such a thing as a free exercise of reason, and what it requires is that reason operates without hidden and unbidden interference of desires.

So a free exercise of reason is opposed, for one thing, to mere rationalization. It calls for a clear awareness of desires, reasons and the facts. But reasons are not something that we simply find, in this world or another. A free exercise of reason, properly understood, involves the ability to share a linguistic practice, that of considering, weighing, and exchanging reasons. There are various ways in which a free exercise of reason can be hindered or impaired. But there is no reason to think that a free exercise of reason must be independent of our linguistic dispositions so as to achieve objective insight into what reasons there are. Reasons of this kind are accommodated by the subjectivist account.

6 Convergence and Its Explanation

There is one indirect objection to the picture of practical reason that emerges on the subjectivist account that we now need to consider. If a free exercise of reason is, essentially, an exercise of the ability of taking certain considerations as reasons, then how can it be that there is such widespread agreement on what reasons there are, and on how to weigh them? To be sure, there is plenty of disagreement about reasons, too, but it seems that there is also a quite substantial area where our judgments about reasons, both theoretical and practical, converge. This fact requires explanation, and it is not obvious what that explanation could be if we take the subjectivist hypothesis seriously.

In reply, I think we should consider first where convergence of our judgments about our reasons is most likely to be found. This might, of itself, suggest the most likely explanation. There is widespread agreement on a number of reasons in the theoretical domain. Here, one might think of reasons based on induction, causal relations, or logical and mathematical inference. There is every reason not to put your hand into the fire, and that is a judgment about reasons that extremely few are likely to dispute. If you have two apples in your bag, and add two apples to them, you have every reason to believe that there will now be four apples in your bag. This judgment, too, is unlikely to be disputed. It is not logically impossible that after putting two and two in, there are only three. But the fact is that there are nearly always four apples when two and two

are added to one bag, and this fact is fused into the foundations of our language game. Adding two apples to two apples in one bag *is what we call* a reason for thinking that there are now four in it.

Is there any parallel to this in the practical domain? Interestingly, there are fewer judgments about reasons for action that no one is likely to dispute. One may say that most agree that, for example, we need powerful outweighing reasons for killing someone, or that there is normally good reason to tell the truth, and that wanton cruelty is wrong. But on closer scrutiny, this widespread agreement leaves us asking what reasons are good reasons for killing someone, what counts as a good reason to lie, and what constitutes wanton cruelty. Even if there is some agreement as to what counts as wanton cruelty, it is not clear that everyone accepts that he or she has no reason to go in for it, if that is what he or she very much desires.

All the same, the fact remains that there are some general judgments in the practical domain about which most people can be brought to agree, and if they continue to differ, their disagreement will often be explicable in terms of further reasons or considerations that seem to provide even more important grounds for making an exception. These reasons or considerations will inevitably be such that some people do, and some do not, accept them as constituting grounds for an exception, but there is at least this agreement, that for making an exception from a morally important principle such as the moral ban on killing, *some*

sufficiently powerful reason must be available to justify the killing, or else moral blame applies.

What explains this agreement, and what explains the fact that it reaches as far as it does, and no further? There seems to me to be only one plausible answer, and that is that no recognizable human form of life could be sustained without some basic moral rules, and that these basic rules are precisely those upon which our moral judgments tend to converge. Thus, if there were no firm rules against the killing of other members of the society, that society would make survival a most difficult affair. At the limit, such a society would be doomed to disintegrate, either because people die as a result of violence, or because too many resources would have to be devoted to making as sure as one could, if necessary by preemptive means, that one does not get killed. Even if common life of some kind should turn out to be possible in such difficult conditions, it could not be prosperous. As Hobbes famously said of the state of nature, life in a society without any form of moral reassurance would be tantamount to times of war. It is worth a moment of reflection to imagine what this life would be, and Hobbes himself gives a most evocative description:

In such condition, there is no place for industry; because the fruit thereof is uncertain: and consequently no culture of the earth; no navigation, nor use of the commodities that may be imported by sea; no commodious building; no instruments of moving, and removing such things as require much force; no knowledge of the face of the earth; no account of time; no arts; no letters; no society; and which is worst of all, continual fear, and

danger of violent death; and the life of man, solitary, poor, nasty, brutish, and short.²⁵

Transposed to a lower key, similar considerations apply to every other basic moral rule and virtue, such as that of truthfulness.²⁶ Here, what is at stake is not life itself, at least not immediately. But once more it is quite clear that no human society could ever prosper without some reliable and stable disposition on the part of most members to tell the truth. Without such a disposition, there could be no trust and no agreement, and no stable transfer of information from one person to the next. It is vital, then, that some moral requirement to tell the truth, and with it some moral reasons in its favour, are in place if human life as we now know it is to be so much as possible.

At the same time, these examples show why agreements about basic moral rules will always have limits. No society can be sustained without some moral ban on killing, but societies may well be able to sustain themselves if they have a ritual such as that of killing the first-born son. Given the natural ties between parents and children, one expects the rationale for this kind of institution to be perceived as being unusually strong, as it presumably was when children were killed for the sake of deities, or to secure, in some magical way, a good for the society that was perceived as higher than the life of individuals.

²⁵ T. Hobbes, *Leviathan*, pt. 1, ch. 13, § 9, p. 84.

²⁶ Cf. B. Williams, *Truth and Truthfulness*, esp. chs. 2, 3 and 5.

It is important that such practices present themselves as justified in their own right, as something that is mandatory, or quite simply 'to be done'. The thought that ultimately higher goods may be at stake comes in only at the level of analysis and explanation. It need not have a place within the practice, and it need not even be a reason that is recognized as relevant to the justification of the practice by those who participate in it. All the same, the rationale for the rule as well as for the exceptions may be found in their securing what is thought of as a higher good for the society.

Similarly, there may be no way of living in a stable, prosperous community without some commitment to the value of truth and a disposition to respect it. The welfare and stability of a society depends on the presence and effectiveness of such a disposition, but once again, the value of truth is not adequately described as being purely instrumental. Moreover, it is plain that a society can be sustained when there are some exceptions to the rule that 'truth be told'. As we know, communal life can be sustained with lots of lies and other methods of deceiving others.

Consequently, we can assume that moral judgment will converge on some basic rule in favour of telling the truth, but that there will be no similar consensus when it comes to morally defensible exceptions to the rule. The relevance of this for the subjectivist is that he can explain, at least in outline and in general terms, why there is convergence in our moral judgments where we find it and where such convergence is likely to be found. The subjectivist need not be at all impressed, therefore, by those who try to infer from the fact that there is widespread moral agreement

that agreement of this kind reflects some common moral insight, or some shared insight into objective reasons. Without having to go beyond the resources available to him as a subjectivist, he can account for our agreement in moral assessment where we find it, just as he can explain in this way why and where agreement ends and disagreement will prevail.

A realist reply to this might be that the explanatory reasons appealed to by the subjectivist are indeed powerful reasons, but that these explanatory reasons reflect what we should already expect them to reflect, namely our responsiveness to reasons. If there is a rationale to be found for basic institutions of morality, and the subjectivist has given the right account of it, the realist can accept the subjectivist account for what it may be worth and incorporate it into moral realism.

The question is whether we can get an independent grip on the idea that some reasons strike us as important, while some others are more difficult to discern. I think there is good reason to be doubtful that we find an adequate epistemology to back up these rather vague ideas. Even if there is, the fact remains that the subjectivist has a compelling and more economical account of the phenomena. If we still have the sinking feeling that the subjectivist leaves out something important, we might be thinking that the conventions and the kinds of linguistic practices to which subjectivists refer are too slender a basis to carry the full weight of our moral attitudes and institutions. Here, the reply must be that we are in the grip of an entirely wrong picture of the nature of the relevant conventions and the kind of linguistic practices on which the subjectivist analysis and

explanation rest. If we think that a convention is something that we can adopt or change or give up just by *fiat*, or that linguistic practices are too easy to abandon and could never have the hold on us that moral reasons evidently have, then this is not so much a reason for rethinking the subjectivist conception of morality than a reason for rethinking our conceptions of conventions and the linguistic practices in which morality is based.

There is an instructive parallel here with the Wittgensteinian view that “essence is expressed in grammar”.²⁷ It seems that essences are something independent, something given to us in a mode of knowledge that is more direct and more authoritative than our knowledge of a linguistic practice. In a remark on the foundations of mathematics, Wittgenstein sets out the problem and replies to it as follows:

I say, however: if you talk about *essence* – you are merely noting a convention. But here one would like to retort: there is no greater difference than that between a proposition about the depth of the essence and one about – a mere convention. But what if I reply: to the *depth* that we see in the essence there corresponds the *deep* need for the convention.²⁸

This is true in mathematics, where the rules are such that a mathematical proposition cannot be refuted by experience. All the same, the necessity involved is that of a rule. But it is not an arbitrary rule. The language game in which the rules of mathematics have their life is basic and most

²⁷ L. Wittgenstein, *Philosophical Investigations*, § 371.

²⁸ L. Wittgenstein, *Remarks on the Foundations of Mathematics*, I-74.

fundamental, so that we will find it difficult to conceive of a world in which our language game of mathematics would no longer have a use. Mathematical essences strike us as particularly deep, but that depth reflects the deep need for the mathematical convention.

The same is true when we consider moral reasons. These reasons are particularly deep and have a great significance for those who care about morality and the aims that it promotes. This does not seem to sit easily with the idea that morality and moral language are ultimately based on linguistic practices. But if we remember the plain fact that institutions of morality are deep and especially important, then it seems that the objection loses much of its initial force. The depth of moral reasons will be due to a deep need for the convention. And what deeper need could there possibly be than that for moral guidance in our actions? In its absence, there could not be a recognizably human form of life at all.

7 The Language of Morality

A striking fact about the language of morality is that it revolves around the notion of a fact. We say, of moral judgments, that they can be *true* or *false*; that there is such a thing as moral *knowledge*, and correspondingly, that there are various forms of moral *ignorance*; that there is something that we *see* to be so, while others are *blind* to it; that we are quite *certain* that some action is the right or wrong one, or that there is at the very least good reason to *believe* this.

Similarly, we have no objection to speaking of moral *errors* and *mistakes*; we say that a person is not *accurate* or not *reliable* in moral judgment, and that some of the moral views a person holds, ourselves included, may be *wrong*. We say that a person may well *miss* something in moral matters, and that such a person may come to *acknowledge* and *appreciate* that something she thought was right was really wrong. We distinguish as a matter of course between *justified* and *unjustified* moral judgments, and all shades in between. Finally, and obviously, we say that certain things that people do *are* right or wrong, that certain things *are* good or bad or worth caring about or not important, and that there *are* certain values that make claims on us, and that moral considerations are such that they *ought* to be respected. It seems plain that this feature of moral thought and practice does not sit easily with the subjectivist account.

There is a special problem about 'ought' because when we say that someone ought to do something, we frequently issue a demand to act in that way. This 'ought' must engage with the beliefs and motivations of the agent if it is the 'ought' of reason, but 'ought' can be used to issue a demand that makes no reference to beliefs or dispositions whatsoever, either on the part of speaker or the addressee. After all, 'You ought to tell her the truth' does not mean 'You ought to tell her the truth unless you do not want to'. That you do not want to we may safely take for granted if this kind of thought occurs to you at all. The 'ought' tells you to be truthful whether you like it or not. How are we to account for such

'categorical' imperatives, if we take seriously the thought that moral reasons, facts and obligations are ultimately based on dispositions?

The problem is not, or not in the first instance, that there is a problem concerning the possibility of 'hypothetical' imperatives.²⁹ These express 'oughts' that are, implicitly or explicitly, relative to some subjective condition, such as the presence of an appropriate motivation. 'If you take an interest in Renaissance art, you ought to go to Florence' is a clear example. This 'ought' is relative to taking some interest in renaissance art, but it leaves entirely open whether or not one ought to take that interest. So this 'ought' is clearly relative to taking such an interest.

There are very many 'oughts' like that, and the problem is not that subjectivists have nothing to say about them, or too little. The problem is that not every 'ought' will fit that mould, and moral 'oughts' provide a clear example. If I say 'You ought to keep your word', then what I mean is not that you ought to keep your word unless you happen to be motivated otherwise. Indeed, the sentence 'You ought to keep your word unless you do not want to' does not as yet have a clear sense, because it is quite obvious that there is only an occasion for this kind of utterance if the person to whom it is addressed would rather do something else, so that a reminder of her moral obligation is appropriate. What we want to say is

²⁹ For the distinction between 'categorical' and 'hypothetical' imperatives, see I. Kant, *Groundwork for the Metaphysics of Morals*, AA 4:414-20. On the possibility of hypothetical imperatives, see C. Korsgaard, 'The Normativity of Instrumental Reason', in: G. Cullity and B. Gaut (eds.), *Ethics and Practical Reason*, pp. 215-54.

that the person ought to keep her word whether she likes it or not. But if this is the situation, it looks as if subjectivism fails. Subjectivism seems to leave no room for the idea that someone ought to keep their word regardless of their present motivations, for it categorically denies that there are any 'oughts' that are not relative to some present motivation.

But this is confused. A first point to remember is that conditional and categorical statements about what someone ought to do have different uses. A statement of the form 'If you have motivation m , you ought to ϕ ' issues a conditional 'ought'. In other words, it does not categorically demand of someone, perhaps as a result of a more general demand on everyone, to ϕ . It merely states that given motivation m , which an agent may or may not have, it would make sense for him to ϕ . The force of 'would make sense' is that of some rational connection, which is why such statements translate fairly easily into statements about reasons for action. In our previous example, the rational connection in question is easy enough to make. Since there is plenty of Renaissance art on display in Florence, someone who takes an interest in Renaissance art has reason to go there to see it. This is what we *call* a reason to go to Florence.

'You ought to be truthful' often has a very different function. It abstracts from present motivations, so it issues a demand that is, in that sense, categorical. It probably derives from a more general demand to be truthful, which amounts to the default principle that one should be truthful unless there is good reason not to. The suggestion is that in the

present situation, there are no such reasons in the offing. The sentence thus expresses a demand to set aside the present inclination, which is, we may presume, an inclination to lie or keep quiet, and to do instead what truthfulness demands.

On the subjectivist account, the addressee of the demand may well have a reason to tell the truth, and if he is not aware of it, the statement may serve as a reminder that he does have such a reason. The agent has a reason to be truthful just in case he has some disposition to tell the truth unless there is good reason not to, or at least, some disposition to see telling the truth as something that is normally a good or mandatory thing to do. Obviously, he may have such a disposition even if at present he is strongly disinclined to act on it. Whether or not he has good reason all things considered to be truthful will depend on other motivations, dispositions and reasons for action he presently has. But the demand to tell the truth need not be conditional in this way. It may be that the addressee has a reason to do what he is told he ought to do, but this need not be so. If he has no such reason, the statement will still issue moral expectations and demands, and as such it will continue to be used in ways that are quite different from the uses of the hypothetical statement 'You ought to tell the truth unless you do not want to'. The latter issues a permission to do something else, while the former closes off that course of action. So on the subjectivist account, a significant difference remains between hypothetical imperatives and statements that make no reference to subjective motivations and are, in that sense, categorical.

The same is true if we adopt the perspective of the person who deliberates. If the subjectivist could no longer have thoughts of the form 'I ought to be truthful', but could only say to himself 'I ought to be truthful unless I am motivated otherwise', his moral standpoint would be largely lost. But he can have thoughts of either form, and if the critic claims that in consistency, he can only have thoughts of the form 'I ought to be truthful unless I am motivated otherwise', then the critic is ascribing to subjectivists a thought they do not have. The subjectivist who thinks that he ought to be truthful has a disposition to be truthful, not a disposition to be truthful unless he happens to be motivated otherwise. To have the thought 'I ought to be truthful', where one does not deceive oneself, is what it is to have that disposition.

What are we to say about the other factual and epistemic terms that are an essential part of moral language? Once again, it is important to remember that their meaning is not given by a moral theory, but based on their role in our linguistic practice. To be clear what factual and epistemic terms mean in the context of our moral discourse, we must look at their ordinary use. When we do, we see quite quickly that there is a characteristic function for such weighty terms as 'know' and 'true' and 'fact', and what that function is can be brought out when these terms are contrasted with some of their common counterparts, which might be used instead.

One thing that is clear is that the grammar of terms such as 'knowing', 'being certain', 'error', 'truth' and 'fact' has a signaling function

that sets them apart from mere expressions of subjective preference. They are used for taking up a moral stance, and that moral stance is one with which others can agree or disagree. If we could only say that some action *seems* to be virtuous, or *appears* to be good or bad or right or wrong, that this is how one *sees* the matter, or that this is the course of action one would personally *prefer*, then it would become quite difficult to take a moral stance towards anything. The moral stance is marked by categorical assertions, which express special commitments to the moral proposition that is being expressed. Someone who says 'I prefer to act like this' merely expresses his personal preference, and so far that is fully compatible with expressions of a different preference. If all we ever said about some moral issue could be phrased in terms like 'I feel ...' or 'My attitude to this is ...', there could be no moral disagreement. As Williams puts it, the situation would be like that of two people on a boat, one of whom says that he feels sick while the other says that he, on the other hand, does not.³⁰ There is evidently moral disagreement, so there must be more to moral language than a set of terms the function of which is confined to reporting moral attitudes.

We can now see that there is a crucial difference between factual and epistemic language and expressions like 'appear' or 'seem'. Expressions of the first type, such as 'This is wrong', or 'She knew that she had to do it' are used to express a firm commitment to a moral judgment. This is different in the case of 'seeming', 'thinking' or 'appearing'.

³⁰ B. Williams, *Morality*, p. 30.

Statements of the latter type, such as 'This seems to me to be wrong', or 'She thought that she had to do it' do not express the same kind of commitment. They remain, as it were, at a safe distance, and they invite critical queries or objections in a way in which the statements of the other type do not. 'It seems wrong' or 'That looks like a good thing' or 'This appears to be all right' all suggest that different ways of looking at the act in question might be just as adequate, or indeed better justified than the initial judgment. There is a need for nuance in expressing moral judgments. It is not surprising, therefore, that there is a range of linguistic tools at our disposal that express those different nuances.

It is interesting to note in this connection that so far, this is entirely in agreement with expressivist analyses of moral thought. The difference is that the subjectivist as I envisage him does not have to deny that there are moral facts, and allows that there is moral knowledge. Expressivists, by contrast, have to claim that, at the underlying level, there are no such things as moral facts, beliefs or knowledge. The subjectivist can get around such elements of error theory, and in this way, he stays in close touch with the phenomena.

At the same time, the subjectivist has gained an important diagnostic insight. If it is correct that terms like 'know' and 'true' and 'fact' and 'seem' work similarly in our moral and our ordinary factual discourse, we have one more explanation of the fact that moral facts can easily appear to be objective, just like the facts of nature are. The distinction between truth and objectivity is drawn at the level of analysis

and explanation, so we should not expect it to manifest itself on the linguistic surface. It is natural to think that moral facts and facts of nature are objective, and it takes analysis and explanation to see there is a crucial difference between these two kinds of fact.

8 The Epistemic Status of Morality

However, there is still a problem about epistemic terms when these are used in moral discourse. As we have seen, the subjectivist does not deny that there can be moral belief and even knowledge. But morality seems to possess a special epistemic status. Moral judgments sometimes carry a particular conviction, which is manifest in moral judgments that strike us as evidently true. The question is how the subjectivist is able to account for that.

A first question to ask is what we mean when we speak of a special epistemic status, or think that a moral judgment strikes us, without doubt, as true. Is it more, or much more, than the fact that we use a certain moral concept confidently, that we have no doubts about the moral judgment that is made, and that this certainty will not be shaken? The subjectivist will be inclined to doubt that.

Second, even if we think of moral judgment as involving some objective insight, that insight cannot be such that it cuts off all questions of justification. It is tempting to think that a moral judgment can strike one immediately as true. But someone who says 'This strikes me immediately

as true', and insists that he is unable to give reasons for his judgment, runs the risk of dogmatism. If no reasons whatsoever for a moral judgment are forthcoming, then it is quite doubtful that the speaker even occupies a moral standpoint. 'I know it to be so' is not an argument, but a refusal to engage in argument. It is an essential feature of the language game that a moral judgment can, in principle, be backed up by justifications.

This does not immediately refute the thought that there is some objective insight, but it helps us to get rid of a tempting misconception of this insight. Moral judgments do not stand in isolation. Moral judgments are normally justified against the background of a context, and to think of a moral judgment as true is to take many other truths for granted. If so, the idea that moral judgments can be justified by some experience or insight must be false.

Our ability to judge in moral matters is much more complicated than any such simplistic model of objective insight can allow. Nevertheless, there are evidently cases where we cannot say more, or much more, by way of justification of our moral judgment than that this is evidently how things stand. 'This was brave', 'She needs help', or 'Cruelty is wrong' might be examples. If you do not understand already why causing pain is bad, what could I possibly say to convince you?

The question is whether this observation provides any evidence for moral realism. Do we really want to say that this is where we find examples of objective moral insight?

Consider again a parallel case, that of judging that this flower here in front of me is red. When I look at the flower, this fact strikes me as obvious, as it would strike anyone who sees colours and distinguishes them as I do. There is no doubt that this flower is red. Nothing could convince me of the contrary. But does this mean that this insight has some special epistemic status? In particular, does it mean that this kind of certainty must 'come from the outside'?

- Clearly not. Of course the flower is something that exists quite independently of me, so in that sense, knowing that this flower is red is to know something about an object that exists independently of me. But my knowledge that this flower is red is not itself objective knowledge. Indeed, it is a paradigm of knowledge that is based on substantial subjective conditions, namely, our specific way of seeing colours and our classifying them in response to our perceptions in the way we do. Our response comes naturally. But this is not evidence for some epistemic faculty that puts us in touch with some mysterious dimension of reality. It is the reflection of our natural response to light reflected by an object and successful linguistic training.

The situation is essentially the same with our ability for moral judgment. The grammar of ethical words is vastly more complex than that of colour words, which form a fairly simple spectrum, or a colour sphere. There are various formal and informal inferential relations between moral terms and judgments, and this is part of the explanation why experiences

of certainty in moral judgment are, for the morally sensible, a fairly rare occurrence.

But there are such experiences, and where they occur, the subjectivist has a compelling explanation for them. He points out that our moral education involves a cultivation of feelings and a development of character, and that to a large extent, it is a form of linguistic training. This training has produced a stable set of dispositions to respond to certain situations by using a range of moral expressions, and when we respond in this way, there is frequently no room for either doubt or hesitation.

Moreover, the subjectivist can explain in this way why the reaction is experienced as one that normally happens of its own accord, not requiring any conscious effort or reflection. But the explanation of phenomena like this does not appeal to a mysterious moral insight into more mysterious realms of moral fact. It appeals to our largely unconscious mastery of a technique, a linguistic practice.

So compared to our abilities of seeing colours, seeing moral facts is an extremely complicated, culturally informed and individually acquired form of seeing. For that reason, we should also expect it to be a form of seeing that is subject to variation and change in response to different cultural, historical, and psychological conditions. Here, the analogy with colours has a clear limit.

But this very difference reinforces the subjectivist account of the phenomena. As with all words, it is true of moral words that what it is to acquire them is essentially to learn to use them correctly. Learning to use

them correctly, however, is by no means the same as to acquire a brute disposition to utter a word when others would utter it too. Rather, it involves an understanding of a range of situations where the word would or would not apply, as well as an understanding of its relations to other words that could be used instead, and generally, of the inferential space that surrounds it.

The inferential space and the range of considerations that affect a moral judgment, sometimes radically, is much wider in the case of moral words than it is with words for colours. Because that space is wider, moral concepts can be related to one another in ways for which there is no parallel in the domain of colour concepts. Moreover, their application is more sensitive to other kinds of fact. For all these reasons, it is not surprising that moral notions are often undetermined, that there will always be grey areas, and room for uncertainty and doubt and disagreement. Nevertheless, at the end of a successful moral education, we have acquired a wide range of moral concepts that we share, at least to a large extent, and from then on, see the world through them.

9 The Charge of Arbitrariness and Power

Suppose we grant that the subjectivist can make sense of the epistemic status of morality and of its stability and independence. As we have seen, subjectivists do not have to think of moral motivation as something that merely befalls us, so that with respect to it, we would be mere bystanders

or passive recipients. Morality is closely bound up with our sense of ourselves, and with our practical identities.

But this fact may not alleviate, but reinforce the worry that our moral confidence is a form of dogmatism, and that moral reasons and the moral point of view in which they play a central role can only be arbitrary if subjectivism is true. If moral discourse, moral argument, and all forms of ethical education turn out to be nothing but an exercise of power, this would question the authority of moral thought itself. At the limit, it would undermine the moral point of view by unmasking it as being based on a pretense. The moral point of view is marked by its rejection of mere power and its unrestrained exercise, granting equal consideration to the interests of all. It would be a most radical dismantling of the moral point of view if its own claim to authority turned out to rest on nothing more than a brute exercise of power.

This is a substantial charge. But so far, it is merely a picture. What exactly is the charge of arbitrariness and power? What exactly is the content of, and the basis for, the challenge? On closer inspection, I think it emerges that the charges are misleading, and that it is not even clear that they are fully intelligible. Either way, they need not worry us.

Consider a case where we all agree that an action is arbitrary, and moreover, that the action is nothing more than a brute exercise of power. Imagine a king like Louis XIV, who fancies himself to be a ruler by the grace of God. He rules in the spirit of '*L'état, c'est moi*', and does not for a moment doubt that this is fully justified. Our king very much desires to

build yet another grand *palais* to manifest his status. There is no money at the moment because many palaces have already been built, so our king comes up with the idea of claiming a new tax from all his people. The people are poor, and it is well known that they strongly object to being treated in this way by their ruler. The king is well aware of this, but looking at the new architectural drawings, he brushes all worries aside. His desire is strong, and all the means towards his end are already at his disposal. His status gives him all the authority he needs. Without hesitation, he puts on his wig, calls for the dictate and proceeds to sign the bill. In doing so, he does not waste time considering the interests of his people, nor is he at all interested what consequences his imposing yet another tax might have. The money comes, the project goes ahead, and that is all he cares about.

An act like this is truly arbitrary, and we understand why for that reason, there is moral blame attached to it. But if this is our model for the charge of arbitrariness and power, we can also see that this charge cannot possibly apply to moral attitudes and arguments across the board.

First, there is no question that the motivations of a moral agent can be reduced to the pursuit of his self-interest. In addition to such selfish motivations that a moral agent undoubtedly has, there are many genuinely altruistic motivations, and these may be, to him, the most important. Of course there is a sense in which even the most altruistic person acts as she sees fit, and does what she wants to do. But there is no

reason to think that what she wants to do is something other than what it appears to be, namely to promote the good of others.

Furthermore, it is not true that our commitment to morality could be arbitrary in the sense of springing from our wills. We do not choose our morality, and what someone finds truly important does not receive that status merely by an act of will. It is the other way around. What we truly care about will shape our wills and find expression in our attitudes and choices. Evidently, we can critically reflect on choices, as much as we can critically reflect on our moral attitudes. But this does not happen in a vacuum. All criticism and revision takes place against the background of beliefs and attitudes that we must simply take for granted. As the case of Nietzsche shows, even the most radical challenge of morality must be based on *values*, values, moreover, that are understood as being even more important than the aims morality promotes. For Nietzsche, this was a radically different ideal of what a man should be. We may not agree with this ideal, and doubt that morality is as hostile to the great man as Nietzsche makes it out to be, but the fact remains that all revising of our moral attitudes must take place against a background of some other attitudes, which are not questioned, and which have not been chosen or adopted arbitrarily.

A similar point can be made about the charge of dogmatism. We describe someone as being dogmatic when he refuses to engage in argument, or finds that there is no further argument he could appeal to. The important point is that someone is charged with dogmatism when he

refuses to engage in argument where we think that arguments and reasons have not yet run out. Similarly, a person is being dogmatic when she sticks to a belief in the absence of good reasons and ignores the reasons that are close at hand. If this is the grammatical environment for charges of dogmatism, then we can see very quickly that it cannot possibly apply across the board. Where someone believes that causing pain is bad, and sticks to that belief in the absence of a further reason, this insistence does not count as dogmatism. After all, where could further reasons for believing this be found? It would be like saying that our belief that there are other human beings who have similar experiences is really a form of dogmatism. It is not even clear what this could mean.

In the light of these reflections, we can see that there may well be individual judgments, attitudes and actions that are justly called dogmatic or arbitrary. They are justly called arbitrary when they are arrived at without due reflection, or when they do not sufficiently take others and their interests into account. An attitude is justly called dogmatic when it is not sensitive to reasons, while there are relevant reasons available that the person does not see. An attitude is not dogmatic when no relevant reasons are available, and questions of justification do not even arise.

The situation is essentially the same when we consider the content of the charge of power. Once again, there are many ways of influencing others that are little more than a brute exercise of power, and this happens frequently where the object of exerting such an influence is to instill moral dispositions. There are far too many moral missionaries, and their

strategies bear witness to their lack of moral sensitivity towards the rights and needs of others, sometimes even their own. But once again, this charge could not possibly apply to moral judgment as a whole. Undue influence or pressure, which attracts the charge of power, presupposes a substantial contrast with a different way of interacting with, and indeed of influencing, others. If everything, including argument and giving reasons, is said to be an exercise of power, we do not even understand what is being said.

It is interesting to note in this connection that the charge of arbitrariness and power is itself *a moral charge*. Where we accuse someone of merely exercising power, or of taking advantage of their privileged position, we occupy the standpoint of those who were wronged and defend their claim to justice. But if that is so, it makes no sense to ask whether moral thoughts and attitudes might be a 'mere' exercise of power. In effect, there would be no moral standpoint left from which we could articulate that claim.

Whatever the charges of arbitrariness and power are said to come to, it is certain that they can only succeed if there is a set of standards against which the attitude or action fails. There must be a substantial contrast. We should act, *not* arbitrarily, *but* with due consideration and respect for others. We should structure our interactions, *not* by exerting power wherever we can, *but* by respecting the rights and the interests of others. It is an empty question whether the standards against which certain actions count as arbitrary or as a brute exercise of power might

themselves be purely arbitrary or such a brute exercise of power. Every contrast would be lost, and that means that the charges are devoid of content.

Still, there is a worry that we might be doing something wrong when we accuse some other party of having done something we think is wrong. After all, if there is no objective moral truth, then it seems that the judgment that what was done was right is equally well justified, at least from the neutral and objective standpoint. How could it be anything but arbitrary, and indeed an exercise of power, if we sanction immoral behaviour on the basis of our attitudes in the way we do?

This is confused. There is of course a sense in which there is a neutral ground, and from that ground all moral attitudes are seen as what they are, namely subjective. There is no right and wrong from that perspective, but so far, this is nothing more than a description of subjectivism. If this insight is supposed to undermine our confidence and our right to take a moral stance when criticizing others, it must involve the idea that we may, in the end, be wrong. This is a thought that no subjectivist can have, unless it means, quite harmlessly, that there is room for doubt and criticism within our subjective moral outlook. As Williams points out, insofar as there is any traceable intellectual link between subjectivism and this kind of moral indifference, it must involve something like the relativist error. The thought must be something like this:

'Because subjectivism is true, I am not justified in protesting.' If this is right, it must be so either because, if subjectivism is true, no one is justified in doing anything; or because, if subjectivism is true, he is specially not justified in protesting. If the first of these, then the inflictors of the injustice are not justified in what they are doing either, nor is this man justified in *not* protesting, and these considerations remove any *basis* he was supposed to have for his indifference.³¹

If the charge is supposed to be that the man is specially not justified in protesting, then this thought cannot come from thinking, further, 'They may be right, and who am I to say that they are wrong'. This thought moves him out of the subjectivist arena by importing the idea that there is some objective standard that he may have missed, and this is a thought that he, as a subjectivist, cannot in consistency have.

It would be a mistake to conclude from this, as it were triumphantly, that we are free to interfere with anyone in any circumstances, or to impose our moral outlook as we please. This would be a form of moral imperialism that would be even harder to tolerate than the stance of moral indifference that seemed to threaten on the other side. The truth is that when it is justified to interfere or to criticize and when we should accept that others continue to see things differently is decided from within our moral practice. Our modern moral outlook provides resources to distinguish, in a reasonable way, between forms of interference that are justified and those that are not, and one of the many relevant considerations on which we can count is that there is such a thing as moral arrogance. Some forms of interference attract moral blame, and the same is

³¹ B. Williams, *Morality*, pp. 40-1.

true of certain forms of ethical indifference. If I happen to find out that my neighbour locks up his children, I may be subject to blame if I fail to bring up the issue, or call the police to let them out. Once again, the relevant distinctions are internal to our moral outlook, and are not affected by our moral theory.

10 Conclusion

These remarks conclude the inference to the best explanation. I have argued, using what seem to me to be the most striking examples, that our subjectivist can make good sense even of the ethical phenomena that initially seem to resist his theoretical analysis. Our subjectivist may even have the better explanation, but if so, this is already more than needs to be established. What the previous arguments have shown is that the subjectivist is clearly at a theoretical advantage. His explanation of morality is coherent and, essentially, more economical than any rival explanation of the realist, insofar as realists can even offer explanations. Unlike realists, subjectivists can make do with resources that a sensible observer of human affairs must allow and use already. In this way, a subjectivist account of ethics can contribute to a properly naturalistic understanding of the human form of life.

There is likely to be a residue, and a sense of unease if we let go of aspirations to substantial moral objectivity. If there is no ethical objectivity, this can be experienced as a loss. Like the loss of religion, or

the loss of an orderly cosmos, the loss of moral objectivity can cause despair, and the feeling of a loss of guidance. There may be a persistent worry that we are left to our own devices, and that these devices might not be enough.

Here, it seems to me important to remember that the loss of moral objectivity can also be experienced as a fundamental liberation. The fact that value is our own creation is a testimony to our independence. We cannot be forced by the world to adopt one set of values rather than another, and that means that our existence is, in yet another sense, free and autonomous.³² In any case, if forms of life are the foundation, it is best for us to face that fact.

In my view, there is reason to be moderately optimistic. Ultimately, there is no such thing as a rational foundation for ethics, and no foundation for it in objective fact. But along with freeing ourselves from the false idea that there can be a rational foundation, we have freed ourselves from the illusion that we need such foundations if ethical life is to continue, and continue to remain significant. All justification must come to an end somewhere, and this is true in the theoretical as well as in the practical domain. There is no *reason* for believing that there are other human beings, or that the earth is more than a few minutes old. But no one *needs* a reason to believe that there are other human beings, and that the earth is more than a few minutes old. There is no *reason* to engage in moral thinking, but it is not obvious why we should *need* a reason to

³² Cf. B. Williams *Ethics and the Limits of Philosophy*, p. 128.

engage in moral thinking. We can see the point of rules and dispositions that protect the moral form of life, which is life as we know it.

Ultimately, there is nothing more than practices to keep us on the rails. The ultimate supports of value are our dispositions.³³ But this is a good place for ultimate supports to be. Hume famously said that “reason is, and ought only to be, the slave of the passions, and can never pretend to any other office than to serve and obey them”.³⁴ There is an important truth in what he says. Reason is based on dispositions. Moreover, reasons run out somewhere. If so, reason has to start from somewhere, and our dispositions are the only basis on which practices like the practice of exchanging reasons can be established and sustained. Far from being threats to rationality, forms of life that lie beyond justification are required if justification is to be so much as possible. Forms of life are not opposed to rational abilities, but lie at their foundation.

Wittgenstein has opened up and taken seriously the possibility that forms of life are the foundation. If this is right, it is a most important insight, and we must accept that even where we least expect it, the contingency of human life comes to the surface. But if Wittgenstein is right, we can both realize that forms of life are the foundation, remain reflectively aware of it, and continue to live humanly. If what I have said is correct, the moral form of life is no exception.

³³ B. Williams, *Ethics and the Limits of Philosophy*, p. 52; and ‘The Primacy of Dispositions’.

³⁴ D. Hume, *A Treatise of Human Nature*, bk. II, pt. III, sect. III, p. 415.

Living With Subjectivism

This insight might, moreover, help us to live better. Life goes better when it is not based on a pretense, and when we are liable to fewer illusions rather than more. A. E. Housman observed:

The house of delusions is cheap to build, but draughty to live in, and ready at any instant to fall.³⁵

His advice in response to the situation is a sound one. Prudence requires that we move our furniture “betimes into the open air”.³⁶

³⁵ A. E. Housman, *Introductory Lecture 1892*, p. 38.

³⁶ *Ibid.*

Bibliography

- Altham, J. E. J., and Harrison, R. (eds.), *World, Mind, and Ethics: Essays on the Ethical Philosophy of Bernard Williams* (Cambridge: Cambridge University Press, 1995)
- Anscombe, G. E. M., *Ethics, Religion, and Politics: The Collected Philosophical Papers*, vol. 3 (Minneapolis: University of Minnesota Press, 1981)
- *Intention* (Oxford: Basil Blackwell, 1957)
- Aristotle, *Nicomachean Ethics*. Translated by S. Broadie and C. Rowe (Oxford: Oxford University Press, 2002)
- Ayer, A. J., *Language, Truth and Logic* (Second edition, London: Victor Gollancz, 1946)
- Bailey, C., *The Greek Atomists and Epicurus: A Study* (Oxford: Clarendon Press, 1928)
- Baker, G. P. and Hacker, P. M. S., *Wittgenstein: Understanding and Meaning. Volume 1 of An Analytical Commentary on the Philosophical Investigations. Part 1, Essays* (Second edition, Oxford: Blackwell, 2005)
- *Wittgenstein: Understanding and Meaning. Volume 1 of An Analytical Commentary on the Philosophical Investigations. Part 2, Exegesis* (Second edition, Oxford: Blackwell, 2005)
- Barrett, C., *Wittgenstein on Ethics and Religious Belief* (Oxford: Basil Blackwell, 1991)
- Bennett, J., *Locke, Berkeley, Hume: Central Themes* (Oxford: Clarendon Press, 1971)
- Bird, G., *William James* (London: Routledge and Kegan Paul, 1986)
- Blackburn, S. and K. Simmons (eds.), *Truth* (Oxford: Oxford University Press, 1999)
- Blackburn, S., *Essays in Quasi-Realism* (New York: Oxford University Press, 1993)
- *Ruling Passions: A Theory of Practical Reasoning* (Oxford: Clarendon Press, 1998)
- *Spreading the Word: Groundings in the Philosophy of Language* (Oxford: Clarendon Press, 1984)
- *Truth: A Guide For the Perplexed* (London: Allen Lane, 2005)
- Charles, D. and Lennon, K. (eds.), *Reduction, Explanation and Realism* (Oxford: Clarendon Press, 1992)
- Child, W., *Causality, Interpretation and the Mind* (Oxford: Clarendon Press, 1994)
- Clack, B. R., *An Introduction to Wittgenstein's Philosophy of Religion* (Edinburgh: Edinburgh University Press, 1999)
- Cohon, R., 'Are External Reasons Impossible?', in: *Ethics* 96:3 (1986), 545–556
- Crisp, R. (ed.), *How Should One Live? Essays on the Virtues* (Oxford: Clarendon Press, 1998)
- Cullity, G., and Gaut, B. (eds.), *Ethics and Practical Reason* (Oxford: Clarendon Press, 1997)
- Dancy, J., *Ethics Without Principles* (Oxford: Clarendon Press, 2004)
- Davidson, D., *Essays on Actions and Events* (Oxford: Clarendon Press, 1980)
- *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984)
- *Problems of Rationality* (Oxford: Clarendon Press, 2004)
- *Subjective, Intersubjective, Objective* (Oxford: Clarendon Press, 2001)
- *Truth, Language, and History* (Oxford: Clarendon Press, 2005)

Bibliography

- Diamond, C., *The Realistic Spirit: Wittgenstein, Philosophy, and the Mind* (Cambridge, MA: MIT Press, 1991)
- ‘Wittgenstein, Mathematics, and Ethics: Resisting the Attractions of Realism’, in: H. Sluga and D. G. Stern (eds.), *The Cambridge Companion to Wittgenstein*, 226-260
- Drury, M. O’C., *The Danger of Words* (London: Routledge and Kegan Paul, 1973)
- Edwards, J. C., *Ethics Without Philosophy: Wittgenstein and the Moral Life* (Tampa: University Presses of Florida, 1982)
- Fogelin, R. J., *Wittgenstein* (Second edition, London: Routledge and Kegan Paul, 1987)
- Foot, P., *Moral Dilemmas and Other Topics in Moral Philosophy* (Oxford: Clarendon Press, 2002)
- *Natural Goodness* (Oxford: Clarendon Press, 2001)
 - (ed.), *Theories of Ethics* (Oxford: Oxford University Press, 1967)
 - *Virtues and Vices and Other Essays in Moral Philosophy* (Second edition, Oxford: Clarendon Press, 2002)
- Forster, M. N., *Wittgenstein on the Arbitrariness of Grammar* (Princeton: Princeton University Press, 2004)
- Frankfurt, H. G., *Necessity, Volition, and Love* (Cambridge: Cambridge University Press, 1999)
- *The Importance of What We Care About: Philosophical Essays* (Cambridge: Cambridge University Press, 1988)
 - *The Reasons of Love* (Princeton: Princeton University Press, 2004)
- Frede, M. and Striker, G. (eds.), *Rationality in Greek Thought* (Oxford: Clarendon Press, 1996)
- Frege, G., *Collected Papers on Mathematics, Logic, and Philosophy*. Ed. B. McGuinness (Oxford: Basil Blackwell, 1984)
- Gaskin, J. C. A., *Hume’s Philosophy of Religion* (Second edition, London: Macmillan, 1988)
- Gerrard, S., ‘A Philosophy of Mathematics Between Two Camps’, in: H. Sluga and D. G. Stern (eds.), *The Cambridge Companion to Wittgenstein*, 171-197
- Gibbard, A., *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Oxford: Clarendon Press, 1990)
- Glock, H.-J., *A Wittgenstein Dictionary* (Oxford: Blackwell, 1996)
- ‘Necessity and Normativity’, in: H. Sluga and D. G. Stern (eds.), *The Cambridge Companion to Wittgenstein*, 198-225
- Hacker, P. M. S., *Insight and Illusion: Themes in the Philosophy of Wittgenstein* (Revised edition, Oxford: Clarendon Press, 1986)
- *Wittgenstein: Connections and Controversies* (Oxford: Clarendon Press, 2001)
 - *Wittgenstein’s Place in Twentieth-Century Analytic Philosophy* (Oxford: Blackwell, 1996)
- Hare, R. M., *The Language of Morals* (Oxford: Clarendon Press, 1952)
- Harman, G. H., *Explaining Value and Other Essays in Moral Philosophy* (Oxford: Clarendon Press, 2000)
- ‘The Inference to the Best Explanation’, in: *The Philosophical Review* 74:1, (January 1965), 88-95
 - *The Nature of Morality: An Introduction to Ethics* (New York: Oxford University Press, 1977)
- Hobbes, T., *Leviathan*. Ed. J. C. A. Gaskin (Oxford: Oxford University Press, 1996)
- Holtzman, S. and Leich, C. (eds.), *Wittgenstein: To Follow a Rule* (London: Routledge and Kegan Paul, 1981)
- Hooker, B. (ed.), *Truth in Ethics* (Oxford: Blackwell, 1996)
- ‘Williams’ Argument Against External Reasons’, in: *Analysis* 47 (1987), 42-44
- Horwich, P., *Truth* (Second edition, Oxford: Clarendon Press, 1998)
- Housman, A. E., *Introductory Lecture 1892* (Cambridge: Cambridge University Press, 1937)
- Hume, D., *A Treatise of Human Nature*. Ed. L. A. Selby-Bigge (Second edition, Oxford: Clarendon Press, 1978)

Bibliography

- *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*. Ed. L. A. Selby-Bigge (Third edition, Oxford: Clarendon Press, 1975)
- *The Natural History of Religion and Dialogues Concerning Natural Religion*. Ed. A. W. Colver and J. V. Price (Oxford: Clarendon Press, 1976)
- James, W., *Pragmatism: A New Name for Some Old Ways of Thinking* (London: Longmans, Green, and Co., 1907)
- *The Varieties of Religious Experience: A Study in Human Nature* (London: Longmans, Green, and Co., 1902)
- Jenkins, M. P., *Bernard Williams* (Chesham: Acumen, 2006)
- Johnston, P., *Wittgenstein and Moral Philosophy* (London: Routledge, 1989)
- Kant, I., *Critique of Pure Reason*. Translated by N. Kemp Smith (London: Macmillan, 1929)
- *Groundwork for the Metaphysics of Morals*. Ed. T. E. Hill and A. Zweig (Oxford: Oxford University Press, 2002)
- *Prolegomena to Any Future Metaphysics That Will Be Able to Come Forward as a Science*. Ed. G. Hatfield (Revised edition, Cambridge: Cambridge University Press, 2004)
- Kolnai, A., *Ethics, Value and Reality: Selected Papers of Aurel Kolnai* (London: Athlone Press, 1977)
- Korsgaard, C. M., *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996)
- ‘The Normativity of Instrumental Reason’, in: G. Cullity and B. Gaut (eds.), *Ethics and Practical Reason*, 215-54
- Kripke, S. A., *Wittgenstein on Rules and Private Language: An Elementary Exposition* (Oxford: Basil Blackwell, 1982)
- Lovibond, S., *Realism and Imagination in Ethics* (Oxford: Basil Blackwell, 1983)
- Mackie, J. L., *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin, 1977)
- *Hume’s Moral Theory* (London: Routledge and Kegan Paul, 1980)
- *Problems from Locke* (Oxford: Clarendon Press, 1976)
- Malcolm, N., *Ludwig Wittgenstein: A Memoir* (Second edition, Oxford: Clarendon Press, 2001)
- McCarthy, T. G. and Stidd, S. C. (eds.), *Wittgenstein in America* (Oxford: Clarendon Press, 2001)
- McDowell, J., *Mind and World* (Cambridge, MA: Harvard University Press, 1994)
- *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998)
- McGinn, C., *The Subjective View: Secondary Qualities and Indexical Thoughts* (Oxford: Clarendon Press, 1983)
- McGinn, M., *Wittgenstein and the Philosophical Investigations* (London: Routledge, 1997)
- McGuinness, B., *Young Ludwig: Wittgenstein’s Life 1889–1921* (Oxford: Clarendon Press, 2005)
- McManus, D. (ed.), *Wittgenstein and Scepticism* (London: Routledge, 2004)
- Millgram, E. (ed.), *Varieties of Practical Reasoning* (Cambridge, MA: MIT Press, 2001)
- ‘Williams’ Argument Against External Reasons’, in: *Noûs* 30:2 (1996), 197-220
- Monk, R., *Ludwig Wittgenstein: The Duty of Genius* (London: Jonathan Cape, 1990)
- Moore, A. W., *Points of View* (Oxford: Oxford University Press, 1997)
- Moore, G. E., *Philosophical Papers* (London: George Allen and Unwin, 1959)
- *Principia Ethica* (Cambridge: Cambridge University Press, 1903)
- *Some Main Problems of Philosophy* (London: George Allen and Unwin, 1953)
- Nagel, T., *Mortal Questions* (Cambridge: Cambridge University Press, 1979)
- *The Last Word* (New York: Oxford University Press, 1997)
- *The Possibility of Altruism* (Oxford: Clarendon Press, 1970)
- *The View From Nowhere* (New York: Oxford University Press, 1986)
- Nietzsche, F., *Beyond Good and Evil*. Translated by R. J. Hollingdale (London: Penguin, 1990)
- ‘On the Genealogy of Morality’ and Other Writings. Edited by K. Ansell-Pearson (Cambridge: Cambridge University Press, 1994)
- Nozick, R., *Philosophical Explanations* (Oxford: Clarendon Press, 1981)

Bibliography

- Pears, D., *The False Prison: A Study of the Development of Wittgenstein's Philosophy* (2 vols., Oxford: Clarendon Press, 1987-8)
- Plato, *Republic*, in: *The Dialogues of Plato*, vol. II. Translated by B. Jowett (Fourth edition, Oxford: Clarendon Press, 1953)
- Platts, M., *Ways of Meaning: An Introduction to the Philosophy of Language* (London: Routledge, Kegan and Paul, 1979)
- Putnam, H., *Ethics Without Ontology* (Cambridge, MA: Harvard University Press, 2004)
- *Renewing Philosophy* (Cambridge, MA: Harvard University Press, 1992)
 - 'Was Wittgenstein Really An Anti-Realist about Mathematics?', in: T. G. McCarthy and S. C. Stidd (eds.), *Wittgenstein in America*, 140-94.
- Quine, W. V., *Pursuit of Truth* (Revised edition, Cambridge, MA: Harvard University Press, 1992)
- Raphael, D. D. (ed.), *British Moralists 1650–1800* (2 vols., Oxford: Clarendon Press, 1969)
- Rhees, R. (ed.), *Ludwig Wittgenstein: Personal Recollections* (Oxford: Basil Blackwell, 1981)
- 'Some Developments in Wittgenstein's View of Ethics', in: *The Philosophical Review* 74:1 (January 1965), 17-26
 - *Wittgenstein's On Certainty: There – Like Our Life*. Ed. D. Z. Phillips (Oxford: Blackwell, 2003)
- Rorty, R., *Contingency, Irony, and Solidarity* (Cambridge: Cambridge University Press, 1989)
- *Philosophy and the Mirror of Nature* (Princeton: Princeton University Press, 1979)
- Runge, P. O., *Farbenkugel: Konstruktion des Verhältnisses aller Mischungen der Farben zueinander und ihrer vollständigen Affinität* (Köln: Tropen-Verlag, 1999)
- Russell, B., *Logic and Knowledge: Essays 1901–1950*. Ed. R. C. Marsh (London: George Allen and Unwin, 1956)
- Sayre-McCord, G. (ed.), *Essays on Moral Realism* (Ithaca: Cornell University Press, 1988)
- Scanlon, T. M., *What We Owe To Each Other* (Cambridge, MA: Harvard University Press, 1998)
- Sellars, W., *Empiricism and the Philosophy of Mind* (Cambridge, MA: Harvard University Press, 1997)
- Sidgwick, H., *The Methods of Ethics* (Seventh edition, London: Macmillan, 1907)
- Sinnott-Armstrong, W. and Timmons, M. (eds.), *Moral Knowledge? New Readings in Moral Epistemology* (New York: Oxford University Press, 1996)
- Sluga, H. and Stern, D. G. (eds.), *The Cambridge Companion to Wittgenstein* (Cambridge: Cambridge University Press, 1996)
- Smith, M., *The Moral Problem* (Oxford: Blackwell, 1994)
- 'The Humean Theory of Motivation', in: *Mind* 96, No. 381 (January 1987), 36-61
- Stevenson, C. L., *Ethics and Language* (New Haven: Yale University Press, 1945)
- *Facts and Values: Studies in Ethical Analysis* (New Haven: Yale University Press, 1963)
- Strawson, P. F., *Analysis and Metaphysics: An Introduction to Philosophy* (Oxford: Oxford University Press, 1992)
- *Freedom and Resentment and Other Essays* (London: Methuen, 1974)
 - *Individuals: An Essay in Descriptive Metaphysics* (London: Methuen, 1959)
 - *Introduction to Logical Theory* (London: Methuen, 1952)
 - *Logico-Linguistic Papers* (London: Methuen, 1971)
 - *Skepticism and Naturalism: Some Varieties* (New York: Columbia University Press, 1985)
 - *The Bounds of Sense: An Essay on Kant's Critique of Pure Reason* (London: Methuen, 1966)
- Stroll, A., *Moore and Wittgenstein on Certainty* (New York: Oxford University Press, 1994)
- Stroud, B., *Hume* (London: Routledge and Kegan Paul, 1977)
- *Meaning, Understanding, and Practice: Philosophical Essays* (Oxford: Oxford University Press, 2000)

Bibliography

- *The Quest for Reality: Subjectivism and the Metaphysics of Colour* (New York: Oxford University Press, 2000)
- *Understanding Human Knowledge: Philosophical Essays* (Oxford: Oxford University Press, 2000)
- von Wright, G. H., *Wittgenstein* (Oxford: Basil Blackwell, 1982)
- Wallace, J., 'How to Argue About Practical Reason', in: *Mind* 99, No. 395 (July 1990), 355-85
 - *Responsibility and the Moral Sentiments* (Cambridge, MA: Harvard University Press, 1994)
- Watson, G., *Agency and Answerability: Selected Essays* (Oxford: Clarendon Press, 2004)
- Wiggins, D., *Ethics: Twelve Lectures on the Philosophy of Morality* (Cambridge, MA: Harvard University Press, 2006)
 - *Needs, Values, Truth* (Third edition, Oxford: Clarendon Press, 1998)
- Williams, B., 'A Critique of Utilitarianism', in: J. J. C. Smart and B. Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973), 75-150
 - *Descartes: The Project of Pure Enquiry* (Harmondsworth: Penguin, 1978)
 - *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press, 1985)
 - *In the Beginning Was the Deed: Realism and Moralism in Political Argument*. Ed. G. Hawthorn (Princeton: Princeton University Press, 2005)
 - *Making Sense of Humanity and Other Philosophical Papers 1982-1993* (Cambridge: Cambridge University Press, 1995)
 - *Morality: An Introduction to Ethics* (Cambridge: Cambridge University Press, 1976)
 - *Moral Luck: Philosophical Papers 1973-1980* (Cambridge: Cambridge University Press, 1981)
 - *Philosophy as a Humanistic Discipline*. Ed. A. W. Moore (Princeton: Princeton University Press, 2006)
 - 'Postscript: Some Further Notes on Internal and External Reasons', in: Millgram, E. (ed.), *Varieties of Practical Reasoning*, 91-97
 - *Problems of the Self: Philosophical Papers 1956-1972* (Cambridge: Cambridge University Press, 1973)
 - *Shame and Necessity* (Berkeley: The University of California Press, 1993)
 - *Truth and Truthfulness: An Essay in Genealogy* (Princeton: Princeton University Press, 2002)
- Wittgenstein, L., 'A Lecture on Ethics', in: *The Philosophical Review* 74:1 (January 1965), 3-12
 - *Culture and Value*. Ed. G. H. von Wright (Revised second edition, Oxford: Blackwell, 1998)
 - *Lectures and Conversations on Aesthetics, Psychology and Religious Belief*. Compiled from Notes taken by Y. Smythies, R. Rhees and J. Taylor. Ed. C. Barrett (Oxford: Basil Blackwell, 1966)
 - *Notebooks 1914-1916*. Ed. G. H. von Wright and G. E. M. Anscombe (Second edition, Oxford: Basil Blackwell, 1979)
 - *On Certainty*. Ed. G. E. M. Anscombe and G. H. von Wright (Oxford, Basil Blackwell, 1969)
 - *Philosophical Grammar*. Ed. R. Rhees (Oxford: Basil Blackwell, 1974)
 - *Philosophical Investigations*. Translated by G. E. M. Anscombe (Third edition, Oxford: Blackwell, 2001)
 - *Philosophical Occasions 1912-1951*. Ed. J. Klagge and A. Nordmann (Indianapolis: Hackett, 1993)
 - *Philosophical Remarks*. Ed. R. Rhees (Oxford: Basil Blackwell, 1975)
 - *Preliminary Studies for the "Philosophical Investigations" Generally Known as The Blue and Brown Books* (Oxford: Basil Blackwell, 1958)
 - *Remarks on Colour*. Ed. G. E. M. Anscombe (Oxford: Basil Blackwell, 1977)

Bibliography

- *Remarks on the Foundations of Mathematics*. Ed. G. H. von Wright, R. Rhees, and G. E. M. Anscombe (Third edition, Oxford: Basil Blackwell, 1978)
 - *Tractatus Logico-Philosophicus*. Translated by D. F. Pears and B. F. McGuinness (London: Routledge and Kegan Paul, 1961)
 - *Wittgenstein's Lectures Cambridge 1932–1935*. From the Notes of A. Ambrose and M. McDonald. Ed. A. Ambrose (Oxford: Basil Blackwell, 1979)
 - *Wittgenstein's Lectures on the Foundations of Mathematics, Cambridge, 1939*. From the Notes of R. G. Bosanquet, N. Malcolm, R. Rhees, and Y. Smythies. Ed. C. Diamond (Chicago: University of Chicago Press, 1989)
 - *Zettel*. Ed. G. E. M. Anscombe and G. H. von Wright (Oxford: Basil Blackwell, 1967)
- Wolf, S., *Freedom Within Reason* (New York: Oxford University Press, 1990)
- Wright, C., *Truth and Objectivity* (Cambridge, MA: Harvard University Press, 1992)

